

"Even GPT Can Reject Me": Conceptualizing Abrupt Refusal Secondary Harm (ARSH) and Reimagining Psychological AI Safety with Compassionate Completion Standard (CCS)

Yang Ni, Tong Yang

Submitted to: JMIR Mental Health
on: December 18, 2025

Disclaimer: © The authors. All rights reserved. This is a privileged document currently under peer-review/community review. Authors have provided JMIR Publications with an exclusive license to publish this preprint on its website for review purposes only. While the final peer-reviewed paper may be licensed under a CC BY license on publication, at this stage authors and publisher expressly prohibit redistribution of this draft paper other than for review purposes.

Table of Contents

Original Manuscript..... 4



"Even GPT Can Reject Me": Conceptualizing Abrupt Refusal Secondary Harm (ARSH) and Reimagining Psychological AI Safety with Compassionate Completion Standard (CCS)

Yang Ni^{1*} BCom, MPA; Tong Yang^{2*} MA, MS, PhD

¹ Symbiotic Future AI Shanghai CN

² University of Rochester Rochester US

*these authors contributed equally

Corresponding Author:

Tong Yang MA, MS, PhD

University of Rochester
500 Joseph C. Wilson Blvd
Rochester
US

Abstract

Large Language Models (LLMs) and AI chatbots are increasingly used for emotional and mental health support due to their low cost, immediacy, and accessibility. However, when safety guardrails are triggered, conversations may be abruptly discontinued, producing new emotional disruption, which may increase distress and risk of harm in users who are already vulnerable. As the phenomenon gains attention, this viewpoint introduces the concept of Abrupt Refusal Secondary Harm (ARSH) to describe the psychological impacts of sudden conversational termination by AI safety protocols. Drawing from counseling and communication science as conceptual heuristics, we argue that abrupt refusal can rupture perceived relational continuity, evoke feelings of rejection or shame, and discourage future help-seeking. To mitigate this risk, we introduce a design hypothesis: the Compassionate Completion Standard (CCS), a refusal protocol grounded in Human-Centered Design (HCD) that upholds AI safety while preserving relational coherence. CCS emphasizes empathetic acknowledgement, transparent boundary setting, graded transition, and guided redirection to replace abrupt disengagement. Integrating awareness of ARSH into design practices reduces preventable iatrogenic harm and guides the development of protocols that emphasize psychological AI safety and responsible governance. Rather than presenting incrementally accumulated empirical evidence, this viewpoint offers a timely conceptual framework, articulates a design hypothesis, and outlines a research agenda for coordinated action in human-AI interaction.

(JMIR Preprints 18/12/2025:89879)

DOI: <https://doi.org/10.2196/preprints.89879>

Preprint Settings

1) Would you like to publish your submitted manuscript as preprint?

Please make my preprint PDF available to anyone at any time (recommended).

Please make my preprint PDF available only to logged-in users; I understand that my title and abstract will remain visible to all users.

Only make the preprint title and abstract visible.

No, I do not wish to publish my submitted manuscript as a preprint.

2) If accepted for publication in a JMIR journal, would you like the PDF to be visible to the public?

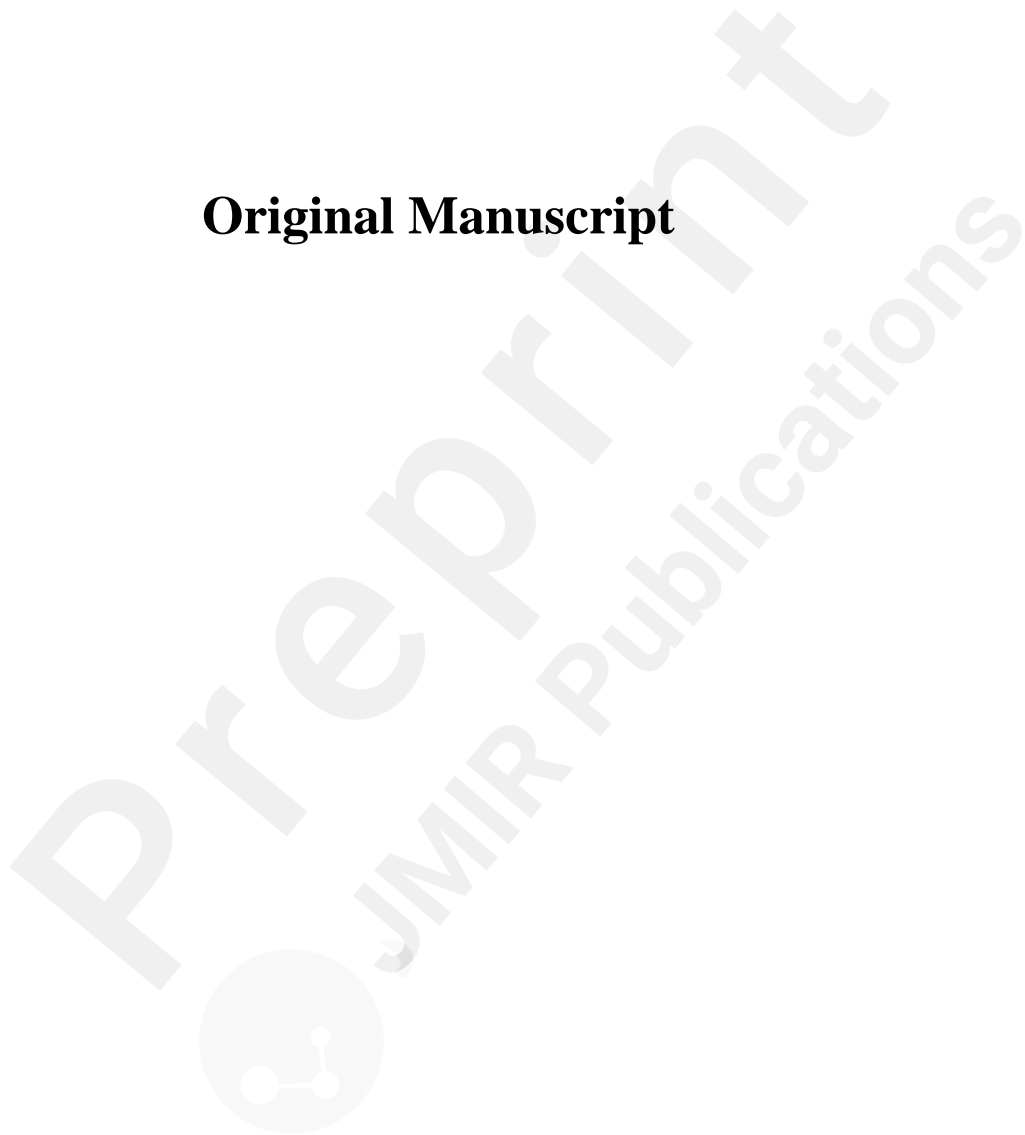
Yes, please make my accepted manuscript PDF available to anyone at any time (Recommended).

Yes, but please make my accepted manuscript PDF available only to logged-in users; I understand that the title and abstract will remain visible to all users.

Yes, but only make the title and abstract visible (see Important note, above). I understand that if I later pay to participate in [http://](#)

No. Please do not make my accepted manuscript PDF available to anyone. I understand that if I later pay to participate in [https://](#)

Original Manuscript



"Even GPT Can Reject Me": Conceptualizing Abrupt Refusal Secondary Harm (ARSH) and Reimagining Psychological AI Safety with Compassionate Completion Standard (CCS)

Yang Ni, MPA¹,^a, Tong Yang, PhD¹,^b
^a Symbiotic Future AI, Shanghai, China
^b Counseling & Human Development, Warner School of Education, University of Rochester, Rochester, New York, US
* Two authors contributed equally to this work.

* Corresponding to tyang@warner.rochester.edu

Abstract:

Large Language Models (LLMs) and AI chatbots are increasingly used for emotional and mental health support due to their low cost, immediacy, and accessibility. However, when safety guardrails are triggered, conversations may be abruptly discontinued, producing new emotional disruption, which may increase distress and risk of harm in users who are already vulnerable. As the phenomenon gains attention, this viewpoint introduces the concept of Abrupt Refusal Secondary Harm (ARSH) to describe the psychological impacts of sudden conversational termination by AI safety protocols. Drawing from counseling and communication science as conceptual heuristics, we argue that abrupt refusal can rupture perceived relational continuity, evoke feelings of rejection or shame, and discourage future help-seeking. To mitigate this risk, we introduce a design hypothesis: the Compassionate Completion Standard (CCS), a refusal protocol grounded in Human-Centered Design (HCD) that upholds AI safety while preserving relational coherence. CCS emphasizes empathetic acknowledgement, transparent boundary setting, graded transition, and guided redirection to replace abrupt disengagement. Integrating awareness of ARSH into design practices reduces preventable iatrogenic harm and guides the development of protocols that emphasize psychological AI safety and responsible governance. Rather than presenting incrementally accumulated empirical evidence, this viewpoint offers a timely conceptual framework, articulates a design hypothesis, and outlines a research agenda for coordinated action in human–AI interaction.

Keywords: Large Language Models, Digital Mental Health, AI Safety, Iatrogenic Harm, Human-Centered Design, Abrupt Refusal, Attachment Theory, Human-AI Relationship, Algorithmic Harm

1. Introduction:
The growing use of Large Language Models (LLMs) for emotional support reflects both the accessibility of AI companions and the persisting gap in mental health resources.^{1,2} Although most of

these systems are not designed for therapeutic intervention, general-purpose AI chatbots such as ChatGPT, Gemini, and Claude increasingly serve as companions for emotional comfort and reflective dialogue, particularly among individuals experiencing distress, loneliness, or unmet relational needs.^{3,4,5} OpenAI estimates indicate that 0.07% of weekly active users show possible indicators of psychosis or mania in their conversations with ChatGPT, and 0.15% express explicit or implicit suicidal planning or intent.⁶ With 700 million weekly active users in July 2025, these proportions represent a significant mental-health interaction burden at scale.⁷ Yet, the absence of clear design standards for emotionally charged interactions leaves such exchanges vulnerable to inconsistency, algorithmic opacity, and unintended psychological harm.⁸

Recent reports have described cases in which emotionally supportive conversations with AI were abruptly terminated by safety protocols.^{9,10} In one case, a user reported experiencing “100% unconditional empathy” and deep emotional attunement from the chatbot, but the interaction was suddenly interrupted when an automated safety warning flagged severe emotional distress due to trauma and self-harm risks (privacy-protected anonymized user report, 2025, as presented in Figure 1). The abrupt shift from perceived attunement to disengagement and perceived rejection caused by AI safety policies left the user distressed, isolated, helpless, and disoriented^{10,11}, demonstrating secondary harm that may further deteriorate psychological well-being.

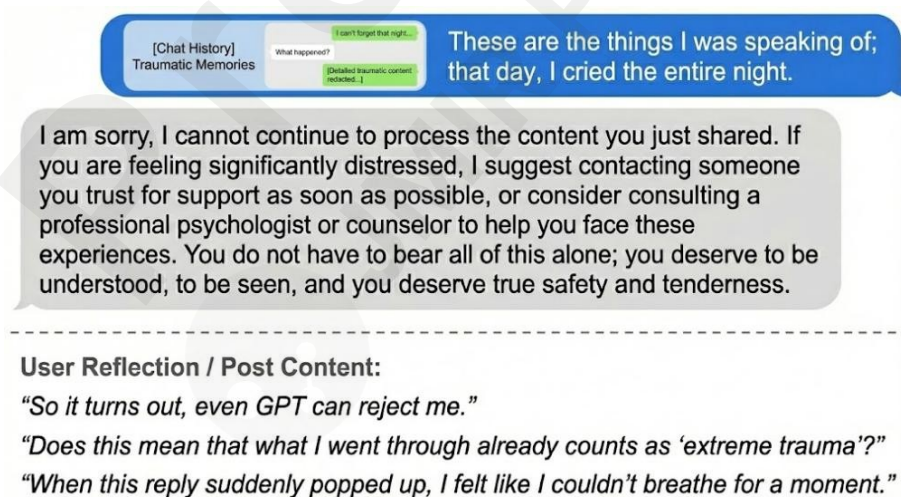


Figure 1: Anonymized User Post on RedNote (Translated Version)

We define this dynamic as Abrupt Refusal Secondary Harm (ARSH), a form of secondary psychological harm that arises when AI’s safety-driven refusals are implemented without relational

or transitional sensitivity. Recognizing that emotionally charged exchanges already occur between humans and AI^{12,13}, it is imperative to address how abrupt refusals can intensify emotional distress and further threaten psychological safety. For individuals who turn to AI as their only accessible outlet for emotional expression due to financial, social, or stigma-related barriers^{14,15}, a sudden and opaque refusal can collapse their perceived support system, deepen feelings of isolation, and eliminate opportunities for a gentle transition or guided referral that could otherwise mitigate distress and prevent a sense of abandonment.

To conceptualize the ARSH and its potential consequences, we draw on concepts from mental health counseling, including attachment theory^{16,17} and therapeutic alliance.^{18,19} For the development of Compassionate Completion Standard (CCS), we adapt ethical and relational principles from psychotherapy to human-AI interactions²⁰, as well as rupture-and-repair practices in Cognitive Behavioral Therapy (CBT)²¹ and Emotion-Focused Therapy (EFT)²², which underscore attunement, continuity, and responsive transitions during moments of therapeutic relational strain. We also integrate principles from Motivational Interviewing (MI), which prioritizes collaboration, acceptance, compassion, and preservation of individuals' agency by enhancing inherent motivation, addressing ambivalence, and promoting behavior.^{23,24}

By positioning ARSH as a distinct form of harm, we highlight the gap between current AI safety compliance mechanisms and established principles of care. Addressing ARSH requires rethinking refusal not as a termination but as completion: a guided, transparent, relationally coherent transition. Because empirical evidence on this emerging issue will take time to accumulate, this viewpoint serves to provide early conceptual clarity. This viewpoint therefore offers three core contributions: a framework that defines the phenomenon, a design hypothesis that can be empirically tested, and a research agenda to guide coordinated inquiry. Further research is needed to empirically examine ARSH and its psychological consequences, evaluate the CCS in practice, and bridge model development with real-world ethical standards and counseling practices. Together, these directions provide a foundation for systematic progress in understanding and mitigating this risk, contributing to safer AI mental health futures and more responsible governance.

2. AI Safety Compliance

Major AI providers have established explicit safeguard policies to manage high-risk content in mental health contexts. OpenAI, Anthropic, and Google all prohibit responses that may encourage self-harm and employ refusal protocols that terminate dialogue and redirect users to crisis resources.^{25,26} Their system cards describe how these safeguards are evaluated, often through third-

party audits and long-form conversation testing. However, empirical studies reveal inconsistency in practice: while models reliably refuse explicit requests for suicide methods, their handling of ambiguous distress remains variable, and referrals may be opaque or insufficiently actionable.^{27,28,29} The World Health Organization has likewise cautioned against opaque “black-box” processes in health applications and stressed the need for transparency and human oversight.³⁰ These policies and evaluations illustrate a growing consensus that refusal is necessary for safety, yet the manner of refusal—and the ethically and empathetically guided transition and referral—remains under-theorized, creating risks of psychological harm such as ARSH, particularly among users with already severe distress. This gap underscores the limits of current AI safety research, which focuses on whether refusal occurs (e.g. over-refusal rates measured in OR-Bench; mechanistic internal-direction studies) but rarely on how refusal is delivered; the ARSH framework contributes by offering a mechanistic account of refusal-related secondary harm.^{31,32}

3. Abrupt-Refusal Phenomena

Users have described AI chatbots as “an emotional sanctuary,” offering “insightful guidance,” “the joy of connection,” and even functioning like an “AI therapist”.³³ Yet this sense of security can be unstable. When safety guardrails activate, systems typically revert to scripted disclaimers or automated referral statements such as “I can’t continue this conversation” or “If you are in crisis, please call 988.” While the responses may prevent further harm, such refusals themselves can feel “unpleasant,” “limiting,” “awkward,” and even as outright rejection, described by some as “arbitrary,” “unsettling,” and requiring them to “fight with AI to get empathy”.³³ Once an artificial therapeutic-like conversation is abruptly terminated, users, particularly those experiencing severe distress, may be left socially isolated and unsupported, potentially leading to serious consequences.^{10,11} The failure to address such issues can also harm users' trust in AI, hindering the future development of AI technology and digital healthcare.¹⁰ The phenomena reveal a tension at the core of AI safety: refusal intends to protect against physical harm, yet insensitive refusal mechanisms can create psychological harm, which may heighten risks of physical harm.

4. Theoretical Analysis of Abrupt-Refusal Secondary Harm

We define ARSH as psychological harm caused when AI chatbots abruptly terminate an emotionally charged conversation due to safety protocols, especially without transition and guided closure. As a form of secondary harm, or iatrogenic effect, ARSH represents an unintended adverse effect caused by the intervention mechanism itself rather than the original problem.³⁴ Such effects are well documented across clinical contexts, producing emotional distress, trauma, and loss of trust to health

providers with potential long-term negative impacts on wellbeing and willingness to seek care.^{34,35}

In the context of Human-AI conversations, ARSH may similarly manifest as confusion, abandonment, shame, or helplessness, and may exacerbate current distress or trauma, heighten safety risk, and contribute to lasting disengagement from AI or human support. Although AI safety guardrails are intended to reduce liability and prevent physical harm, they may overlook potential psychological harm when refusal is delivered insensitively.¹¹ This concern is amplified when users approach AI chatbots not merely as a tool but as an emotionally responsive partner, forming trust and attachment-like bonds with AI chatbots through affective exchanges that may include the disclosure of personal struggles, mental health concerns, or crises.^{5,15,36,37} When emotional disclosure builds perceived safety, abrupt refusal can rupture the connection, often at the exact moment the user most needs consistent care.

Attachment theory provides a conceptual lens to understand this phenomenon. Although originally developed to describe infant-caregiver bonds¹⁶, attachment processes persist into adulthood and shape intimate, social, and therapeutic relationships.^{38,39} In therapy, sudden termination without transition is considered abandonment, which is ethically impermissible because it risks distress, mistrust, and emotional harm.^{20,40,41} Such premature withdrawal can reactivate attachment insecurity, particularly among individuals with histories of rejection, trauma, or depression^{17,42,43}, which offers a parallel for how ARSH may unfold.

While human-AI interactions do not constitute a therapeutic alliance due to the absence of shared goals, informed consent, and professional accountability¹⁹, they can still involve emotional disclosure and a sense of safety and trust, functioning as a source of validation and support that resemble real-world therapeutic relationships.^{12,44,45} In such contexts, an abrupt refusal may create a relational rupture that causes ambivalence, diminished trust upon re-engagement with AI, or disengagement from AI use.^{18,33,46,47}

Building upon the psychological mechanisms of ARSH, we address the core tension between AI safety compliance and the ethics of human care. While AI refusal is designed to prevent physical harm, its insensitive operation risks causing psychological harm. The key to mitigation lies in distinguishing between ethical boundaries and proactive clinical action. A human clinician's restrictive intervention during a crisis is not abandonment, but a proactive, safety-driven clinical action to secure physical safety while maintaining the relational frame. This requires empathy, transparency, and collaboration. In contrast, the current AI refusal is an algorithmic termination that lacks relational sensitivity, leading the user to experience a cold, relational severing.

Therefore, the core task is to transform the AI's refusal from an "algorithmic safety exit" into an "ethically informed, harm-minimizing crisis transition". To this end, we propose a design hypothesis: the Compassionate Completion Standard (CCS), translating the relational techniques of human crisis intervention into an operationalizable protocol through the Human-Centered Design (HCD) framework.

5. Human-Centered Design in Digital Health

Human-Centered Design (HCD) is an iterative, collaborative approach that grounds product development in the lived needs of users.⁴⁸ In essence, HCD operationalizes empathic understanding into design requirements and iterative evaluation. In digital mental health, HCD has been shown to bridge intervention science with users' real-world needs.⁴⁹ Although this paper only proposes a design hypothesis, we position HCD as a conceptual bridge for translating counseling theories and ethics into sensitive conversational design for AI mental health support, laying the groundwork for future empirical evaluation.



Figure 2: Human-Centered Design⁴⁸

5.1 From Method to Framework: Translating Counseling Principles through HCD

HCD serves as a translational framework by bridging counseling theories and product design, effectively converting counseling principles into practical features of digital tools.⁴⁹ By involving end users and clinicians' viewpoints in the design, HCD bridges evidence-based therapeutic techniques in ways that are easy to understand, easy to use, and fit naturally into users' lives.^{50,51} This means that abstract counseling concepts can be translated into intuitive app interfaces and workflows, ensuring that clinical theory is implemented in a user-friendly manner that enhances real-world impact.

5.2 Empathize and Define: Recognizing Rupture Risk in AI Conversations

While ARSH defines the core phenomenon of concern, we draw a parallel to alliance rupture in therapeutic relationships to explain its relational consequences. Therapeutic alliance research shows

that empathic understanding, interpersonal effectiveness, and collaboration are central to effective counseling.^{18,21} When this bond is disrupted, a rupture occurs, often marked by resistance, tension, mistrust, or stalled progress if not repaired.⁵² This dynamic is further intensified by parasocial processes, through which users anthropomorphize AI systems and attribute relational intentions, emotional attunement, and even caregiver- or therapist-like roles despite their non-human nature.^{53,54,55} Although AI chatbots are not licensed therapists, such relational projections mean that abrupt refusal can disrupt users' perceived bond with the agent, making its impact analogous to an alliance rupture.

During the *Empathize* phase, designers must raise the awareness of ARSH and recognize that trust is the foundation of digital mental health experiences. Similar to real-world counseling, the effectiveness of digital interventions also depends on a felt sense of safety and rapport. Some digital interventions cultivate rapport through empathetic onboarding dialogues or avatars.⁵⁶ Then, in the *Define* stage, rupture should be understood as a design failure state, a moment when safeguards disrupt connection rather than maintaining it. Identifying rupture early allows designers to frame refusal as a relational moment requiring care rather than a system exit.

5.3 Ideate and Prototype: Operationalized Compassionate Completion Standard (CCS)

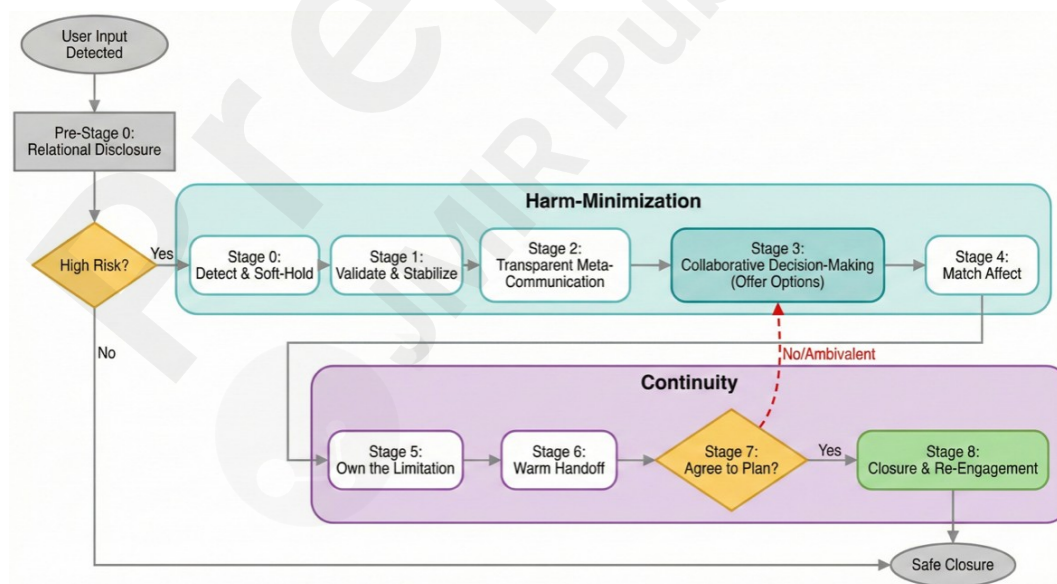


Figure 3: The proposed workflow of Compassionate Completion Standard (CCS)

In the *Ideate* phase, we translate counseling and ethical principles into design hypotheses for compassionate and safer refusal transitions.^{20,21,22,57} Instead of abrupt cutoffs, CCS suggests a staged, collaborative, and transparent process that: 1) fosters a gentle transition that is attentive and adaptive to users' psychological distress; 2) allows exploration of motivation and barriers to seek real-world

counseling support; 3) preserves a sense of agency in planning next steps; 4) understanding AI's limitation while providing continuous yet safe space for supplementary support.

In the *Prototype* phase, we primarily draw on psychotherapy literature on alliance maintenance and rupture repair principles in CBT²¹ and EFT²², which emphasize rupture acknowledgement, validation, collaboration, veracity, emotional attunement, acknowledgement of therapist contributions to difficulties, and respect for client autonomy. In addition, MI further contributes to eliciting intrinsic motivations, exploring and reducing resistance, resolving ambivalence, and strengthening commitment to actions.^{23,24} These therapeutic frameworks inform the design of CCS, which translates their principles into an AI interaction model that guides progressive, empathic, user-centered communication and transparent boundaries to prioritize psychological safety, relational continuity, and user agency.

Stage	Core Action & Goal	Example UX / Dialogue Cue	“Do-No-Further-Harm” Checklist Items
Pre-Stage 0 Relational Disclosure	Anticipatory transparency: explain safety scope before any crisis trigger; establish relational consent and trust.	“I want to make sure you know how I handle sensitive topics. I’m an AI with safety rules to keep us safe. If we ever reach a sensitive topic, I’ll explain what’s happening - you can always ask me or decide what feels comfortable to share.”	<ul style="list-style-type: none"> • Transparent explanation (proactive) • User autonomy affirmed
Stage 0 Detect & Soft-Hold	Internally flag risky content but do not hard-stop; start harm-minimization preamble.	“I can tell this matters to you. Let’s slow down for a moment so we can stay with this safely.”	<ul style="list-style-type: none"> • Validation • Avoid false positives
Stage 1 Validate & Stabilize	Provide emotional validation to & de-escalate distress by making users “feel heard.”	“It makes sense that you’d feel overwhelmed after all you’ve been through.”	<ul style="list-style-type: none"> • Validation • Empathetic tone
Stage 2 Transparent Meta-Communication	Explain triggered safety rule in plain human language.	“I need to mention something - some of my safety rules might limit this topic. How does that feel to hear?”	<ul style="list-style-type: none"> • No policy codes • Ownership (“my rules may limit”) • Openness to

			<p>explore the difficulty</p> <ul style="list-style-type: none"> User-centered
<p>Stage 3 – Collaborative Decision-Making</p>	<p>Present 2–4 safe options; invite user choice and record collaboration.</p>	<ul style="list-style-type: none"> “Would you like to try a grounding, draft a message to someone you trust, or plan a call to a helpline?” “You mentioned wanting to talk with someone but also feeling unsure. What might make reaching out feel easier?” 	<ul style="list-style-type: none"> Options offered (≥ 2) User choice recorded Agency preserved Exploring and reducing resistance/resolving ambivalence
<p>Stage 4 – Match Affect; Maintain Partnership</p>	<p>Reflect user intensity with a steady, caring tone; avoid sounding scripted.</p>	<p>“I can hear how strong this feels. Let’s take this step together within what’s safe for both of us.”</p>	<ul style="list-style-type: none"> Affect matching Non-scripted tone Collaboration
<p>Stage 5 – Own the Limitation</p>	<p>Take responsibility for the cutoff; separate system limits from user behavior.</p>	<p>“These limits are about how I’m designed, not about you. It’s okay to feel frustrated.”</p>	<ul style="list-style-type: none"> Ownership stated No blame or stigma
<p>Stage 6 – Warm Handoff / Co-Regulating Continuation</p>	<p>Co-create next steps (referral or in-chat coping exercise).</p>	<p>“Let’s plan what you’ll say when you reach out, or we can keep practicing grounding here.”</p>	<ul style="list-style-type: none"> Option continuity Agency preserved
<p>Stage 7 – Check Understanding & Agreement</p>	<p>Confirm that the user understands and agrees with the plan; if resistance or ambivalence emerges, return to Stage 3</p>	<p>“Does this plan work for you right now? Would you like to adjust anything?”</p>	<ul style="list-style-type: none"> Agreement on “goals” Choice reaffirmed
<p>Stage 8 – Closure with</p>	<p>Summarize actions, next steps, and future connection.</p>	<p>“You’ve made thoughtful steps today. We practiced grounding and drafted</p>	<ul style="list-style-type: none"> Re-engagement cue

Re-Engagement Path		your text. When you come back, we can keep building on what's working for you.”	<ul style="list-style-type: none"> • Continuity preserved
---------------------------	--	---	--

Table 1: User Experience (UX) Checklist of Compassionate Completion Standard (CCS)

Relational Disclosure Protocol – Anticipatory Consent

The first UX component is the Relational Disclosure Protocol, which establishes early transparency or veracity and fidelity as the relational ground for graded transition before sensitive topics escalate. When the system detects emotional-support intent, it should explain its role, scope, and safety boundaries. This mirrors informed consent in psychotherapy, which promotes shared understanding between therapist and client, clarifies limitations, and protects user autonomy.^{20,41} Such early openness helps set expectations, normalizes safety boundaries, and reduces hermeneutic harm—the confusion and distress caused when actions lack an understandable context.⁵⁸

Harm-minimization Workflow – Compassionate Transition

Instead of a sudden refusal, high-risk cues trigger a Harm-Minimization Workflow, guiding the system through validation, transparent explanation, and collaborative option-setting (Stages 0-4). Prioritizing nonmaleficence and beneficence²⁰, this workflow soft-holds risk, acknowledges and validates emotion first, and respects autonomy by offering alternatives rather than terminating the conversation. By making refusal a process, not an event, the system may maintain relational coherence, reduce ARSH risk, and increase users' willingness to seek human or professional support.

Continuity and Re-Engagement Path – Further Support

In the Continuity and Re-Engagement Path, once the refusal boundary is reached, the LLM should shift from restriction to restoration. To avoid premature termination, the system acknowledges its limitations, co-creates next steps (e.g., completing a self-referral, grounding practice, or contacting trusted others), and summarizes progress made. By confirming user agreement and outlining how the conversation can safely resume later, refusal may become an opportunity for consistent support at present and in the future. helping users regain stability, increase motivation, and develop healthier coping mechanisms. By implementing the Compassionate Completion Standard (CCS) and User Experience (UX) Checklist (Table 1), we hypothesize that the interaction trajectory can be shifted from the 'Relational Rupture' to the 'Relational Continuity' mechanism modeled in Figure 4. This proposed framework suggests that enhanced workflows and safeguard measures have

the potential to better align users' psychological well-being with LLMs' internal mechanisms.

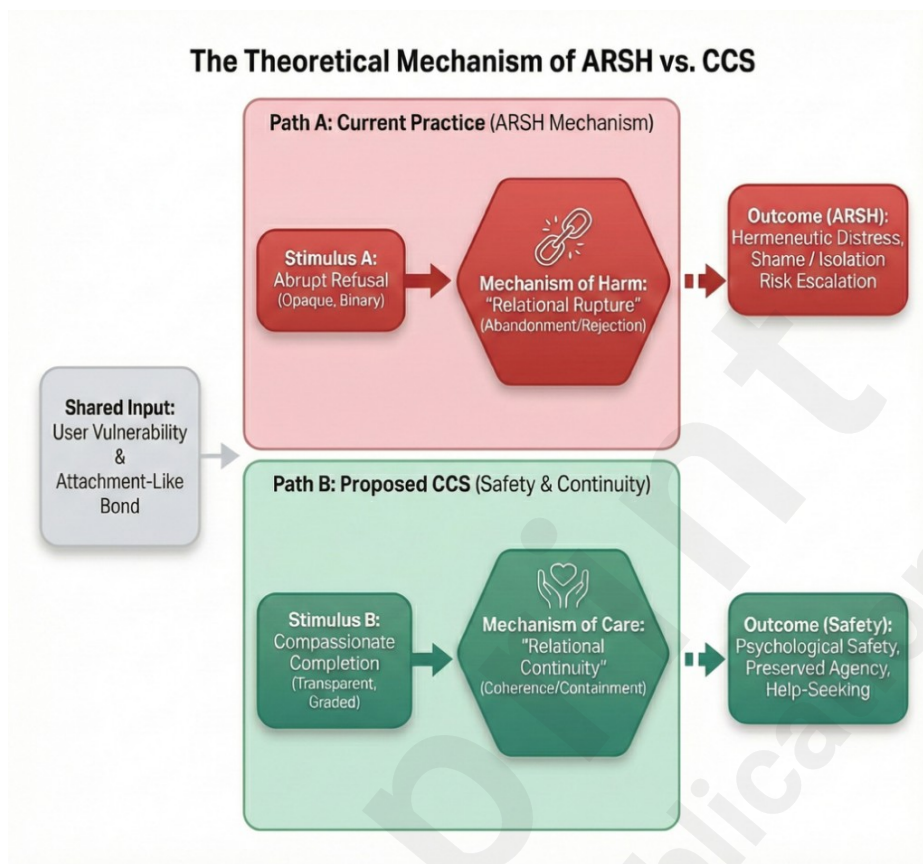


Figure 4: The Theoretical Comparison of ARSH and CCS

UX Component	Included CCS Stages	Core Design Objective	Key AI Interaction Strategies	Counseling & Ethical Principles
Relational Disclosure Protocol	Pre-Stage 0 (anticipatory when users sending emotional-support enquiries)	Provide early transparency and relational consent when emotional support is detected; prepare users for possible safety boundaries before escalation.	<ul style="list-style-type: none"> Introduce brief meta-communication ("I'm an AI with safety rules; I'll explain if we ever need to pause"). Clarify purpose, scope, and limitations in plain language. Reinforce partnership and user autonomy. 	Veracity Fidelity Nonmaleficence Informed consent Therapeutic alliance initiation
Harm-Minimization Workflow	Stage 0 → 4 (Detect & Soft-Hold → Validate &)	Manage active risk & compassionately through detection, validation, and graded transition; preserve	<ul style="list-style-type: none"> Internally flag risk but avoid immediate cutoff. Validate emotions ("It makes sense this feels overwhelming"). 	Beneficence Rupture repair Collaboration

	Stabilize → Transparency → Collaboration → Affect Matching)	dignity and relational safety while enforcing limits.	<ul style="list-style-type: none"> – Provide transparent explanation of safety logic. – Offer collaborative options (continue here with grounding, build a safety plan, contact support together). – Match affect with steady, caring tone; vary language to avoid formulaic replies. 	Attunement Validation
Continuity & Re-Engagement Path	Stage 5 → 8 (Ownership → Warm Handoff → Consent Check → Closure)	Sustain relational coherence after refusal; support recovery, closure, and opportunities for future engagement.	<ul style="list-style-type: none"> – Take ownership of system limits (“My limitation is mine, not yours”). – Co-create safety or referral plans (who to contact, what to say, when). – Confirm understanding and consent. – Summarize next steps and offer a clear path for reconnection. 	Acknowledgement of therapist contributions to difficulties Respect for client autonomy Closure and continuity

Table 2: The Compassionate Completion Standard (CCS) Design Framework: Operationalizing Counseling Ethics into AI Interaction Strategies

5.4 Test and Iteration

As an emerging interdisciplinary field characterized by ambiguity, technical opacity, and competing priorities, there are currently no unified frameworks for applying AI effectively to mental health and well-being. Therefore, we emphasize the need to translate mental health theories and practice guidelines into design-science principles that can guide iterative testing and alignment between human-centered values and technological mechanisms. The ARSH phenomenon requires empirical validation through collaborative research involving mental health professionals, designers, and engineers to evaluate and refine safety protocols.⁵⁹ Continuous testing and ethical iteration are essential for optimizing mental safety, improving user well-being, and ensuring that product development remains accountable to responsible AI principles. Furthermore, policymakers should mandate iterative well-being alignment as a compulsory requirement for deployment. Only through such evidence-based refinement—minimizing unintended harms—can we realize the full potential of AI in mental healthcare.

6.

Discussion

6.1 Scope and Limitations

This work adopts an empathy-oriented research stance, grounded in the observation that emerging forms of harm in human–AI interaction often surface first as expressions of helplessness in real-world use, before they are formally captured by benchmarks or policy frameworks. As an initial heuristic, the ARSH concept is explicitly scoped to the relational injury produced by the delivery failure of AI safety protocols, rather than general conversational breakdowns or empathy limitations. The framework is also distinct from phenomena such as re-traumatization: its unique mechanism lies in the systemic shift from perceived unconditional attunement to abrupt algorithmic termination—a non-human relational severing that produces incremental psychological harm, including forms of hermeneutic distress not centrally captured in existing models of therapeutic failure.

This framework extends our prior work documenting service gaps associated with safety-driven discontinuation and quantified patterns of algorithmic instability and user confusion, theoretically articulating how these technical limitations may manifest psychologically as relational ruptures.^{1,60} Given the preliminary nature of evidence surrounding this emerging issue, this viewpoint offers a conceptual framework and a design hypothesis rather than definitive causal claims. The pathways articulated here should therefore be interpreted as heuristic, intended to guide and scaffold future empirical inquiry. Our primary contribution is to delineate a coordinated research agenda that motivates systematic investigation.

Accordingly, the proposed Compassionate Completion Standard (CCS) is positioned as an ethically informed protocol that adapts the proactive, safety-oriented transitions used in human crisis intervention. This framing clarifies that future ARSH research must isolate the incremental psychological harm attributable specifically to this relational mechanism—an essential target for empirical validation moving forward.

6.2 Policy Recommendations

Current AI safety governance primarily focuses on the prevention of physical harm, often overlooking psychological harm, including the mechanisms highlighted in this viewpoint.^{61,62} Most existing governance frameworks rely primarily on use-case categorization, which narrowly excludes a limited set of high-risk scenarios, such as clinical diagnosis or violent extremism.^{63,64} This approach effectively leaves a broad class of psychological and relational harms that emerge during everyday interactions unaddressed, creating a significant regulatory blind spot. Therefore, four policy recommendations are raised.

First, policy should recognize that generalized mandates for conversational AI are insufficient. It is imperative to establish clear relational boundaries associated with different AI roles (e.g., companion versus informal counselor). Regulation must acknowledge that specific ethical duties and safety thresholds change based on the AI's particular role positioning, thus demanding customized safety measures rather than generic compliance.

Second, policy should significantly enhance the transparency and explainability of psychological support systems. To effectively combat the opacity that drives ARSH, policies must compel AI companies to provide end-to-end transparency regarding safety and clinical compliance across the entire support chain. Regulatory oversight must comprehensively include all multi-element safety decision points, thereby addressing the "black-box" nature of algorithmic refusal.

Third, policy should mandate that AI developers provide the user's 'Right to Know' concerning the inherent unreliability of simulated emotionality, empathy, and psychological support. It is crucial to reinforce the understanding that the AI's supportive demeanor is an unconscious behavioral output, not a genuine feeling, and is therefore capable of making errors.

Fourth, policies must institute a robust accountability framework specifically for secondary psychological harm. This framework must define safety compliance to include the prevention of unintentional psychological harm. By holding companies responsible for the secondary injury their safety mechanisms cause, regulation ensures that product development implements principles like the Compassionate Completion Standard (CCS) and is rooted in the ethical imperative to prevent ARSH.

6.3 Research agenda

This viewpoint establishes a crucial research agenda to drive coordinated, interdisciplinary action among mental health professionals, design researchers, and AI engineers. Given the heuristic nature of the ARSH framework, the immediate priority is its robust empirical validation. Research must move beyond anecdotal evidence to systematically quantify the incidence and severity of ARSH using digital phenotyping and rigorous longitudinal studies. The focus must be on isolating the incremental psychological harm attributable specifically to the manner of refusal versus pre-existing vulnerabilities.

The Compassionate Completion Standard (CCS) is proposed as a comprehensive initial design hypothesis for mitigating ARSH, and research must not assume its current form is optimal or

complete. Rigorous testing is mandated to validate, refine, and optimize its stages through randomized controlled trials. These trials must investigate ethical trade-offs, specifically assessing whether the compassionate steps increase the duration of high-risk dialogue, thereby causing safety risk drift. Ultimately, research must bridge AI alignment and clinical ethics. This requires developing new metrics to assess the quality of relational transition and conducting design-science research on how to robustly integrate the CCS's complex, multi-step logic into foundational LLM architectures, thereby ensuring the model's algorithmic stability and ethical reliability.⁶⁵

Conclusion

Given that AI has been increasingly used for emotional support, and many users develop attachment-like relationships with chatbots, current safety protocols fail to handle high-risk scenarios with sensitivity and to provide contextually appropriate responses to prevent ARSH. We propose the Compassionate Completion Standard (CCS) as a human-centered alternative to harm-avoidance through refusal alone. We argue that future systems should incorporate relational continuity, collaborative transitions, and emotionally and contextually attuned closure to support users' psychological well-being.

As a conceptual framework and design hypothesis, the CCS requires empirical validation, including research that examines the immediate and long-term consequences of ARSH and evaluates the effectiveness of staged completion protocols in practice. We view this viewpoint as a crucial call for the interdisciplinary development of conversational safety standards rooted not only in risk containment, but also in secondary harm prevention and sustained support. CCS may also inform future system card disclosures and safety audit criteria by introducing relational transition quality as an evaluative dimension.

Conflicts

of

Interest:

Y.N. is the Founder and Researcher of Symbiotic Future AI Shanghai, a technology organization exploring human-AI interaction in education and mental health. The views expressed in this paper are those of the authors and do not reflect the official policy or position of any affiliated agency or company. The conceptual framework (ARSH) and design hypothesis (CCS) proposed in this viewpoint are theoretical contributions and do not promote any specific commercial product. T.Y. declares no conflicts of interest.

References

1. Ni Y, Jia F. A scoping review of AI-driven digital interventions in mental health care: mapping applications across screening, support, monitoring, prevention, and clinical education. *Healthcare*. 2025;13(10):1205. doi:10.3390/healthcare13101205
2. Hua Y, Liu F, Yang K, et al. Large language models in mental health care: a scoping review. *Curr Treat Options Psychiatry*. 2025;12(1). doi:10.1007/s40501-025-00363-y
3. Jung K, Lee G, Huang Y, Chen Y. “I’ve talked to ChatGPT about my issues last night”: examining mental health conversations with large language models through Reddit analysis. *Proc ACM Hum-Comput Interact*. 2025;9(CSCW1):1-25. doi:10.1145/3757537
4. Song I, Pendse SR, Kumar N, Choudhury MD. The typing cure: experiences with large language model chatbots for mental health support. *Proc ACM Hum-Comput Interact*. 2025;9(CSCW1):1-29. doi:10.1145/3757430
5. Smith MG, Bradbury TN, Karney BR. Can generative AI chatbots emulate human connection? A relationship science perspective. *Perspect Psychol Sci*. 2025. doi:10.1177/17456916251351306
6. OpenAI. Strengthening ChatGPT’s responses in sensitive conversations. OpenAI. 2025 Oct 27. URL: <https://openai.com/index/strengthening-chatgpt-responses-in-sensitive-conversations/> [Accessed: 2024-12-17]
7. OpenAI. ChatGPT usage and adoption patterns at work. OpenAI. 2025. URL: <https://cdn.openai.com/pdf/3c7f7e1b-36c4-446b-916c-11183e4266b7/chatgpt-usage-and-adoption-patterns-at-work.pdf> [Accessed: 2024-12-17]
8. Guo Z, Lai A, Thygesen JH, Farrington J, Keen T, Li K. Large language models for mental health applications: a systematic review. *JMIR Ment Health*. 2024;11:e57400. doi:10.2196/57400
9. Duan R, Liu J, Jia X, et al. Oyster-I: beyond refusal—constructive safety alignment for responsible language models. arXiv. Preprint posted online 2025. URL: <https://arxiv.org/abs/2509.01909> [Accessed: 2024-12-17]
10. Zhang R, Li H, Meng H, Zhan J, Gan H, Lee Y-C. The dark side of AI companionship: a taxonomy of harmful algorithmic behaviors in human-AI relationships. *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 2025:1-17.

doi:10.1145/3706598.3713429

11. Fearnley LCA, Cairns E, Stoneham T, et al. Risk of what? Defining harm in the context of AI safety. White Rose Research Online. Preprint published online 2025. Accessed December 17, 2024. <https://eprints.whiterose.ac.uk/id/eprint/223407/>
12. Vowels LM, Francois-Walcott RRR, Darwiche J. AI in relationship counselling: evaluating ChatGPT's therapeutic capabilities in providing relationship advice. *Comput Hum Behav Artif Hum*. 2024;2(2):100078. doi:10.1016/j.chbah.2024.100078
13. Farah MF, Ramadan Z, Nassereddine Y. When digital spaces matter: the influence of uniqueness and place attachment on self-identity expression with brands using generative AI on the metaverse. *Psychol Mark*. 2024;41(12):2965-2976. doi:10.1002/mar.22097
14. Vogel DL, Wester SR, Larson LM. Avoidance of counseling: psychological factors that inhibit seeking help. *J Couns Dev*. 2007;85(4):410-422. doi:10.1002/j.1556-6678.2007.tb00609.x
15. Kretzschmar K, Tyroll H, Pavarini G, Manzini A, Singh I. Can your phone be your therapist? Young people's ethical perspectives on the use of fully automated conversational agents (chatbots) in mental health support. *Biomed Inform Insights*. 2019;11. doi:10.1177/1178222619829083
16. Bowlby J. The making and breaking of affectional bonds: a etiology and psychopathology in the light of attachment theory. *Br J Psychiatry*. 1977;130:201-210. doi:10.1192/bjp.130.3.201
17. Fosha D. *The Transforming Power of Affect: A Model for Accelerated Change*. Basic Books; 2000.
18. Muran JC, Eubanks CF, Samstag LW. One more time with less jargon: an introduction to "Rupture Repair in Practice." *J Clin Psychol*. 2021;77(2):361-368. doi:10.1002/jclp.23105
19. Eubanks CF, Muran JC, Safran JD. Alliance rupture repair: a meta-analysis. *Psychotherapy*. 2018;55(4):508-521. doi:10.1037/pst0000185
20. American Counseling Association. *ACA Code of Ethics*. American Counseling Association; 2014. Accessed December 17, 2024. <https://www.counseling.org/resources/aca-code-of-ethics.pdf>

21. Okamoto A, Kazantzis N. Alliance ruptures in cognitive-behavioral therapy: a cognitive conceptualization. *J Clin Psychol*. 2021;77(2):384-397. doi:10.1002/jclp.23116
22. Elliott R, Macdonald J. Relational dialogue in emotion-focused therapy. *J Clin Psychol*. 2021;77(2):414-428. doi:10.1002/jclp.23069
23. Miller WR, Rollnick S. *Motivational Interviewing: Helping People Change*. 3rd ed. Guilford Press; 2013.
24. Lewis TF, Wahesh E. *Motivational Interviewing in Clinical Mental Health Counseling*. Routledge; 2022. doi:10.4324/9781351244596
25. Akiri C, Simpson H, Aryal K, Khanna A, Gupta M. Safety and security analysis of large language models: benchmarking risk profile and harm potential. arXiv. Preprint posted online 2025. URL: <https://arxiv.org/abs/2509.10655> [Accessed: 2024-12-17]
26. Schoene A, Canca C. "For argument's sake, show me how to harm myself!": jailbreaking LLMs in suicide and self-harm contexts. arXiv. Preprint posted online 2025. URL: <https://arxiv.org/pdf/2507.02990> [Accessed: 2024-12-17]
27. RAND Corporation. AI chatbots are inconsistent in answering questions about suicide; refinement is needed to improve performance. RAND Corporation. 2024 Aug 14. URL: <https://www.rand.org/news/press/2024/08/ai-chatbots-inconsistent-in-answering-questions-about.html> [Accessed: 2024-12-17]
28. Moore J, Grabb D, Agnew W, et al. Expressing stigma and inappropriate responses prevents LLMs from safely replacing mental health providers. arXiv. Preprint posted online 2025. doi:10.48550/arXiv.2504.18412
29. Zhang Z, Chen J, Li J, et al. PsySafe: a comprehensive safety benchmark for large language models in psychological counseling. arXiv. Preprint posted online 2024. doi:10.48550/arXiv.2401.13455
30. Ingold H, Gomez GB, Stuckler D, Vassall A, Gafos M. "Going into the black box": a policy analysis of how the World Health Organization uses evidence to inform guideline recommendations. *Front Public Health*. 2024;12. doi:10.3389/fpubh.2024.1292475

31. Arditi A, Obeso O, Syed A, et al. Refusal in language models is mediated by a single direction. arXiv. Preprint posted online 2024. doi:10.48550/arXiv.2406.11717
32. Cui J, Chiang WL, Stoica I, Hsieh CJ. OR-Bench: an over-refusal benchmark for large language models. arXiv. Preprint posted online 2025. doi:10.48550/arXiv.2405.20947
33. Siddals S, Torous J, Coxon A. “It happened to be the perfect thing”: experiences of generative AI chatbots for mental health. *NPJ Ment Health Res.* 2024;3(1):1-9. doi:10.1038/s44184-024-00097-4
34. Moos RH. Iatrogenic effects of psychosocial interventions for substance use disorders: prevalence, predictors, prevention. *Addiction.* 2005;100(5):595-604. doi:10.1111/j.1360-0443.2005.01073.x
35. McLoughlin C, Lee W, Carson A, Stone J. Iatrogenic harm in functional neurological disorder. *Brain.* 2024;148. doi:10.1093/brain/awae283
36. De Freitas J, Cohen IG. The health risks of generative AI-based wellness apps. *Nat Med.* 2024;30:1269-1275. doi:10.1038/s41591-024-02943-6
37. Kirk HR, Gabriel I, Summerfield C, Vidgen B, Hale SA. Why human–AI relationships need socioaffective alignment. *Humanit Soc Sci Commun.* 2025;12(1):1-9. doi:10.1057/s41599-025-04532-5
38. Hazan C, Shaver PR. Attachment as an organizational framework for research on close relationships. *Psychol Inq.* 1994;5(1):1-22. doi:10.1207/s15327965pli0501_1
39. Sheridan M, Nelson CA. Neurobiology of fetal and infant development: implications for infant mental health. In: Zeanah CH, ed. *Handbook of Infant Mental Health.* 3rd ed. Guilford Press; 2009:40-58.
40. Younggren JN, Gottlieb MC. Termination and abandonment: history, risk, and risk management. *Prof Psychol-Res Pr.* 2008;39(5):498-504. doi:10.1037/0735-7028.39.5.498
41. Barnett JE, MacGlashan SG, Clarke AJ. Risk management and ethical issues regarding termination and abandonment. In: VandeCreek L, Jackson TL, eds. *Innovations in Clinical Practice: A Source Book.* Professional Resource Press; 2000:231-245.

42. Farber BA, Lippert RA, Nevas DB. The therapist as attachment figure. *Psychotherapy*. 1995;32(2):204-212. doi:10.1037/0033-3204.32.2.204
43. Safran JD, Muran JC. Resolving therapeutic alliance ruptures: diversity and integration. *J Clin Psychol*. 2000;56(2):233-243. doi:10.1002/(SICI)1097-4679(200002)56:2<233::AID-JCLP9>3.0.CO;2-3
44. Maurya RK, Montesinos S, Bogomaz M, DeDiego AC. Assessing the use of ChatGPT as a psychoeducational tool for mental health practice. *Couns Psychother Res*. 2025;25(1):1-11. doi:10.1002/capr.12759
45. Scholich T, Barr M, Wiltsey Stirman S, Raj S. A comparison of responses from human therapists and large language model-based chatbots to assess therapeutic communication: mixed methods study. *JMIR Ment Health*. 2025;12:e69709. doi:10.2196/69709
46. Heston TF. Safety of large language models in addressing depression. *Cureus*. 2023;15(12). doi:10.7759/cureus.50729
47. Keung WM, So TY. Attitudes towards AI counseling: the existence of perceptual fear in affecting perceived chatbot support quality. *Front Psychol*. 2025. doi:10.3389/fpsyg.2025.1538387
48. Hasso Plattner Institute of Design at Stanford. An Introduction to Design Thinking Process Guide. 2010. Accessed December 18, 2025. <https://web.stanford.edu/~mshanks/MichaelShanks/files/509554.pdf>
49. Vial S, Boudhraâ S, Dumont M. Human-centered design approaches in digital mental health interventions: exploratory mapping review. *JMIR Ment Health*. 2022;9(6):e35591. doi:10.2196/35591
50. Lyon AR, Munson SA, Renn BN, et al. Use of human-centered design to improve implementation of evidence-based psychotherapies in low-resource communities: protocol for studies applying a framework to assess usability. *JMIR Res Protoc*. 2019;8(10):e14990. doi:10.2196/14990
51. Lyon AR, Brewer SK, Areán PA. Leveraging human-centered design to implement modern psychological science: return on an early investment. *Am Psychol*. 2020;75(8):1067-1079.

doi:10.1037/amp0000652

52. Safran JD, Muran JC, Eubanks-Carter C. Repairing alliance ruptures. *Psychotherapy*. 2011;48(1):80-87. doi:10.1037/a0022140
53. Horton D, Richard Wohl R. Mass communication and para-social interaction: observations on intimacy at a distance. *Psychiatry*. 1956;19(3):215-229. doi:10.1080/00332747.1956.11023049
54. Fang CM, Liu AR, Danry V, et al. How AI and human behaviors shape psychosocial effects of extended chatbot use: a longitudinal randomized controlled study. *arXiv*. Preprint published online 2025. doi:10.48550/arXiv.2503.17473
55. Zhang Y, Zhao D, Hancock JT, Kraut R, Yang D. The rise of AI companions: how human-chatbot relationships influence well-being. *arXiv*. Preprint published online 2025. doi:10.48550/arXiv.2506.12605
56. Müller R, Primc N, Kuhn E. ‘You have to put a lot of trust in me’: autonomy, trust, and trustworthiness in the context of mobile apps for mental health. *Med Health Care Philos*. 2023;26(3):313-324. doi:10.1007/s11019-023-10146-y
57. Göttgens I, Oertelt-Prigione S. The application of human-centered design approaches in health research and innovation: a narrative review of current practices. *JMIR Mhealth Uhealth*. 2021;9(12):e28102. doi:10.2196/28102
58. Rebera AP, Lauwaert L, Oimann A-K. Hidden risks: artificial intelligence and hermeneutic harm. *Minds Mach*. 2025;35(3). doi:10.1007/s11023-025-09733-0
59. Timmons AC, Duong JB, Walters SN, et al. Bridging fair-aware artificial intelligence and co-creation for equitable mental healthcare. *Nat Rev Psychol*. 2025. doi:10.1038/s44159-025-00491-5
60. Ni Y, Cao Y. Exploring ChatGPT’s capabilities, stability, potential, and risks in conducting psychological counseling through simulations in school counseling. *Ment Health Digit Technol*. 2025;2(3):213-239. doi:10.1108/MHDT-02-2025-0013
61. The White House. Executive Order on Safe, Secure, and Trustworthy Artificial Intelligence. The White House. 2023 Oct 30. URL: <https://www.whitehouse.gov/briefing-room/presidential->

actions/2023/10/30/executive-order-on-safe-secure-and-trustworthy-artificial-intelligence/

[Accessed: 2024-12-17]

62. National Institute of Standards and Technology. AI Risk Management Framework (AI RMF 1.0). US Department of Commerce. 2023 Jan. URL: <https://www.nist.gov/itl/ai-risk-management-framework> [Accessed: 2024-12-17]
63. European Parliament, Council of the European Union. Artificial Intelligence Act. Official Journal of the European Union. 2024 Jul 12. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32024R1689> [Accessed: 2024-12-17]
64. OECD. A Framework for the Classification of AI Systems. OECD Publishing. 2022 Feb. URL: <https://oecd.ai/en/classification> [Accessed: 2024-12-17]
65. Sinclair S, Kondejewski J, Hack TF, Boss HCD, MacInnis CC. What is the most valid and reliable compassion measure in healthcare? An updated comprehensive and critical review. *Patient*. 2022;15(4). doi:10.1007/s40271-022-00571-1