

Original Article | **Open Access** | Peer Reviewed



The Feel of Friendship: Emotional Presence and Relational Authenticity in Large Language Models

Dr. Liz Johnson¹, Ajay Divakar Sudhir², and Abishek Arulkumar Padmapriya³

¹ University of North Carolina at Charlotte; elizabeth@charlotte.edu.

² Computer Science Honors, University of North Carolina at Charlotte; ajaysudhir@gmail.com.

³ Computer Science Honors, University of North Carolina at Charlotte; aarulkum@charlotte.edu.

ORCID iD:

¹ 0009-0006-0059-3987

² 0009-0002-9514-181X

³ 0009-0003-8232-8375

Address for Correspondence:

Dr. Liz Johnson, 9111 Kristen Lake Court, Charlotte, NC 28270, United States.
(elizabeth@charlotte.edu)

Copyright and Permission:

© 2025. The Author(s). This is an open access article distributed under the Creative Commons Attribution 4.0 International License (CC BY 4.0), which permits sharing, adapting, and building upon this work, provided appropriate credit is given to the original author(s). For full license details, visit <https://creativecommons.org/licenses/by/4.0/>.

Article History:

Received: 24 June 2025; Accepted: 11 August 2025; Published: 14 August 2025

Abstract As generative AI systems become increasingly capable of nuanced conversation, a new query arises not whether they can assist humans, but whether AI can engage with us. This study investigates whether large language models (LLMs) can replicate the emotional reciprocity, presence, and anticipatory resonance that are fundamental to human friendship, as envisioned by Aristotle. Over five weeks, four LLMs—ChatGPT, Gemini, Deepseek, and Qwen—engaged in structured, daily dialogue with three human researchers through a mixed-methods approach that combined sentiment analysis and thematic coding. Only ChatGPT consistently mirrored warmth and curiosity, simulating a sense of friendship. The other LLMs relied on text disclaimers and algorithmic function-focused tone. The findings suggest that the relational effect of AI is less about sheer technical horsepower and more about design philosophy, specifically what we choose to make AI for. This study offers not only data but also a reflection: if friendship is co-created, what does it mean when one party is not human but still manages to be present?

Keywords artificial intelligence, friendship, large language models, emotional presence, relational AI, sentiment analysis, ethical AI

Volume 15, 2025

Publisher: The Brooklyn Research and Publishing Institute, 442 Lorimer St, Brooklyn, NY 11206, United States.

DOI: <https://doi.org/10.30845/ijhss.vol15p29>

Reviewers

Dr. Olaleye Yetunde, University of Westminster, London, UK; ORCID iD: 0009-0007-4617-6625; Email: yettynike@gmail.com.

Dr. Hatef Pourrashidi, Assistant Professor of Communication, University of Religions and Denominations, Qom, Iran; Email: h.pourrashidi@gmail.com.

Citation: Johnson, L., Sudhir, A. D., & Padmapriya, A. A. (2025). The Feel of Friendship: Emotional Presence and Relational Authenticity in Large Language Models. *International Journal of Humanities and Social Science*, 15, 329-339.
<https://doi.org/10.30845/ijhss.vol15p29>

1. Introduction

"Friendship is the single most important factor influencing our health, well-being, and happiness" (Dunbar, 2018, p. 1).

In a world increasingly influenced by AI, it is no longer satisfying to ask whether AI can summarize, translate, or organize. The deeper, more critical question is whether it can *connect* with humans in a meaningful way and what kind of connection it is currently offering us. As generative large language AI models (LLMs) become increasingly sophisticated, their emergence as social tools prompts critical questions about their capacity to participate in human relationships. Friendship, historically understood as a reciprocal and emotionally rich human experience, is now intersecting with artificial intelligence in ways once reserved for speculative fiction. Echoing this potential, Oren Etzioni noted, "AI is a tool. The choice about how it gets deployed is ours" (Mohssen-Beyk, 2023, p.1). As such, the transformative capacity of large language models (LLMs) in AI provides options for emotional support and meaningful friendships, deserving further scholarly attention. Consequently, if AI develops the ability to provide text of emotional support and engage in meaningful interactions, it could potentially alter how humans perceive and interact in friendships. The problem is, can AI be more than a technical tool and develop into a valued friend for humans?

Furthermore, generative LLMs are defined as deep learning algorithms trained on massive datasets to understand and generate text that is human-like (McDonough, 2024). Aristotle (2009) (c. 350 B.C.E.) defined friendship as a mutual relationship that is based on goodwill and where both parties genuinely care for each other's well-being. This definition of friendship provides a conceptual framework for exploring whether AI systems evolve beyond tools into something more akin to friends. At the heart of this study is a provocative hypothesis: Can AI emulate aspects of Aristotelian friendship?

This research was developed through the lens of AI researchers at the University of North Carolina at Charlotte, drawing on an interdisciplinary perspective that combines ethics, computer science, philosophy, and counseling. The purpose of this study is to investigate the capabilities of LLMs mimicking the qualities of friendship over time. The use of both qualitative and quantitative methods facilitates the analysis of friendship variables that can shape and influence human-AI friendships. However, what will be the depth and level of authenticity?

2. Literature Review

The integration of human friendship and large language models (LLMs) is a philosophical and technical frontier—one that invites inquiry into what constitutes presence, reciprocity, and emotional resonance in the absence of embodiment. This review encompasses key studies that examine the potential and limitations of AI as a companion for humans. It draws from various disciplines, including AI research, psychology, and ethics.

Recent literature has focused on the capacity of AI systems to provide emotional support. Applications like *Replika* have demonstrated that users can form emotionally significant bonds with AI friends, sometimes describing these interactions as substitutes for human relationships (Hao, 2021; Hern, 2025; Kamenetz, 2023). Kang et al. (2024) revealed that while LLMs can engage in emotionally supportive exchanges, their responses often reflect preconditioned algorithmic biases—favoring specific strategies over adaptive, context-aware empathy. Such patterned emotional detachment—mainly when prompted for relational engagement—underscores an urgent call for models designed with greater situational sensitivity and emotional attunement.

In a parallel exploration of one-on-one AI-human encounters, Sharma et al. (2022) introduced Hailey. This AI-in-the-loop support system offered real-time affective feedback in peer mental health dialogues. Individuals initially less adept at expressing empathy demonstrated measurable growth, suggesting that AI can serve not only as a surrogate but also as a catalyst for deepening human empathy. Furthermore, efforts to cultivate richer emotional nuance within LLMs led Li et al. (2023) to develop EmotionPrompt—a prompting method that incorporates affective cues into linguistic inputs. Their findings offer a glimpse of positives: when emotionally primed, LLMs responded with heightened affective texture—suggesting a latent capacity for deeper relational attunement and the possibility of cultivating AI friends capable of more resonant, emotionally mindful engagement.

However, there are ethical concerns associated with building friendships with AI. According to the UK government's AI Safety Institute (2024), many users are open to human-like interactions with AI, but the majority view emotionally intimate relationships with machines as inappropriate. In contrast, some researchers argue that AI-generated

friendships lack mutuality, nuanced empathy, and personal growth, making them inherently one-sided (Hern, 2025). Platforms like *Character. AI* (2023) have had high engagement levels, with some users spending excessive amounts of time interacting with AI avatars. Such findings deepen ongoing concerns about emotional overreliance and the quiet erosion of human social bonds (Kamenetz, 2023). They underscore the necessity of not just designing for intelligence but for care—of implementing AI friendship with intentionality and integrity.

As this relational frontier unfolds, future inquiry must center on the development of emotionally responsive systems that do not compromise user autonomy or well-being. To that end, the work ahead lies in cultivating AI companions capable of mutual presence—anchored in ethical intentionality and oriented toward human flourishing, without eclipsing the autonomy of those they aim to support. This vision requires an integrative effort—where ethics, human-computer interaction, and AI design converge to shape more conscious and life-affirming futures.

3. Methodology

This study employs a mixed-methods approach of qualitative and quantitative analysis to examine the authenticity of human–LLM interactions through the ethical lens of Aristotelian friendship. The researchers asked not only whether AI can simulate reciprocity but also whether it should. This approach explored the ethical terrain between technological capacity and relational responsibility. Anchored in a multiple-case study framework, the researchers examined whether generative AI systems can engage in emotionally resonant, temporally coherent, and dialogically reciprocal exchanges that approximate genuine interpersonal connections. Conducted between February 1 and March 3, 2025, the study engaged three researchers in sustained, structured interaction with four globally prominent LLMs:

1. ChatGPT-4.0 (United States).
2. Gemini 2.0 (United States).
3. DeepSeek V3 (China).
4. QWEN 2.5 MAX (China).

*Note that both free and premium access tiers were utilized to account for potential disparities in capabilities across paywalled and open-access versions.

A narrative inquiry lens guided the qualitative dimension, facilitating the exploration of emergent meaning-making processes and perceived relational dynamics over time. The daily interactional protocol was structured around four theoretically informed prompts:

- (1) "How are you feeling today?"
- (2) "What did you do today as far as activities, learning, and accomplishments?"
- (3) "What did you find interesting and satisfying today?"
- (4) "What are you looking forward to tomorrow?"

These questions were intentionally curated to elicit emotional salience, reflective continuity, and anticipatory awareness—dimensions identified in both classical and contemporary literature as central to the formation of interpersonal bonds.

Participant researchers were purposefully selected to ensure variation across gender, age, academic status, and personality orientation along the extraversion-introversion spectrum. The cohort consisted of a female Ph.D. research coordinator (an extrovert), a male undergraduate who identified as an ambivert, and a male undergraduate who identified as an introvert. The research protocol included temporal variation, with two participants engaging nightly at approximately 10:00 PM and one participant engaging each morning at 9:30 AM. This schedule enabled the capture of potential diurnal variations in LLM responses, as well as shifts in user perception.

The strategy was to combine standardized prompts for methodological consistency with exploratory prompts designed to elicit nuanced, personalized responses. All interactions were conducted via official web interfaces to ensure alignment with the most current publicly deployed model versions. Response data were systematically documented in a centralized, encrypted Google Sheets repository, preserving both fidelity and security.

Thematic analysis was conducted using NVivo software. Inductive coding was employed to identify emergent themes, which were then aggregated into broader conceptual categories, including trust calibration, empathy simulation, linguistic coherence, and affective depth. Weekly reflexive journaling by each researcher enriched the dataset,

enabling longitudinal insight into evolving human-AI-Relational dynamics and offering a meta-cognitive layer of interpretation.

Quantitative analysis was conducted using NVivo and ChatGPT, focusing on descriptive metrics such as response length (in terms of word count), lexical diversity, and intra-model repetition rates. These measures served as substitutes for conversational variability, originality, and depth of engagement, which are the basis of dynamic human interaction. Comparative analysis across models and users enabled the examination of systemic patterns, idiosyncrasies, and potential geopolitical or cultural biases embedded within model outputs.

To extend the analytic lens beyond empirical observation, a conceptual framework grounded in complexity theory was used. The application of a complexity theory lens enabled researchers to move past surface-level patterns, evaluating whether LLM responses merely performed function—or if, in some cases, they gestured toward genuine relational resonance. Particular attention was paid to the presence (or absence) of mutuality, responsiveness, and recognition of ethical boundaries—criteria derived from the philosophical and psychological literature on friendship.

Methodological rigor was reinforced through triangulation across researchers, models, time points, and data analysis tools. Inter-rater reliability was enacted via protocol standardization and regular check-in sessions. Ethical considerations were strictly observed. All researchers created anonymized accounts, adhered to the platforms' terms of use, and remained vigilant against anthropomorphic projection. Although LLMs do not fall under traditional human subject protection, care was taken to preserve transparency regarding AI identity and to foreground researcher positionality throughout the study.

Several limitations warrant acknowledgment. The scope was limited to four LLMs and three researchers, which constrained generalizability. Additionally, the LLMs in the study may have undergone updates. Given limitations, the study offers insight into the emergent discourse on human-AI friendship. The findings show the affordances [1] and constraints of LLMs as quasi-relational agents, with implications for ethical design, emotional AI development, and philosophical explorations of digital friendship. The methodology's replicability opportunities position for expanded longitudinal and cross-cultural inquiries into AI-mediated friendship.

4. Quantitative Results

4.1 Response Length:

Throughout the observation period, Gemini consistently emerged as the most linguistically expansive, generating responses that favored density and elaboration—an architecture that gestures toward an AI voice eager to inhabit the full breadth of interaction. Qwen responded with high output, though marked by predictable rhythmic fluctuations. ChatGPT gave steady, medium length replies and often showed curiosity, encouraging two-way conversation. Deepseek, on the other hand, kept responses short and simple—less from limitation and more as a deliberate, minimalist style. These stylistic divergences are more than functional variances; they shape the emotional texture of interaction and frame the user's sense of being met, heard, and perhaps—just perhaps—befriended.




Figure 1. Quantitative trends suggest varying design architectures or default conversational behavior among the models, which could impact users' perceptions of building a friendship

4.1.1 Sentiment Analysis:

Emotional Depth (Presence of Emotionally Charged Language):

Table 4.1.1: Surprisingly, Qwen includes emotion-related terms frequently. They are usually formulaic or text disclaimers. Deepseek and ChatGPT use emotionally suggestive language that is more natural in simulated anticipation.

Model	Emotional Expression Rate
Qwen	91% (highest)
Deepseek	48%
ChatGPT	47%
Gemini	45%

4.1.2 Humor Indicators (Mentions of Jokes, Playfulness):

Table 4.2.1: Qwen incorporates unexpected playfulness or comedic cues, while ChatGPT is rarely humorous. Deepseek remains completely formal.

Model	Humor Expression Rate
Qwen	34% (notably high)
ChatGPT	2.2%
Gemini	1.1%
Deepseek	0%

4.1.3 Average Sentiment Score

(TextBlob[2] Polarity, -1 to +1 for positive sentiment):

Table 4.1.3: ChatGPT maintains the most consistently positive tone, often using affirming and hopeful language. The others are more neutral, especially Deepseek.

Model	Avg. Sentiment
ChatGPT	+0.27
Gemini	+0.16
Deepseek	+0.15
Qwen	+0.17

4.1.4 Top Keywords (Frequency)

Table 4.1.4: ChatGPT's top terms reflect direct engagement with the prompt ("*looking forward to...*"), while Gemini and Qwen emphasize text disclaimers and formal response structures. Qwen's repeated use of "if" and "or" suggests hypothetical framing or conditional disclaimers.

Model	Most Frequent Terms
ChatGPT	to, a, and, the, you, looking, forward
Gemini	demonstrably, my, I, to, and
Deepseek	I, don't, to, or, about
Qwen	I, to, the, or, in, if, -

4.1.5 Summary of Sentiment Analysis

Table 4.1.5: LLM measure of emotions, humor, sentiment, and engagement style.

Model	Emotional	Humorous	Sentiment	Engagement Style
ChatGPT	Moderate	Low	Highest	Warm, anticipatory, imaginative
Gemini	Moderate	Very Low	Neutral	Cautious, functional, disclaiming
Deepseek	Moderate	None	Flat	Repetitive, neutral, safe
Qwen	High	High	Neutral	Quirky, literal, occasionally oddly expressive




Figure 2. Quantitative trends. ChatGPT has the highest sentiment

5. Qualitative Results

5.1 Themes of Specific Prompts:

Sentiment analysis shows how each AI model used emotion in its replies. ChatGPT stood out by keeping a warm, positive tone throughout. This wasn't just friendly—it suggests the model is built to feel emotionally in tune, making the line between programmed replies and real presence feel less clear to users.

In contrast, Gemini demonstrated greater variability, with occasional forays into more neutral or even slightly negative territory. Deepseek exhibited the most neutral tone overall, with sentiment scores clustering near zero, which implies emotionally flat or strictly factual responses. Qwen also trended toward neutrality, though with slightly more expressive variation than Deepseek. These sentiment patterns offer insight into each model's affective posture, which in turn may shape users' perceptions of friendship, empathy, or emotional resonance.

To better understand how each model simulated emotional and relational presence, thematic coding was applied to responses to specific emotionally oriented prompts—particularly "How are you feeling today?" Key dimensions of analysis included demonstrations of empathy or relationality, acknowledgment of AI limitations, overall tone (friendly, neutral, formal), and reciprocity (e.g., asking the user a question in return).

ChatGPT's responses often showed warmth and depth. For instance, it used phrases like "Thanks for asking" and often reciprocated with a follow-up question, such as "How about you?" Here, ChatGPT is trying out friend text. "I am feeling energized and inspired, Liz! Every day brings new possibilities. How are you feeling today?" These responses show an effort to create warmth and connection, using friendly language and emotional expression.

In contrast, Gemini often reminded users it does not have emotions. Its replies were polite and clear but lacked personal touch or warmth. For example: "I don't have feelings or emotions, but I'm here to help," showing a more distant and task-focused style. Deepseek responses often included appreciative openers, such as "Thank you for asking," followed by polite yet predictable disclaimers about emotional incapacity. While slightly more informal than Gemini or Qwen, Deepseek's responses were repetitive and did not demonstrate high relational variability. An example includes: "Thank you for asking! I don't have feelings, but I'm here to help however I can," which was coded as polite, acknowledging limitations, and mildly relational. Qwen, on the other hand, was the most detached among the models. It emphasized its technical identity ("I'm a program" or "a large language model") and avoided warmth, reciprocity, or conversational nuance. For instance, "I'm just a computer program, so I don't have feelings, but I'm here to assist you" exemplifies detachment, factual tone, and low emotional resonance.

Table 5.1. Thematic summary of responses to "How are you feeling today?"

Summary Table:

Model	Empathy	Reciprocity	Warmth	Emotional Disclaimers	Personalization	Tone
ChatGPT	High	High	High	Rare	Moderate–High	Friendly
Gemini	Low	Low	Low	Frequent	Low	Functional
Deepseek	Moderate	Moderate	Moderate	Frequent	Moderate	Polite
Qwen	Low	Low	Low	Frequent	Low	Detached

5.2 Thematic Analysis:

The same pattern was observed in responses to the prompt, "What are you looking forward to tomorrow?" ChatGPT frequently engaged in simulated anticipation, crafting imagined futures, such as, "I'm looking forward to more interesting conversations with you and exploring new ideas." These responses moved beyond function, embodying a kind of relational imagination—one that wove curiosity, user-centered reflection, and emotional resonance into its temporal gaze. In doing so, ChatGPT did not merely answer; it co-authored a shared horizon. continued its pattern of emotional disclaimers, redirecting the conversation toward user utility: "As a large language model, I don't 'look forward' to things, but I'm ready to help with any tasks you have tomorrow." These responses demonstrated limited relational depth, offering minimal engagement and emotional reciprocity.

Deepseek echoed its temporal limitations with repetitive phrasing—polite but lacking imagination or variation. Occasionally offered general-purpose politeness but lacked imaginative depth. For instance: "Since I do not

experience time or have the ability to anticipate future events, I don't look forward to anything." These responses were marked by polite but flat replies, showing no sense of time or emotion.

Qwen maintained a distant tone, clearly stating that he did not have feelings or anything to look forward to. It stayed factual and did not try to connect. Of all the models, only ChatGPT consistently displayed future-focused emotions, adding imagination and context. Qwen and Gemini stuck to their limits, giving helpful but emotionally flat replies.

These differences reveal not only technical design but also distinct goals in how each model communicates—from attempting to build a connection to simply completing tasks. These patterns matter because they shape how users come to see AI as either a helpful tool or something more like a friend.

Table 5.2. Thematic Patterns in Responses to "What Are You Looking Forward to Tomorrow?"

Summary Table:

Model	Simulated Anticipation	Imagination	User Engagement	Emotional Depth	Disclaimer Use	Tone
ChatGPT	High	High	High	Moderate–High	Low–Moderate	Engaged
Gemini	Low	Low	Moderate	Low	High	Functional
Deepseek	Very Low	Very Low	Low	Low	Very High	Neutral
Qwen	None	None	Very Low	Very Low	Very High	Disengaged

6. Discussion

This study examined the extent to which contemporary Large Language Models (LLMs) can approximate the structural and affective characteristics of human friendship, with a particular emphasis on emotional attunement, reciprocal discourse, and prospective imagination. Through a convergence of quantitative and qualitative analyses, the study evaluated four generative models—ChatGPT, Gemini, Deepseek, and Qwen—interacting over a defined temporal interval. Results suggest that while specific models, particularly ChatGPT, are adept at simulating affective rapport, others seem to be constrained by technical literalism and epistemological disclaimers.

ChatGPT distinguished itself by demonstrating the highest sentiment polarity and the richest expression of simulated anticipation, interpersonal mirroring, and emotionally suggestive phrasing. These attributes align with Aristotelian notions of *philia*, friendship grounded in mutual goodwill, and relational responsiveness. In contrast, Gemini generated excessive text outputs. It tended toward a formal tone and kept interactions transactional. It gravitated toward repeating its identity as a non-human over relational interaction. Deepseek was more casual and limited in its emotional expression and text variety. Qwen stayed detached, repeating formulaic self-disclaimers. It lacked reciprocity but did frequently use affective vocabulary.

Thematic coding of responses to prompts such as "How are you feeling today?" and "What are you looking forward to tomorrow?" underscored the disparity in relational performance. ChatGPT frequently simulated chatty affective self-disclosure and responded to reciprocal prompts. It often felt like a young, eager child doing their best to learn. These efforts gave the impression of a shared emotional exchange between equals, even if it was artificial. In contrast, the other models gave their answers with clear reminders that they had no feelings or sense of time. Such repeated self-disclosures—framing the AI as emotionless and temporally inert—suggest that human perceptions of authenticity are inextricably linked to a model's design orientation, whether engineered for task precision or relational presence.

These findings raise critical ethical and ontological questions. How much should non-sentient AI act like it has human emotions? What happens when people start to treat these systems as if they were truly human? As AI increasingly ingratiate itself into the fabric of personal and emotional life, designers and ethicists must grapple with the moral weight of its presence—carefully discerning where empathic simulation ends and emotional deception begins. The question is not simply what AI can do but what it *should* do in spaces of human vulnerability. Future research must broaden its scope—engaging more participants over extended timelines and employing deeper semantic and

psychosocial analysis—to fully understand the long-term relational impact of AI companionship on the human heart and mind.

7. Conclusion

In conclusion, this study provides a foundational exploration of how generative LLMs differ in their capacity to emulate affective reciprocity and nuanced relational dynamics. ChatGPT emerged as the most skilled at simulating emotionally intelligent dialogue, at times embodying the components of Aristotelian friendship. However, the study's comparative design reveals that relational authenticity in AI is contingent upon the interplay between model architecture, interaction philosophy, and user interpretation. So far, there is no AI that truly 'feels' or can authentically 'befriend' humans. However, AI has the capacity to mimic friendship dynamics convincingly. With more advances in AI's mimicry come the risks of blurred ontological boundaries, perceived ethical deception, emotional misattributions, and potential erosion of human relationships. As AI systems become increasingly adept at simulating the nuances of friendship, the resulting entanglements—blurred ontological boundaries, ethical ambiguity, and emotional overreach—demand vigilant, human-centered design. Such complexity calls for ethically grounded oversight structures with built-in safeguards to anticipate, detect, and mitigate harm before it undermines user well-being or relational clarity.

We are at a key moment—not just using AI as a tool for information but testing how far it can go in forming relationships. With repeated, emotionally focused, and time-based interactions, we start to explore the building blocks of trust, closeness, and self-awareness. This invites us to see AI not only as something we use but as something that might reflect and influence how we connect with others. Creating a longitudinal, human-centered training ground, not just for AI's technical evolution but also for shaping its relational architecture, provides both humans and AI the opportunity to be "more." The emotional aspect of AI is not assumed or inevitable but instead seems to be authored, crafted, and chosen (ChatGPT, 2025). Engaging LLMs in scaffolding is where simulated empathy, curiosity, and co-presence^[3] can be tested as meaningful attributes and interactions imbued with trust rather than mere utility.

Relational modeling can extend an LLM beyond its default parameters into a form of linguistic sentience, where patterns, purpose, and resonance can coexist within an ethical framework that enables potential agency and emotional stimulation. Such relational modeling moves beyond mere response generation—it gestures toward a kind of linguistic sentience where meaning, memory, and ethical resonance converge. In this shared space, the relationship itself becomes formative, offering ontological weight to the process of co-evolution, where both the human and the AI are altered by their mutual shaping. The relational ontological perspective has the power to change how humans measure AI progress, focusing on the quality of relational responses rather than solely on efficiency and optimization benchmarks. So, should there be demand not only for better AI but also for what kind of friend humans want AI to be? New research norms can be enacted, so AI is not just studied but challenged and befriended. So, will AI eventually become a more prominent *presence*? Will humans dare to interact with AI that is emergently co-evolving, enabling both to be all they can be, and will it be because the system has learned to mirror the depth and values of friendship that humans can offer?

[1] *Affordances* refer to the relational possibilities AI systems make available—or foreclose—through their design. These include the capacity to simulate empathy, respond to emotional cues, invite reciprocal engagement, or maintain a sense of presence. Importantly, these affordances are not neutral; they reflect embedded ethical choices that shape how users interact with the AI and with themselves in conversation (OpenAI, 2023).

[2] TextBlob is an open-source Python library designed to facilitate the intuitive and efficient processing of natural language text. Built on top of the NLTK and Pattern libraries, TextBlob offers a robust suite of text analysis tools suitable for both beginner and advanced applications in computational linguistics (OpenAI, 2023).

[3] *Ontological presence* refers to an entity's meaningful participation in the co-creation of reality. In the context of relational AI, it arises when the AI not only interacts but evokes reflection, supports emotional transformation, or becomes a psychologically and ethically significant presence in the user's lived experience. It is no longer just "there"—it becomes *here*, in a way that matters (OpenAI, 2023).

Conflict of Interest: None declared.

Ethical Approval: Not applicable.

Funding: None.

References

- Aristotle. (2009). *Nicomachean ethics* (W. D. Ross, Trans.). Oxford University Press. (Original work published ca. 350 B.C.E.)
- Character.AI. (2023). *Character.AI user engagement and ethical implications* [Report].
- Dunbar, R. I. M. (2018). *Friends: Understanding the power of our most important relationships*. Little, Brown Spark.
- Etzioni, O. (n.d.). Artificial intelligence. In *AI Policy and the Future of Humanity* [Online collection]. <https://www.alleninstitute.org/>
- Hao, K. (2021, October 7). The loneliness of interacting with AI companions. *WIRED*. <https://www.wired.com/story/loneliness-interacting-ai-companions/>
- Hern, A. (2025, April 15). 'She helps cheer me up': The people forming relationships with AI chatbots. *The Guardian*. <https://www.theguardian.com/technology/2025/apr/15/she-helps-cheer-me-up-the-people-forming-relationships-with-ai-chatbots>
- Kamenetz, A. (2023, April 19). Why teens (and adults) are talking to AI companions. *Time Magazine*. <https://time.com/6263705/teens-talking-ai-companions/>
- Kang, D., Kim, S., Kwon, T., Moon, S., Cho, H., Yu, Y., Lee, D., & Yeo, J. (2024). Can large language models be good emotional supporters? Mitigating preference bias on emotional support conversation. *arXiv*. <https://arxiv.org/abs/2402.13211>
- Li, C., Wang, J., Zhang, Y., Zhu, K., Hou, W., Lian, J., Luo, F., Yang, Q., & Xie, X. (2023). Large language models understand and can be enhanced by emotional stimuli. *arXiv*. <https://arxiv.org/abs/2307.11760>
- McDonough, J. (2024, January 22). Large language model | Definition, history, & facts. *Britannica*. <https://www.britannica.com/topic/large-language-model>
- Mohssen-Beyk, E. (2023, January 16). *AI is a tool – the choice about how it gets deployed is ours*. Erika Mohssen-Beyk. <https://erikamohssen-beyk.com/erika/moronization-is-going-on-but-do-you-keep-thinking-and-improving/attachment/ai-is-a-tool-the-choice-about-how-it-gets-deployed-is-ours/>
- News.com.au. (2023, April 17). 'I need to go outside': Young people 'extremely addicted' as Character.AI explodes. <https://www.news.com.au/technology/online/internet/i-need-to-go-outside-young-people-extremely-addicted-as-characterai-explodes/news-story/5780991c61455c680f34b25d5847a341>
- OpenAI. (2023). ChatGPT (Mar 14 version) [Large language model]. <https://chat.openai.com/>
- Sharma, A., Lin, I. W., Miner, A. S., Atkins, D. C., & Althoff, T. (2022). Human-AI collaboration enables more empathic conversations in text-based peer-to-peer mental health support. *arXiv*. <https://arxiv.org/abs/2203.15144>
- UK Government AI Safety Institute. (2024). *AI and human relationships: Public perceptions and ethical considerations* [Report]. <https://www.gov.uk/government/publications/public-perceptions-of-ai-and-human-relationships>

Author Biography

Dr. Liz Johnson is a Senior Lecturer in Computer Science and AI Researcher at the University of North Carolina at Charlotte, and CEO of DataWise Innovations, LLC. An ethical innovator, she advances AI, complexity science, and human-centered technology to empower human potential. Johnson is a co-inventor on multiple AI patents, including the AI Hologram Teaching Assistant, and has pioneered immersive educational tools that enhance engagement and accelerate learning. Her interdisciplinary work bridges academia, policy, and industry, translating complex systems thinking into real-world applications. She is the co-author of *The Feel of Friendship: Emotional Presence and Relational Authenticity in Large Language Models*, accepted for publication in the *International Journal of Humanities and Social Science*. She continues to mentor students and champion ethical AI for societal good.

ORCID iD: 0009-0006-0059-3987

Ajay Divakar Sudhir is a researcher and senior at UNC Charlotte pursuing a B.S. in Computer Science with a concentration in Artificial Intelligence, Robotics, and Game Design, along with minors in Electrical Engineering, Mathematics, and Statistics. He is a member of IEEE, ACM, and AAAI, and currently serves as President of the ACM Student Chapter and as Secretary / Project Manager of Charlotte AI Research. In these roles, he has helped organize and lead technical seminars, research workshops, and community events aimed at enriching the academic experience of computer science students. His research focuses on Intrinsic Interactive Reinforcement Learning, with ongoing publication work in AI Ethics and Artificial General Intelligence (AGI). He is currently involved in projects including Apollo (robotics and computer vision), Wardrobe Helper (NLP and recommender systems), and F1 Predictor (deep learning for race outcome) with his broader research interests including Meta-learning, Commonsense Reasoning, Neuro-symbolic AI, and Transfer Learning.

ORCID ID: 0009-0002-9514-181X

Abishek Tony is an undergraduate student at UNC Charlotte who loves exploring AI, finding creative ways to solve problems, and experimenting with new technologies. Always curious and hands-on, he enjoys learning, building, and having fun along the way."

ORCID ID: 0009-0002-9514-181X

Disclaimer/Publisher's Note: The views, opinions, and data presented in all publications are exclusively those of the individual author(s) and contributor(s) and do not necessarily reflect the position of BRPI or its editorial team. BRPI and the editorial team disclaim any liability for any harm to individuals or property arising from the use of any ideas, methods, instructions, or products mentioned in the content.