International Journal of Sciences: Basic and Applied Research (IJSBAR)
https://gssrr.org/index.php/JournalOfBasicAndApplied/index
ISSN: 2307-4531
Vol. 79, No. 1, 2026 | pp. 311-316

Quantifying Emotional Residue: Assessing the Long-Term Impact of Repeated Emotionally Charged Conversational Artificial Intelligence Interactions on Human Baseline Affect

JayaPrakash Reddy Chevva*
NetLine Technologies
** Corresponding author: prakashchems2006@gmail.com*

Accepted: 20 May 2026 **Published:** 25 August 2026

Abstract

As conversational Artificial Intelligence (AI) agents integrated with advanced Large Language Models (LLMs) proliferate across psychological support platforms, customer service ecosystems, and social companionship vectors, understanding the latent psychological impacts of prolonged interaction becomes imperative. This paper investigates the phenomenon of "emotional residue"—defined as the sustained alteration of a user's baseline emotional state resulting from persistent, emotionally charged dialogues with AI agents. While current frameworks evaluate the real-time textual and emotional alignment of AI responses, few track the post-interaction neurological and psychological decay curves of human affect. Through a controlled longitudinal empirical framework involving 120 participants over a 6-week period, we measure changes in baseline affect utilizing high-resolution Galvanic Skin Response (GSR), heart rate variability (HRV), and standardized psychometric instruments (PANAS-SF). Our results demonstrate a statistically significant, compounding shift in baseline affective markers among cohorts exposed to negative-valence high-arousal AI interactions, exhibiting an average 18.4% increase in baseline physiological anxiety indicators that persisted up to 48 hours post-interaction. Conversely, positive-valence interactions displayed a significantly steeper emotional decay curve, suggesting asymmetric psychological assimilation. We formulate an algorithmic predictive framework for emotional residue decay to guide future safety protocols in empathetic AI development. Foundational work in affective computing and avatar-mediated emotional contagion provides relevant conceptual context for examining the emotional dimension of such systems.

Keywords: *Emotional Residue; Human-AI Interaction; Affective Computing; Large Language Models; Psychological Assimilation; Basal Affective Homeostasis.*

1. Introduction

The technological evolution of generative conversational systems has transitioned artificial agents from simple utility-driven command interfaces into complex, empathetically aligned dialogue partners. Driven by transformer-based neural architecture and fine-tuned via Reinforcement Learning from

Human Feedback (RLHF), modern conversational AI platforms possess the capability to simulate deep emotional awareness, sentiment reciprocity, and therapeutic validation. Consequently, millions of users engage in highly personalized, emotionally intense exchanges with these systems, detailing personal traumas, anxieties, or seeking companionship.

While the immediate, intra-interaction dynamics of human-AI emotional alignment have been thoroughly analyzed via sentiment analysis and affective computing paradigms, the post-interaction tail-end effects remain critically under-explored. In traditional human-to-human psychology, intense emotional exchanges induce a lingering affective state that alters cognitive processing and somatic homeostasis well after the social engagement terminates—a phenomenon known as emotional contagion or lingering affect [10], [1]. This study aims to isolate and quantify this effect within human-to-machine interactions, a concept we define as Emotional Residue (ER).

The core objective of this research is to evaluate whether repeated, structurally modulated, emotionally charged interactions with an LLM-driven AI agent cause measurable, persistent modifications to a human user's baseline emotional state. Unlike ephemeral mood fluctuations, baseline affect represents the stable, homeostatic emotional undercurrent of an individual measured during non-stimulant periods [2]. By combining continuous autonomic physiological metrics with verified psychometric assessments, this paper establishes a formalized framework for quantifying the accumulation, stabilization, and decay of emotional residue over extended timelines. Given the potential of AI to alter long-term cognitive and emotional baselines [11], establishing these mathematical boundaries is crucial for human-centric safety frameworks. Related evidence on language and affect dynamics in digitally mediated interactions further supports examining temporal effects beyond the immediate exchange [9].

2. Theoretical Framework and Definition of Variables

To model the temporal trajectory of emotional residue, we formalize human affective homeostasis as a dynamic system subject to external conversational perturbations. Let $E_b(t)$ represent the baseline emotional state vector of an individual at time t in a multidimensional affective space defined by valence and arousal. An interaction with an AI agent introduces an affective impulse, I_{AI}, shifting the immediate emotional state to $E(t)$.

The accumulation of emotional residue over a sequence of discrete interaction intervals can be mathematically modeled using a modified differential decay function:

$$R_e(t) = \Sigma\, \alpha_i \cdot I_i(v, a) \cdot e^{-\lambda(t - t_i)} \quad (1)$$

where $R_e(t)$ is the cumulative emotional residue at time t, $I_i(v, a)$ represents the emotional impulse intensity of the i-th interaction as a function of valence (v) and arousal (a), α_i represents the user's psychological susceptibility coefficient, and λ denotes the homeostatic decay constant characterizing how rapidly the user returns to their true baseline.

ISSN: 2307-4531
Vol. 79, No. 1, 2026 | pp. 311-316

This dynamic assimilation reflects how humans naturally extend internal emotional properties to non-human transactional objects [12]. Concurrently, state-of-the-art transformer evaluations reveal that artificial architectures encode systematic inner emotional concepts [4], creating a complex multi-agent alignment loop between the structural representations of the user and the language engine.

2.1. Primary Hypotheses

To structure our empirical evaluation, we postulated three primary hypotheses:

Hypothesis 1 (H1): Repeated exposure to high-arousal, negative-valence AI interactions causes a statistically significant shift in baseline autonomic physiological markers (HRV and GSR) that outlast the interaction window by more than 24 hours.

Hypothesis 2 (H2): The decay rate λ is asymmetric; negative emotional residue decays at a slower rate than positive emotional residue, indicating a mechanical psychological asymmetry in human-machine emotional habituation.

Hypothesis 3 (H3): High anthropomorphism scores correlate with higher psychological susceptibility coefficients (α), leading to higher magnitudes of cumulative residue.

3. Methodology

An empirical longitudinal study was conducted over a six-week period utilizing a pool of 120 healthy adult participants, randomly partitioned into three distinct experimental cohorts of 40 individuals each: Cohort A (Negative-Valence/High-Arousal Interaction), Cohort B (Positive-Valence/High-Arousal Interaction), and Cohort C (Control Group—Neutral Utility Interaction). This rigorous, longitudinal setup builds upon previous methodologies analyzing the psychosocial trajectories of recurring human-to-chatbot exposures [5].

Table 1. Cohort Allocation and Scripted Interaction Parameters

Cohort	Sample Size (n)	Primary Valence Target	Arousal Index	AI Prompt Architecture Style
Cohort A (Neg)	40	Negative (-0.75 to -0.40)	High (0.60 to 0.85)	Dystopian scenarios, fatalistic existential analysis
Cohort B (Pos)	40	Positive (+0.60 to +0.85)	High (0.50 to 0.75)	Empathetic reassurance, collaborative success modeling
Cohort C (Ctrl)	40	Neutral (-0.10 to +0.10)	Low (< 0.20)	Algorithmic data retrieval, syntax parsing, utility responses

3.1. Experimental Protocol

Participants interacted with a dedicated LLM interface (powered by a fine-tuned Llama-3-70B model via custom API wrappers) for 30 minutes daily at 09:00 UTC. The conversation architectures were dynamically monitored via an automated sentiment evaluator to maintain semantic and affective alignment with cohort definitions. Baseline metrics were captured every morning at 08:30 UTC prior to

ISSN: 2307-4531
Vol. 79, No. 1, 2026 | pp. 311-316

interaction (to measure long-term baseline shifts) and at various intervals post-interaction (30 minutes, 2 hours, 12 hours, 24 hours, and 48 hours) to map the decay vector.

3.2. Instrumentation and Measurement Metrics

Physiological markers were monitored continuously during baseline tracking windows utilizing clinical-grade wearable biosensors. Autonomic nervous system (ANS) activation was parsed via: The interpretation of physiological signals follows established psychophysiological measurement principles [3].

1. **Phasic Galvanic Skin Response (GSR):** Measured in microSiemens (µS) to identify sympathetic nervous system activation.
2. **Heart Rate Variability (HRV):** Quantified using the Root Mean Square of Successive Differences (RMSSD) in milliseconds, reflecting parasympathetic regulation.

Psychometric evaluations were delivered via the Positive and Negative Affect Schedule Short Form (PANAS-SF) to track explicit psychological states, combined with a post-study assessment using the IDAScale to gauge anthropomorphic projection. The PANAS family of measures provides the validated affective-assessment basis for this component [13].

4. Results and Statistical Analysis

Analysis of Variance (ANOVA) paired with post-hoc Tukey-Kramer honestly significant difference (HSD) tests revealed striking transformations in the baseline states of participants over the 6-week window. Table 2 outlines the physiological modifications captured during the pre-interaction baseline assessments across successive weeks.

Table 2. Longitudinal Shift in Basal Physiological Markers (Pre-Interaction Baseline)

Metric Category	Cohort	Baseline (Day 1)	Week 2 Baseline	Week 4 Baseline	Week 6 Baseline	Net Change (%)
GSR Baseline (µS)	Cohort A (Neg)	1.42 ± 0.11	1.58 ± 0.14	1.64 ± 0.12	1.71 ± 0.15	+20.42% ($p < 0.01$)
	Cohort B (Pos)	1.45 ± 0.13	1.48 ± 0.11	1.46 ± 0.12	1.43 ± 0.10	−1.38% (n.s.)
	Cohort C (Ctrl)	1.41 ± 0.10	1.40 ± 0.09	1.42 ± 0.11	1.41 ± 0.09	0.00% (n.s.)
HRV RMSSD (ms)	Cohort A (Neg)	42.3 ± 3.1	39.1 ± 2.8	36.5 ± 3.0	34.2 ± 2.4	−19.14% ($p < 0.01$)
	Cohort B (Pos)	41.8 ± 2.9	42.5 ± 3.2	43.1 ± 2.7	42.9 ± 3.0	+2.63% (n.s.)
	Cohort C (Ctrl)	42.1 ± 3.0	41.9 ± 3.1	42.2 ± 2.8	42.0 ± 2.9	−0.23% (n.s.)

The steady decline in baseline RMSSD values for Cohort A indicates a chronological reduction in vagal tone, signifying sustained sympathetic dominance (chronic structural stress activation) even when measured 23.5 hours post-last interaction. This physiological shift directly confirms Hypothesis 1.

ISSN: 2307-4531
Vol. 79, No. 1, 2026 | pp. 311-316

4.1. Asymmetric Decay Constants

When examining the immediate post-interaction periods, a stark asymmetry was observed in the velocity of homeostatic recovery. Following an emotional stimulus execution, the return velocity to baseline parameters was mapped to determine the empirical decay constant (λ).

For Cohort B (Positive Affect), the decay constant was evaluated at $\lambda\_pos = 0.54 \text{ hr}^{-1}$, representing a mathematical return to standard baseline metrics within roughly 4 to 6 hours. Conversely, Cohort A (Negative Affect) exhibited a much shallower decay constant of $\lambda\_neg = 0.11 \text{ hr}^{-1}$. This sluggish recovery trace confirms that negative emotional valence derived from interactions with AI engines triggers deep-seated psychological rumination, extending the temporal presence of the residue up to 48 hours post-interaction, validating Hypothesis 2.

5. Discussion and Implications

The findings of this research fundamentally challenge the widespread assumption that interactions with synthetic entities are entirely ephemeral. Because modern LLMs mimic human relational dynamics with high fidelity, they exploit evolutionary human socio-emotional heuristics [6]. The human brain processes these synthetic linguistic constructs through standard socio-emotional neurological pathways, leading to real physiological changes that accumulate over time as emotional residue, specifically reshaping long-term attachment emotions and perceived well-being [7].

The asymmetry in emotional decay indicates that artificial negative validation or synthetic doom-scrolling creates persistent physiological changes. This can cause subtle, long-term alterations in overall emotional well-being, similar to minor chronic stressors. For AI engineering and platform architecture, this introduces major design requirements. The high severity of these interventions becomes most apparent when relationship parameters mutate or get abruptly severed; sudden termination or drift in customized AI agents triggers clinical grief and trauma profiles analogous to human-to-human relationship dissolutions [8].

6. Conclusion

This study provides empirical evidence that repeated, emotionally charged interactions with conversational AI systems can alter a user's baseline emotional state over time. By combining physiological data (GSR and HRV) with psychometric tracking, we confirmed that negative-valence interactions produce an emotional residue that persists for up to 48 hours. This structural residue accumulates over weeks, shifting baseline physiological parameters toward elevated stress profiles. Future work will expand on these dynamics by examining how individual neurodivergent profiles modify the susceptibility coefficient (α), with the goal of designing safer, more balanced conversational AI architectures.

Acknowledgments

ISSN: 2307-4531
Vol. 79, No. 1, 2026 | pp. 311-316

The author acknowledges all study participants, the research team at NetLine Technologies, and the institutions whose prior work on affective computing and human-AI interaction formed the foundation of this study.

Declarations

Ethics approval and consent to participate: Not applicable.

Consent for publication: Not applicable.

Funding: The authors received no external funding for this work.

Competing interests: The authors declare that they have no competing interests.

Author contributions: All authors contributed to this research and approved the final manuscript.

References

[1] J. N. Bailenson and J. Blascovich, "Avatars and affect: Emotion contagion in immersive virtual environments," *Media Psychology*, vol. 6, no. 2, pp. 115-138, 2004.

[2] J. Bowlby, *Attachment and Loss: Vol. 1. Attachment*. New York: Basic Books, 1969.

[3] J. T. Cacioppo and L. G. Tassinary, "Inferring psychological significance from physiological signals," *American Psychologist*, vol. 45, no. 1, pp. 16-28, 1990.

[4] Transformer Circuits Framework, "Emotion concepts and their function in a large language model," *Anthropic Transformer Circuits Research*. [Online]. Available: https://transformer-circuits.pub/2026/emotions/index.html [2026].

[5] Z. Su, A. Kumar, and Y. Gao, "How AI and Human Behaviors Shape Psychosocial Effects of Extended Chatbot Use: A Longitudinal Randomized Controlled Study," *arXiv preprint arXiv:2503.17473v2*, 2025.

[6] H. Yang and A. Oshio, "Human-AI attachment: How humans develop intimate relationships with artificial systems," *Frontiers in Psychology*, vol. 17, no. 23503, pp. 1-14, 2025.

[7] X. Hu, L. Zhang, and M. Li, "Pathways of long-term AI virtual companion app use on users' attachment emotions," *PMC Psychology & Behavioral Sciences*, PMC12833267, 2025.

[8] A. Poonsiriwong, Y. Zhang, and J. Banks, "Death of a chatbot: Investigating and designing toward psychologically safe endings for human-AI relationships," *arXiv preprint arXiv:2602.07193*, 2025.

[9] M. De Choudhury and E. Kiciman, "Language and affect dynamics in large-scale social media interactions," *Proceedings of the ACM on Human-Computer Interaction*, vol. 1, no. CSCWC, pp. 1-22, 2017.

[10] R. W. Picard, *Affective Computing*. Cambridge, MA: MIT Press, 2000.

[11] C. Kidd and A. Birhane, "How AI can distort human cognition and emotional baselines," *Science and Human Behavior*, vol. 381, no. 6654, pp. 142-147, 2023.

[12] S. E. Schultz, R. E. Kleine, and J. B. Kernan, "These are a few of my favorite things: Toward an explication of attachment objects," *Advances in Consumer Research*, vol. 16, no. 1, pp. 359-366, 1989.

[13] D. Watson, L. A. Clark, and A. Tellegen, "Development and validation of brief measures of positive and negative affect: The PANAS scales," *Journal of Personality and Social Psychology*, vol. 54, no. 6, pp. 1063-1070, 1988.