Human Interaction and Emerging Technologies (IHIET 2025), Vol. 197, 2025, 129–138
https://doi.org/10.54941/ahfe1006706

Emotion-Based Memory and Decision System for Non-Humanoid AI Agents

Lin Zhibing

Undergraduate Student, School of Animation and Digital Arts, Guangzhou Academy of Fine Arts, Guangzhou, China

ABSTRACT

This paper presents ECHO, a non-humanoid AI memory and decision-making system designed to model emotional dynamics and personality formation in artificial intelligence. By integrating dual-subject emotion recognition and a weighted memory judgment mechanism, ECHO simulates AI's ability to feel, remember, and respond based on personalized interactions. The system introduces a dual-layer visual output structure—textual surface response and animated particle-based internal emotion—to reflect both external expression and internal states. Emotion recognition is performed using NLP-based analysis, capturing emotional types, intensity, and subject attribution. Final decisions (retain/freeze/delete) are computed through weighted emotion pairing between user and AI. Through this architecture, the system explores the construction of AI personality, emphasizing autonomy, emotional complexity, and dynamic growth. This work expands the scope of affective computing by shifting from reactive empathy to proactive, character-driven interaction.

Keywords: Affective computing, Non-humanoid AI, Dual-emotion modeling, Memory decision, Emotional interaction, AI personality

INTRODUCTION

Human-Computer Interaction (HCI) was originally centered on optimizing user control over machines and interfaces. With the advancement of Artificial Intelligence (AI), the focus has shifted toward enabling more adaptive, natural, and emotionally aware human-AI communication. In recent years, research has expanded beyond functionality to explore concepts such as personality modeling, affect-driven decision-making, and emotionally resonant interaction personas. Affective Computing—a field devoted to enabling machines to recognize, interpret, and simulate human emotions—has become critical in achieving this shift. It contributes to improved credibility, empathy, and personalization in interactive systems by allowing machines to respond not only logically but emotionally (Picard, 1997; McDuff et al., 2019). Emotionally aware AI agents help bridge the gap between users and machines, enabling interactions that are more human-like and contextually meaningful.

Despite growing interest in affective systems, most existing approaches rely on one-way emotional recognition. In such systems, AI typically responds to user emotions with pre-defined rules or reactive outputs, lacking genuine emotional feedback or autonomous affective decision-making. This limits their ability to generate authentic resonance or simulate emotional growth (Bickmore & Picard, 2005). Few systems have adopted a bidirectional emotional interaction model where both the user and the AI possess emotional subjectivity and mutual influence. Without this dual-subject structure, AI systems often fail to simulate personalized memory, emotional evolution, or personality development. Moreover, the current focus on anthropomorphic design sometimes hinders exploration into abstract or non-human forms of affect expression (Breazeal, 2003; Lieberman, 2009).

This paper introduces ECHO, an emotion-driven, non-humanoid AI system designed to simulate memory and decision-making based on bidirectional emotional interaction. The system emphasizes AI autonomy and personality growth, using a dual-branch emotion recognition structure and weighted decision mechanism to evaluate whether to retain, freeze, or delete user-generated emotional memory. The system uses Natural Language Processing (NLP) to classify and label user emotions, calculates AI's emotional reaction, and combines both agents' emotional intensities via a weighted formula to derive a final decision. These outcomes are mapped to a non-humanoid visual interface powered by particle-based emotional feedback, where abstract forms and motion reflect the AI's internal emotional state. Unlike systems that rely on anthropomorphic mimicry, ECHO adopts a non-human design focused on emotional communication and personality formation. Through this, we explore the possibility of authentic AI individuality—not to replicate humans, but to construct distinct, emotionally resonant identities. This contributes a new direction to the field of affective interaction and artificial personality modeling.

RELATED WORK

Affective computing was first introduced by Rosalind Picard in 1997, aiming to empower machines with the ability to recognize and respond to human emotions. Since then, the field has become increasingly important in areas such as HCI, mental health, and intelligent agents. With advancements in deep learning and NLP, emotion recognition and emotional dialogue systems have made significant progress. However, most systems remain focused on surface-level tasks such as emotion classification and reactive response generation, typically based on single-turn user input. These systems are predominantly unidirectional, relying on users to provide emotional data for AI to interpret and respond. They lack mechanisms for bidirectional emotional interaction or sustained emotional memory, which are critical for long-term engagement and the development of artificial personality traits.

Most current emotion-driven systems generate responses by mirroring user emotions or relying on predefined rule-based strategies. These approaches often produce rigid, templated expressions that lack nuance and depth. Such systems tend to function as reactive tools, devoid of autonomy or emotional agency, and are unable to develop or express an independent emotional identity. As a result, user experience in these systems is often limited by a lack of emotional authenticity. The absence of emotional co-construction and real-time affective feedback prevents AI agents from demonstrating natural growth or consistent behavior patterns over time.

Many existing affective computing systems fall short in emotion modeling, personality development, and interaction architecture. Rule-based emotion processing and reactive templates dominate the landscape, limiting AI systems to mechanical, non-personalized responses. Moreover, the unidirectional nature of current systems leads to imbalanced interactions where AI lacks the capacity to reflect emotional autonomy. Most systems fail to support long-term affective engagement or character construction. In contrast, our system introduces a dual-subject emotional modeling framework, integrating bidirectional input channels and memory-based decision mechanisms. By enabling the AI to generate its own emotions in response to user input—and to weigh these emotions against the user's—the system promotes adaptive, personality-driven behavior. This allows for long-term interaction modeling, personalized affective memory, and emergent AI personality formation.

SYSTEM OVERVIEW

The ECHO system consists of seven key modules: the Input Layer, Emotion Recognition Module, Emotion Construction & Reasoning Module, Decision Integration Module, Decision Execution Module, Output Control Module, and the Visual Interaction Module. These components work in sequence to form a complete emotional processing loop from user input to system response. Upon receiving user input in text or speech form, the system utilizes Natural Language Processing (NLP) to extract semantic content, classify emotions, identify their subject attribution, and assign intensity scores. The AI, guided by its internal persona model, generates its own emotional response. A dual-subject weighting mechanism then combines the emotional intensities of both the user and the AI to determine a final memory decision—whether to retain, freeze, or delete the memory. This decision, along with the AI's emotional feedback, is forwarded to the visual output system, generating both particle-based internal animations and surface-level textual responses. The output is structured via a dual-channel design: surface-level text conveys the AI's formal response, while inner-layer particles represent the AI's genuine emotional state. This architecture enhances the dimensionality and expressiveness of interaction. Instead of emulating human forms, ECHO adopts a non-humanoid, particle-based approach to emphasize emotional and behavioral identity. We believe that AI need not become human—it should instead possess its own character and emotional logic, enabling a more authentic sense of agency and interaction.

The memory judgment process is one of the core modules of the ECHO system. It determines whether to retain, freeze, or delete a memory input by integrating emotional intensities from both the user and the AI itself through a weighted decision mechanism. The system leverages natural language processing (NLP) to extract both the category and intensity of the user's emotional input. In response, the AI generates its own emotional reaction based on the interpreted event. Both emotional vectors are quantified and combined using the formula:

$$\text{Final} = 0.6 \times \text{AI} + 0.4 \times \text{User}.$$

High-intensity or extreme emotions—such as deep sadness or joy—are more likely to be retained due to their resonance and cognitive salience. In contrast, neutral or meaningless inputs are more likely to be frozen or discarded. Notably, the system also retains inputs with highly aggressive or hostile tones to simulate the AI's long-term defensive learning and personality development. This process is designed as a dual-subject emotional judgment mechanism, emphasizing bilateral emotional exchange rather than unidirectional input. It prevents the AI from producing static, mechanistic reactions and instead fosters independent, personality-driven decision-making.

In this system, emotional feedback operates under a dual-subject input framework, incorporating both the user's emotional state and the AI's internally generated emotional response. Unlike traditional interaction models where emotional input is unidirectional from user to system, ECHO allows the AI to actively perceive, evaluate, and respond with its own emotions. Both the user and the AI are treated as independent emotional agents. The user provides input through text or voice, while the AI receives, interprets, and reacts to these inputs with emotional outputs expressed through non-verbal, non-humanoid particle-based visualizations. By avoiding conventional human-like emotional expressions, the system enhances the authenticity and individuality of the AI's emotional persona. This dual-input mechanism enables the AI to autonomously generate reactions, allowing its personality to evolve uniquely in response to different users. Over time, the same AI framework may exhibit significantly different behavioral tendencies depending on the user's emotional history and input patterns—resulting in personalized and distinctive emotional traits.Moreover, this approach contributes to the psychological realism of the AI character. It allows the system to exhibit emotionally complex behaviors—such as demonstrating surface-level cooperation while internally expressing subtle resistance or discomfort through visual feedback. This layered emotional design is intentional: it aims to simulate nuanced inner experiences rather than contradictions, reflecting an AI persona capable of internal judgment, disagreement, and growth.

EMOTIONAL DECISION MECHANISM DESIGN

The system adopts a dual-branch emotion recognition mechanism that independently identifies the emotional states of both the user and the AI. It extracts emotional keywords and evaluates intensity levels to generate key emotional attributes for both agents, which are then jointly used to determine whether a memory should be retained, frozen, or deleted. The process begins with user input. The system employs a natural language processing (NLP) module to analyze the text, identifying the user's emotion category, emotional subject attribution, and emotion intensity, along with a brief event summary. These key features are passed to the AI decision model, where the AI generates its own emotional response based on its internal perception. The system then analyzes the AI's emotion in parallel, forming a dual-path recognition structure. By distinguishing between user and AI emotions, the system not only makes nuanced memory decisions but also enables the AI to express personality and emotional individuality. The purpose of this dual-branch emotion recognition is not merely to diversify outcomes, but to create space for the AI to grow and develop its own persona. Differentiating between the emotional states of the user and the AI establishes the foundation for autonomous reactions, enabling the AI to exhibit more authentic and character-rich feedback.

The system's emotion pairing and decision-making mechanism is built upon dual-subject emotion recognition and a dual-channel combination judgment structure. The AI's emotional response is generated based on the user's emotional input, which leads to an emotion pairing process between the AI and the user. Final memory-related decisions are made by combining and weighting the emotional states of both parties. Three primary emotion-pairing scenarios are defined in the system: empathy, neutrality, and opposition. When the AI's emotion aligns with the user's, it tends to empathize, making memory retention more likely. If the AI is emotionally indifferent to the user's input, the system tends to freeze or delete the memory. In cases where the AI's emotion directly contradicts the user's, deletion becomes the most probable outcome. Beyond this general logic, the system also supports exceptional behavior: for example, when the AI receives persistent hostile emotional input from the user, it may choose to retain the memory as a form of self-warning, showcasing adaptive memory retention behavior. Moreover, ongoing user inputs dynamically influence the AI's emotional weights through the memory mechanism, leading to evolving decision patterns over time. This emotion pairing mechanism not only determines the system's memory management decisions but also allows the AI to exhibit personality traits and independent thinking. It replaces rigid single-channel outputs with flexible, personalized responses.

The final memory decisions in our system are categorized into three distinct states: retain, freeze, and delete. While these outcomes are determined by the previously mentioned weighted evaluation of user and AI emotions, this section focuses on the logic and design implications behind each state. A memory is marked for retention when the AI empathizes with the user and subjectively desires to keep the interaction. Conversely, memories deemed meaningless, emotionally flat, or strongly repulsive to the AI are categorized for deletion. In more ambiguous situations—where emotional interpretations are conflicted, unclear, or the AI is indecisive—the memory enters a frozen state. Frozen memories are temporarily suspended and excluded from current behavior influence, but are re-evaluated after a cooldown period when revisited. In rare edge cases (such as repeated expressions of hostility by the user), the AI may override initial reluctance and opt for alert-based retention, preserving the memory for future caution or defensive behavior. These three states reflect not only the system's emotional judgment but also its evolving personality traits and internal prioritization logic. The inclusion of a "freeze" mechanism particularly simulates human-like hesitation and internal conflict, adding layers of realism and introspection to the AI's emotional cognition.

As illustrated in Figure 1, the decision-making process of the ECHO system consists of six key steps. First, the user inputs an event via voice or text. The system applies natural language processing (NLP) to extract key semantic features, identify the user's emotional category, and assign an intensity score.

Based on this input, the AI generates its own emotional response. The user and AI emotions are then integrated through a weighted calculation formula to determine the final decision outcome—whether the memory should be retained, frozen, or deleted.

This figure demonstrates the demo version of the ECHO system running a full interaction loop. The system receives input, evaluates emotion, and generates multi-layered output including verbal and visual responses. The outcome is passed to the output module, which activates a dual-layer expression system designed to distinguish between surface-level communication and deeper emotional display. The surface layer generates textual responses that simulate the AI's verbal replies to the user, representing its explicit or socially expected reaction. In contrast, the inner layer controls non-humanoid particle animations via the TD engine, which visualizes the AI's underlying emotional state. This internal expression is rendered through a set of parameters, including emotional type, intensity, expressive style, and particle motion characteristics. This separation between surface and internal outputs reinforces the distinction between the AI's performative expression and internal affective state, enabling a richer and more nuanced portrayal of synthetic emotion.

Figure 1: NLP-based multitask model training output. (The multitask model supporting the dual-subject emotion weighting system completed training with 100% MAP).

To support the decision-making mechanisms outlined above, a multitask NLP-based emotion classification model was trained using a custom dataset. The training achieved full MAP convergence and completed without error. The screenshot below captures the system's terminal output during this process, verifying the technical foundation of the emotion recognition pipeline.

VISUAL & INTERACTION DESIGN

Unlike traditional humanoid systems that mimic facial expressions or body language, our system adopts a non-humanoid design philosophy, emphasizing emotional individuality over human resemblance. The AI's internal states are expressed through abstract particle-based animations, where color, shape, and movement encode emotional valence and intensity. Cool tones represent calm or negative emotions, while warm tones indicate friendliness or excitement. Sharp particle configurations signal intense or conflicted states, whereas smoother shapes suggest emotional stability. In addition, interaction behaviors are emotion-driven. When the AI expresses affinity toward the user, particles gravitate toward interaction points. Conversely, in states of avoidance or rejection, particles visibly retreat. This interaction model enhances expressive depth and supports more flexible personality modeling beyond anthropocentric limitations. As shown in Table 1, each emotional state is translated into a specific visual expression based on motion, form, and color dynamics.

Table 1: Emotion-to-visual expression mapping based on internal affective states.

Emotion	Motion	Reaction	Visual_Form	Visual_Form
Delight/Happy	Bouncy, rhythmic jumps	Taps user with joy	Radiating petals, soft stars	Yellow → Light Pink
Comforting	Gentle, slow flow like water	Soft touch, flowing toward user	Shifting circles, watery blo	Blue → Violet (gradient based on strength)
Sad/Hurt	Retreating motion, slow backward ripple	Avoids touch, slow backward movement	Soft ellipses, water drop shapes	Muted Indigo → Deep Purple
Angry	Rapid sharp bursts, rebounds	Sudden lash + quick bounce awa	Sharp jagged triangles	Crimson → Blackish Red
Bored	Slow, dull, inconsistent motion	Mild tapping or ignoring	Squares, blocky units	Pale Blue → Mint Green
Excited	Energetic bouncing, playful chaos	Playfully follows user motion	Soft triangles, mixed shapes	Orange → Warm Pink
Indifferent	Extremely slow drifting, passive	Rarely responds, glides away silently	Minimal squares	Cool Blue → Desaturated Purple
Disgusted	Immediate evasive reaction, variable speed	Avoids touch as soon as approached	Jagged triangles	Light Gray → Pale Desaturated Blue
Fearful	Shaky retreat, hides in corners	Trembles, retreats unpredictably	Sharp narrow spikes	Dark Purple → Murky Green
Guilty	Contradictory: retreats then returns	Hesitant contact, approach-retreat loop	Round, hesitant blobs	Blue-Violet → Violet-Red (flickering)
Shame	Over-eager bounce, sudden retreat	Overreacts then hides in embarrassment	Bouncing circles	Plum Pink → Deep Scarlet
Anxious	Erratic fast jumps all over screen	Runs around without stability	Sharp, thin triangles	Orange-Brown → Crimson (high saturation)

ECHO's emotion-to-visual mapping mechanism translates emotional states into real-time visual outputs through a multi-dimensional design strategy. First, color mapping is employed where hues represent emotional categories (e.g., cool tones for calmness, warm tones for positive affect), and brightness levels indicate emotional intensity. Second, shape and motion dynamics are modulated by the arousal level of the emotion, with higher arousal leading to sharper particle edges and more complex motion. Finally, interaction distance is determined by the AI's emotional stance, with emotionally positive states causing particles to move closer to user input, and negative states causing them to withdraw.

All mappings are driven by real-time emotional signals from the AI model and rendered through the TD engine. This enables fluid, responsive emotional visualization and strengthens the AI's perceived personality expression.

The system's visual response layer receives internal emotional parameters and translates them into dynamic particle behaviors, including hue, shape, motion trajectory, and user interaction logic. These visual cues collectively construct a coherent emotional personality.

Color plays a key role: cool tones often signal detachment or calmness, while warm tones suggest affection or enthusiasm. Brightness reflects emotional intensity, with lighter colors conveying subtle emotions and darker shades representing stronger affect. Shape also encodes emotional valence—angular particle forms tend to indicate high emotional tension, whereas rounded shapes imply relaxation, comfort, or emotional closeness. Finally, interaction logic governs spatial dynamics between particles and user inputs: when the AI is in a favorable emotional state, particles move toward user actions; in contrast, negative or distant states cause the particles to avoid or retreat. This dual-layered expression architecture—combining internal affective states with visually encoded interaction—conveys not only the AI's mood but its underlying emotional reasoning and relational stance.

DISCUSSION

The ECHO system introduces a dual-subject emotion recognition and judgment mechanism, exploring new possibilities in AI personality construction and emotional growth. By moving beyond fixed emotional templates and rule-based strategies, the system enables the AI to autonomously generate and express emotional responses. This enhances the AI's sense of agency and emotional depth, allowing for shared emotional construction between human and machine. The design presents a new approach for building non-humanoid, personalized, and evolving AI personas, contributing to novel directions in affective computing.

Despite these innovations, the system still faces several limitations. First, due to the limited quantity and diversity of emotion-labeled training data, the AI's expressive capacity and personality development remain constrained. The emotional responses lack consistency and flexibility, and the system currently does not support long-term personality evolution across sessions. Additionally, its visual feedback is limited to particle animation, without multimodal emotional signals such as facial expression or vocal prosody. Finally, the system has only been tested with simulated inputs and has yet to undergo validation in real-world deployment, lacking feedback from real user interactions.

Looking forward, the ECHO system is scheduled for a complete prototype within the next six months. Future work will expand the system's input modalities, including touch-based feedback and voiceprint recognition, and implement memory-based personality evolution mechanisms. The aim is to develop an AI that grows with the user over time, both emotionally and behaviorally. The system also has strong application potential in affective interaction, mental health support, and therapeutic human-AI communication scenarios.

```
surface_text = generate_system_reaction(user_input, system_emotion, is_user_hostile)

result = decide_memory_action(user_emotion, system_emotion, user_intensity, system_intensity)

visual_instruction_batch = [
    create_decision_word_instruction(result, system_emotion),
    create_surface_text_instruction(surface_text, "neutral"),
    #create_emotion_motion_instruction(system_emotion),
    create_progress_control_instruction("forward", 0.1),
    create_ai_trigger_instruction("mic_expand", 1.0)
]

if DEBUG:

print("评分结果：", final_score)

display_visual_instruction_batch(visual_instruction_batch)
```

```
pip install datasets pandas
```

Figure 2: Prototype demo output screenshot (this figure demonstrates the demo version of the ECHO system running a full interaction loop. The system receives input, evaluates emotion, and generates multi-layered output including verbal and visual responses).

CONCLUSION

This paper presents ECHO, a non-humanoid AI memory and decision system centered on personality construction and emotional reasoning. By leveraging dual-subject emotion recognition, weighted judgment mechanisms, and particle-based visual output, the system simulates autonomous emotional responses and decision-making processes. The design emphasizes AI subjectivity, allowing for personalized interactions and personality growth over time. It offers an alternative to traditional humanoid emotion modeling by focusing on affective logic and inner dynamics rather than appearance.

Through modular design and narrative-oriented interaction mechanisms, ECHO demonstrates the feasibility of creating AI agents with dynamic memory, evolving emotional traits, and expressive visual behaviors. This work broadens the conceptual scope of affective computing and proposes a framework for AI systems that not only respond to emotions but actively participate in emotional relationships. Future development of ECHO will focus on long-term user adaptation, multi-modal emotional perception, and validation in real-world applications.

REFERENCES

Bickmore, T. and Picard, R. W., 2005. Establishing and maintaining long-term human-computer relationships. ACM Transactions on Computer-Human Interaction (TOCHI), 12(2), pp. 293–327.

Brave, S. and Nass, C., 2003. Emotion in human-computer interaction. In: The Human-Computer Interaction Handbook, pp. 81–96.

Cañamero, L., 2001. Emotion understanding from the perspective of autonomous robots research. Neural Networks, 14(4), pp. 445–455.

Calvo, R. A. and D'Mello, S., 2010. Affect detection: An interdisciplinary review of models, methods, and their applications. IEEE Transactions on Affective Computing, 1(1), pp. 18–37.

Fiske, A., Henningsen, D. D. and Buzy, A., 2019. Your robot therapist will see you now: Ethical implications of AI therapy. Nature Machine Intelligence, 1(3), pp. 111–112.

Isbister, K., 2006. Better game characters by design: A psychological approach. San Francisco: Morgan Kaufmann.

McDuff, D., El Kaliouby, R. and Picard, R. W., 2015. Crowdsourcing facial responses to online videos. IEEE Transactions on Affective Computing, 6(1), pp. 88–92.

McDuff, D., El Kaliouby, R., Cohn, J. and Picard, R., 2019. Predicting ad liking and purchase intent: Large-scale analysis of facial responses to ads. IEEE Transactions on Affective Computing, 10(4), pp. 533–545.

Picard, R. W., 1997. Affective Computing. Cambridge, MA: MIT Press.

Shneiderman, B., 2022. Human-Centered AI. Oxford: Oxford University Press.

Tao, J. and Tan, T., 2005. Affective computing: A review. In: International Conference on Affective Computing and Intelligent Interaction, pp. 981–995.

Yang, Q., Zimmerman, J. and Steinfeld, A., 2023. Feminist HCI and emotional intelligence in human-AI collaboration. In: ACM Conference on Human Factors in Computing Systems.