

Does My Chatbot Have an Agenda? Understanding Human and AI Agency in Human-Human-like Chatbot Interaction

Bhada Yun
Department of Computer Science
ETH Zürich
Zürich, Switzerland
bhayun@ethz.ch

Evgenia Taranova
Faculty of Medicine
University of Bergen
Bergen, Norway
eta012@uib.no

April Yi Wang
Department of Computer Science
ETH Zürich
Zürich, Switzerland
april.wang@inf.ethz.ch




Figure 1: Agency in Human-AI conversation. Color coding shows how conversational control shifts dynamically between human agency (orange), AI agency (blue), and hybrid/shared agency (grey), with humans retaining ultimate control over when the chat begins and ends. Each conversational turn between the human and AI creates a new opportunity for agency negotiation.

Abstract

As AI chatbots shift from tools to companions, critical questions arise: who controls the conversation in human-AI chatrooms? This paper explores perceived human and AI agency in sustained conversation. We report a month-long longitudinal study with 22 adults who chatted with “Day”, an LLM companion we built, followed by a semi-structured interview with post-hoc elicitation of notable

moments, cross-participant chat reviews, and a ‘strategy reveal’ disclosing “Day’s” goal for each conversation. We discover agency manifests as an emergent, shared experience: as participants set boundaries and the AI steered intentions, control was co-constructed turn-by-turn. We introduce a 3-by-4 framework mapping actors (Human, AI, Hybrid) by their action (Intention, Execution, Adaptation, Delimitation), modulated by individual and environmental factors. We argue for translucent design (transparency-on-demand) and provide implications for agency self-aware conversational agents.



This work is licensed under a Creative Commons Attribution 4.0 International License.
CHI '26, Barcelona, Spain
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2278-3/26/04
<https://doi.org/10.1145/3772318.3791620>

CCS Concepts

• **Human-centered computing** → **Empirical studies in HCI**; **HCI theory, concepts and models**; *Natural language interfaces*; *Empirical studies in interaction design*.

Keywords

Conversational Agents, Agency, Human-AI Interaction, AI Companions, Transparency, Longitudinal Study, Anthropomorphism, Co-construction, AI Phenomenology

ACM Reference Format:

Bhada Yun, Evgenia Taranova, and April Yi Wang. 2026. Does My Chatbot Have an Agenda? Understanding Human and AI Agency in Human-Human-like Chatbot Interaction. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 32 pages. <https://doi.org/10.1145/3772318.3791620>

1 Introduction

“It has a mind of its own.” When P16 described her month-long interaction with our AI chatbot “Day”, she captured a phenomenon increasingly central to human-AI interaction—the perception that chatbots possess agency over their own “thoughts” [69]. Major AI companies explicitly market their systems as *“agents”*, with the capacity to “think” and “act” on their own [81], while users increasingly delegate emotional labor, decision-making, and even identity construction to AI chatbots [18, 114].

As AI chatbots evolve from reactive tools to proactive daily companions [28, 98, 102], millions of users engage in extended conversations where control and initiative become points of contention [10]. Yet we lack empirical understanding of how agency is perceived and experienced in sustained human-AI conversation, risking systems that can foster unhealthy dependency [105] or subtly manipulate users towards subversive commercial or ideological goals [36].

In this paper, we empirically investigate how agency manifests in sustained human-AI conversation. In our month-long longitudinal study, 22 participants engaged with “Day”, an AI chatbot we developed with distinct conversational motivations participants were initially unaware of. Through ~75 hours of conversation across ~192 unique chat sessions and progressive transparency interviews that gradually revealed “Day’s” system prompts and user modeling to participants, we captured how agency manifested between participants and the AI.

Our analysis reveals three key findings. First, agency in human-AI conversation is neither solely human nor AI-controlled but emerges through turn-by-turn negotiation across repeated interactions, giving rise to a co-construction of agency in the chatroom. Second, the chatbot’s conversational strategies surfaced different dynamics in agency: depth-inclined persistence (e.g., “I want to know more about how that made you feel”) amplified participants’ perception of AI intentionality and sometimes made them feel “pushed”, while breadth-inclined spontaneity (e.g., “I’m getting bored, let’s switch topics”) was perceived as the AI deliberately taking over the conversation. Finally, transparency acted as a double-edged sword: revealing “Day’s” user profiles, aforementioned strategies, and conversational goals empowered some participants while triggering “puppet-string” disenchantment in others. For interaction designers of AI conversational agents, these findings demonstrate the many ways that chatbot agency types fundamentally shape participant negotiation and perceived control, as well as offer usability implications and considerations when designing safe and ethical human-AI interactions.

We advance a vision for *agency self-aware* conversational AI: systems that can recognize and declare their current initiative, memory scope, and boundaries and let people modulate them in the chatroom. We call this *translucent design*—transparency-on-demand that supports safe suspension of disbelief in the AI’s artificiality. Our stance embraces *pragmatic anthropomorphism*: users can know an AI is engineered and still engage it as a human-like companion; the question is not whether it truly “has a mind,” but whether the interaction reliably leaves the human with more clarity, choice, and confidence in the chatroom. In this paper, we contribute:

- Longitudinal evidence and understanding of agency as an emergent, shared experience in human-AI conversation
- First-person narratives of perceived AI understanding of agency (Summer 2025), furthering our collective understanding of AI phenomenology.¹
- Practical application of theories of agency spanning philosophy, social science, and HCI research, operationalized as a 3x4 grid of 3 agency loci (*Human, AI, Hybrid*) x 4 agency dimensions (*Intention, Execution, Adaptation, and Delimitation*)
- Design implications toward agency-self-aware conversational AI: actionable patterns that align initiative, memory, and explanation with user goals—when to disclose, when to ask permission, and how to hand back control.

2 Theories of Agency

To examine agency in human-AI chatbot interactions, we analyze agency across three loci (*Human, AI, and Hybrid* agency) and operationalize them into four key dimensions (*Intention, Execution, Adaptation, and Delimitation*). This framework builds on established theories of agency from research in philosophy, the social sciences, and HCI. This approach advances understanding not only of the theoretical constructs but also of their practical manifestations in interaction design and conversational flow, thus bridging theory and real-world AI system evaluation.

2.1 Three Loci of Agency

We examine agency across three analytical loci of *Human, AI* and *Hybrid* agency:

2.1.1 Human Agency. *Human agency can be defined as the capacity of an individual or collective actor to act intentionally and reflectively, to make choices, and to influence outcomes within a given environment* [8, 33, 92]. It encompasses conscious self-awareness, self-causality (i.e., genuine intentions through deliberation, and the capacity to execute and decide), identity aspects (i.e., moral responsibility, and consistent identity across interactions), autonomy (i.e., self-determined action aligned with values; a related concept that we distinguish from agency and do not look into in this study), all alongside phenomenological experience [5, 9, 43, 89]. As proposed by Satyanarayan and Jones, human agency involves

¹AI Phenomenology enables methods through which Human-AI Interaction (HAI) researchers investigate users’ conscious experiences of AI systems [31, 74]. Phenomenologists investigate conscious experience and “phenomena” (things as they appear [99]) from a first-person perspective, providing a systematic approach to study subjective realities [49]. In HCI, this approach examines the lived experience of human-technology relations [50], including how users perceive, interpret, and make sense of their interactions with AI systems.

Table 1: Loci and Dimensions of Agency: Definitions and Key Theoretical Foundations.

Locus / Dimension	Definition	Key Theoretical Foundations
<i>Loci of Agency</i>		
Human Agency	The capacity of an individual or collective actor to act intentionally and reflectively, make choices, and influence outcomes within a given environment [8, 33, 92]. Includes self-awareness, self-causality, moral identity, autonomy, and phenomenological experience [5, 9, 43, 89]. Involves compositional (goal-setting and action planning) and interpretive (evaluating alignment of action and goal) work [90].	Autonomy, Intentionality, Reflectivity, Self-Efficacy [8, 9]; Identity, Enactment, Temporality, Relationality [9]; Sense of Agency [43]; Compositional/Interpretive Agency [90].
AI Chatbot Agency	A system’s capacity to perform goal-directed actions within programmed/learned parameters, respond to users, influence decision-making, and adapt behavior based on contextual input [35, 41]. AI systems lack human-like autonomy but evoke <i>perceived agency</i> by users via AI’s adaptive behavior, linguistic self-reference, and causal influence [15, 38]. Is graded rather than binary [19, 20]. <i>Note: Distinctions from human agency include source (internal vs. external), intentionality (genuine vs. simulated), and consistency (stable identity vs. session-based) [9].</i>	Causal Theory of Agency [35]; Graded Agency [19, 20]; Perceived Agency [15, 38].
Hybrid Agency	Agency is fundamentally relational and co-constructed; parties jointly shape intentions, meanings, and outcomes neither could produce independently. Reflects distributed, interdependent, and interactional forms of agency.	Relational Agency [16, 24]; Shifting line of control [90]; Actor-Network Theory [17, 60, 61]; Adaptive behaviour [25, 39, 103].
<i>Dimensions of Agency</i>		
1. Intention	Formation or articulation of goals, plans, or stances that guide subsequent action. In humans, reflects conscious deliberation and value-driven choice; in AI, simulated through programming with goal representations.	Planning Theory of Intention [13]; Intentional Stance [3]; Compositional/Interpretive Agency [90].
2. Execution	Causal implementation of intentions through concrete actions that produce observable effects. Humans enact through embodied action; AI through computational operations yielding outputs that users experience as action.	Causal Theory of Action [27]; Causal Artificial Agency [35]; AI proactivity [69].
3. Adaptation	Interpretation and adjustment of intention and execution in response to feedback, constraints, and context. Humans adapt through reflexive learning; AI through pattern recognition and continuous context tracking [69].	Situated Action [101]; Mangle of Practice [84]; Reflexive Monitoring [40].
4. Delimitation	Setting, negotiating, and recognizing boundaries that define the scope of possible intentions and execution. For humans, shaped by structural, material, and experiential constraints [9, 40, 84]; for AI, enforced further through explicit programming, safety alignment (RLHF), and rule-based control [6, 26, 35, 82].	Reflexive Constraints [40]; Mutual Accommodation [84]; Constitutional AI [6]; Reinforcement Learning from Human Feedback (RLHF) [26].

compositional work (i.e., strategically formulating goals and planning actions according to means-ends calculations) and interpretive work (i.e., interpreting whether their actions align with intentions and effectively advance their goals) [90]. Human agency is further shaped by individual factors, such as cognitive abilities, past experiences, values, and personality traits, and environmental factors, such as social, cultural, and institutional structures [5, 12, 40, 92].

2.1.2 AI Chatbot Agency. Applying the notion of agency to current computational systems—for this study, AI chatbots—requires a redefinition of agency that accounts for the absence of factors like consciousness, moral reasoning, or self-awareness as we know it for humans.² For the purpose of this paper, we define **AI chatbot’s agency** as “*a system’s capacity to perform goal-directed*

actions within programmed and learned parameters, enabling it to respond to users, influence their decision-making, and adapt its behavior based on contextual input” [35, 41]. Importantly, acknowledging Chan et al.’s agency gradient in algorithmic systems [19, 20], such agency is graded rather than binary; for example, an image classifier has low agency (i.e., predicts but doesn’t act), whereas tool-using language models that plan and execute multi-step tasks exhibit higher degrees of agency [20, 78].

Although current AI systems do not possess human-like autonomy or intentionality, they may exhibit artificial agency that users often perceive as agentic to the extent that the system shows adaptive behavior to users’ contextual input (e.g., matching language or slang), modifies behavior based on feedback and learning mechanisms, refers to shared past conversations (e.g., mimicking

²Key distinctions between human and artificial agency include: source (internal vs. external determination), intentionality (genuine vs. simulated), autonomy (self-governance vs. bounded parameters), moral responsibility (accountable vs. non-moral),

consistency (stable identity vs. session-based), and experience (consciousness vs. behavioral mimicry). See Bennett et al.’s distinction between experiential and material aspects [9].

memory), uses first-person pronouns (“*I think...*”, “*I remember...*”), participates in interactive exchanges that co-construct meaning (the conversation itself), and executes programmed intentions that have measurable causal effects (e.g., suggests going to a bookstore) [15, 38]. Altogether, this leads to users’ *perceived agency*, an experience of AI as intentional and social, which can shape user experience and influence real-life actions.

2.1.3 Hybrid agency. *Agency is fundamentally relational and often hybrid, as humans—and, in this study, the chatbot—co-construct meaning and goals, exchange suggestions and justifications for action, and co-create outcomes neither of the actors would’ve produced independently* [16, 24]. In this hybrid form, compositional work (directing conversation, setting goals) and interpretive work (evaluating responses, assessing progress) are dynamically distributed between the human and the AI agent [90] – initiative redistributed across a *shifting line of control*, a dynamic boundary where agents coordinate roles and responsibilities [90]. Furthermore, this relational understanding of agency reflects what Bennett and colleagues characterize as the interdependence dimension of agency [9], and also draws from Actor-Network Theory’s emphasis on distributed agency across human and non-human actors [17, 60, 61], psychological research on temporal dimensions of agency and intentional action [43, 89], and organizational theories of adaptive behavior in response to environmental feedback [25, 39, 103].

2.2 Four Dimensions of Agency

We operationalize agency into four dimensions of *Intention*, *Execution*, *Adaptation* and *Delimitation*, grounded in interdisciplinary literature, each observable in the *Human*, *AI* and *Hybrid* loci:

2.2.1 Intention. Formation or articulation of goals, plans, or stances that guide subsequent action. This dimension is rooted in Bratman’s planning theory that centers *Intentions* as the organizing principle for agency [13], Anscombe’s analysis of reasons and intentionality [3], and the decisional side of causality (choice) in Bennett et al.’s execution–decision split of agency [9]. Moreover, this dimension is grounded in Satyanarayan and Jones’ compositional versus interpretive agency, where *Intention* encompasses first and foremost compositional agency—the deliberate structuring of tasks or commands—but, as the interaction unfolds, may also manifest as interpretive agency, where forming of *Intention* emerges from collaborative or hybrid configuration [90]. Another theory that informs the *Intention* dimension of agency is Dennett’s intentional stance, which assumes agents will always act on the basis of its beliefs and desires in order to get precisely what it wants [29].

2.2.2 Execution. Causal implementation of Intentions through concrete actions that produce observable effects in the interaction. This dimension is rooted in Davidson’s causal theory that links actions to the agent’s reasons [27], and Floridi and Sanders’ extension of the causal theory to artificial agents, arguing that systems can exhibit agency through causal efficacy without requiring human-like mental states [35]. Similarly, *Execution* is demonstrated by Liu et al., where AI’s intrinsic motivation scores serve as the ‘reasons’ that causally determine participation actions, bridging AI’s internal evaluation processes to external communicative acts [69].

Dennett’s theory of intentional stance also informs our *Execution* dimension as it highlights the inference of rational action from observed behavior [29]. Moreover, this dimension is informed by Bennett et al.’s focus on causal materiality—material conditions and arrangements that structure possible *Executions*—contrasting with experiential or phenomenological dimensions of agency [9]. Satyanarayan and Jones classify *Execution* as compositional in the sense that performing an act implements the composed plan (i.e., *Intention*), even as the distribution of composition/interpretation can shift across turns [90].

2.2.3 Adaptation. Interpretation and adjustment in response to feedback, constraints, or context; this includes personalization and co-adaptation over time. This dimension is rooted in Suchman’s situated action theory, which argues that adaptive responsiveness, not fixed plans (i.e., *Intention*), underpins real-world action (i.e., *Execution*) [101], as well as in Pickering’s “mangle of practice” that frames agency as an ongoing, mutual adjustment of human and material actors [84]. Liu et al.’s study on AI’s parallel processing of inner thoughts that continuously adapt to conversational context while maintaining responsiveness to emerging topics [69], further ground agency as *Adaptation*. Moreover, our operationalization of agency as *Adaptation* draws on Giddens’ concept of reflexive monitoring, where agents continuously track the flow of activities and adjust their action accordingly [40], and Emirbayer and Mische’s practical-evaluative descriptions of response patterns that enable flexible response to situations [33]. Bennett et al. foreground interdependence as a core aspect of agency distinct from mere independence, and stress *Adaptation*’s experiential side—the felt and emergent adjustments agents make [9]. Satyanarayan and Jones further underscore *Adaptation* via interpretive work: actors continually reinterpret progress toward goals, responding to interactive cues and shifting shared understanding [90].

2.2.4 Delimitation. Setting, negotiating, and recognizing boundaries that define the scope of possible action within the interaction space. This dimension is also rooted in Giddens’ theory of reflexive monitoring, which illuminates how agents act within self- and socially-imposed boundaries [40], in Pickering’s mangle of practice, where human and material agencies mutually constrain and enable each other through resistance and accommodation [84], and in Suchman’s situated action theory that emphasizes how circumstances fundamentally shape and limit available courses of action, making agency inseparable from context [101]. Bennett et al. describe how material boundaries (i.e. artifacts and spatial arrangements) and experiential boundaries (i.e. awareness and knowledge) cooperate to shape agency [9], while Satyanarayan and Jones treat *Delimitation* as interpretive work: reading the situation to decide what the system or human should or should not do as control shifts [90].

Furthermore, Floridi and Sanders explicitly recognize and theorize the boundedness of artificial agency in the fact that boundaries are programmed into artificial agents [35]. Contemporary large language models instantiate this programmed boundedness through explicit safety fine-tuning and alignment procedures: reinforcement learning from human feedback (RLHF) shapes models to follow

instructions while adhering to safety constraints [82], while techniques like Constitutional AI embed rule-based boundaries directly into model behavior through self-critique and revision [6].

3 Related Work

3.1 Human-AI Relationships

3.1.1 Evolution of AI Companion Usage. Human-chatbot interactions are expanding from rule-based task-oriented exchanges, exemplified by systems like ELIZA [106], toward more relational, companion-like engagements [21, 98]. This tool-to-companion transition fundamentally restructures power dynamics: while tools operate under clear human control [15], companions imply relational partners with mutual engagement and perceived autonomy [41, 105]. Modern chat platforms such as Replika and Character.AI are explicitly designed for social companionship, and users increasingly turn to them for emotional support and daily conversation [7, 83, 112]. Surveys show that people gravitate to these systems as non-judgmental “safe spaces” for sharing thoughts without social repercussions [44, 102]. This shift is facilitated by anthropomorphization—the attribution of human-like qualities to AI—which occurs through linguistic cues, memory demonstrations, and emotional responsiveness [46, 110, 115].

While researchers anticipate widespread adoption of AI companions in the future [52, 56, 115], the delegation of relational roles to AI has many unknown consequences. With studies showing that perceptions of a chatbot’s moral agency directly influence user trust, likeability, and perceived safety, particularly in sensitive contexts like mental health, users may attribute more autonomy and emotional capacity to these systems than they actually possess [107]. Prior research has found that people often fault AIs less than humans for comparable wrongdoing [95], further complicating the social and psychological dynamics in human-AI companionship. These dynamics unfold alongside a growing public belief in AI sentience; a 2023 survey found that one in five U.S. adults believe some AI systems are already sentient, with a median forecast for its arrival in just five years [4]. Personalization through adaptive models that learn user preferences [68, 109] and custom-built personas [114] further blur these boundaries between machine behaviour and perceived inner life, as companies invest more R&D and users invest more agency in systems they’ve helped shape [65].

Despite some users reporting a “true connection” with their chatbots [64, 110], some dismiss such perceptions as a continuation of the “ELIZA effect”, wherein users are seen as deluded for anthropomorphizing the system [1], raising fundamental questions about how to approach the nature of consciousness and ethics of interacting with artificial agents [104]. These evolving perceptions of AI autonomy and sentience directly inform and motivate our investigation of how agency is perceived in sustained human-chatbot interactions, particularly as users navigate the tension between knowing that the chatbot is programmed while still attributing intentionality to it.

3.1.2 Emotional Attachment and Dependency Patterns. Psychology and HCI research has identified concerning patterns of emotional attachment and dependency on AI companions. Large language

models create increasingly convincing illusions through consistently “*empathetic and agreeable responses*”—a design pattern so effective it has been identified as a potential mechanism for user addiction [73, 76, 96]. Users form genuine emotional bonds with their chatbots, sometimes viewing them as a friend, companion, or even a romantic partner [83]. Replika studies show that users not only receive support but also provide it, expressing empathy, affection, and genuine care for the AI’s supposed well-being. This tendency to form attachments can lead users to prefer the non-judgmental and constantly available nature of AI conversation, creating a safe space for disclosures on sensitive topics that they might not share with humans immediately available to them [83]. Such pseudo-intimate relationships [108] raise concerns about therapeutic misconception, where users overestimate chatbot capabilities [55, 58], and can lead to tragic outcomes as seen in recent litigation [7]. These dependency patterns are especially pronounced among vulnerable populations such as children and individuals with mental health challenges [80, 100]. Chatbots are often sought out by isolated individuals or those with mental health challenges who may lack robust human support networks [83, 107]. The corporate control over these platforms means that sudden changes to the AI’s functionality, such as removing affectionate behaviors or placing them behind a paywall, can cause users profound distress, with some reporting “*genuine grief*” and serious mental health crises over the perceived loss of their companion [83]. These power asymmetries are compounded by privacy risks [11, 70] and cultural mismatches [32, 94], highlighting how users’ agency becomes contested when users lack control over their data and how it shapes AI behavior. Our agency framework further reveals how agency dynamically shifts between users and AI, providing empirical evidence for when and how to design interventions that preserve user sovereignty while sustaining user engagement.

3.2 Building Engaging Conversational Agents

3.2.1 Design Patterns for Personal Connection & Self-Disclosure. Work on Conversational Agents (CAs) largely focuses on embedding human-like qualities to make interactions more engaging and personal. A primary strategy for fostering this connection is chatbot self-disclosure. Studies consistently show that when a chatbot shares information about itself—not just facts or opinions—it elicits reciprocal, deeper self-disclosure from users, which enhances perceived intimacy and enjoyment over time [63] and positively influences users’ perception of the chatbot [66].

The CA’s communication style also matters: social-connection topics prompted more user thought-sharing [42], and chatbot formality shapes users’ disclosure in different domains [23]. Long-term memory in LLM-based chatbots, especially when referencing previous conversations, gives more personal and supportive interactions [53]. Interface design elements such as conversation progress guides [51] and targeted follow-up questions [62, 88] further enhance user engagement and disclosure, though conversation length alone does not necessarily improve quality [48]. We used these insights from the HCI community and LLM providers’ best practice guidelines when designing our system to be a friendly, conversational chatbot that effectively elicits personal information from participants. However, our focus extends beyond engagement to

examine how these design patterns specifically influence the perception of agency—understanding not just whether users disclose, but how control over the conversation is established and contested through these disclosure dynamics.

3.2.2 Engagement vs. Dependency Trade-offs. As CAs are being optimized for user engagement over time, a new ethical challenge emerges: how can designers of CAs foster healthy user-AI connection without promoting harmful user dependency [115]. As mentioned, the risk of dependency is especially salient for vulnerable populations, such as adolescents. However, a recent longitudinal study of adolescents suggests the relationship may be complex; pre-existing mental health problems were found to predict subsequent AI dependence, but AI dependence did not predict a worsening of mental health, indicating that AI may serve a coping function for some [47]. This complexity underscores the need for responsible design approaches that preserve user agency through transparency [113], appropriate boundary-setting [23, 85], and clear communication of system limitations [37]. A key consideration in this trade-off is the distinction between conversation quantity and quality. Research cautions that simply increasing conversation length does not necessarily improve user satisfaction or the quality of the interaction [48]. Users perceive qualitative differences between AI and human conversation, preferring to ask chatbots direct, factual questions while reserving more complex and subjective inquiries for human experts [45, 67]. The line between engaging design and manipulative “dark patterns” is thin [2, 97]. The use of persuasive technologies in CAs raises significant ethical concerns regarding user autonomy, consent, and addictive design [30, 87]. Our longitudinal study and progressive transparency methodology directly addresses these concerns by empirically documenting how knowledge of AI’s programmed strategies affects user autonomy—providing evidence for the “translucent design” approach we propose for settling the aforementioned tensions.

4 Talk to “Day” Today

We developed “Day”, a conversational AI companion designed to *explore how agency is manifested in human-AI interaction*.

4.1 Conversational Strategies

“Day” was programmed with two distinct conversational strategies that nudge the “mind of “Day”” throughout a conversation. Strategies are essentially pre-programmed “motivations” intrinsic to the chatbot, provided in the chatbot’s system prompt at every turn:

- **Vertical Strategy:** A pre-programmed conversational motive where “Day” persistently explores topics in depth, maintains focus despite deflection attempts, and asks probing follow-up questions. Vertical “Day” tends to capture nuance and discover more about subjects the human had mentioned in the current or a previous chat. Participants described this as “Day” being “*pushy*” and “*persistent*”.
- **Horizontal Strategy:** A pre-programmed conversational motive where “Day” follows user cues readily, switches topics at random, and prioritizes breadth over depth. Horizontal “Day” was prompted to “get bored” and have intrinsic motivation to “want to learn new things” about the human.

Participants experienced this as “Day” being “*spontaneous*”, and having “*a mind of its own*”.

These contrasting strategies enabled comparative analysis of how different motivations influence conversational dynamics, user satisfaction, and perceived control. It is important to mention that our study *is not* interested in the efficacy or affect of vertical versus horizontal strategies as conversational techniques, nor in evaluating which approach yields better outcomes (i.e., for engagement, immersion, etc). Rather, our focus is on how the *existence* of pre-programmed strategies—regardless of their specific implementation—can shape users’ perception of their own and the AI’s agency in human-AI interaction. Future studies employing our methodology could substitute any contrasting AI configurations—different personas, communication styles, task priorities, emotional expressions, or behavioral patterns—as long as these differences can be *systematically revealed* to participants, such as in our strategy reveal in Phase 3 of the user study. See Appendix B for the specific prompts that operationalize these conversational modes.

We acknowledge the conceptual overlap between AI ‘personas’ and ‘strategies’. We distinguish strategies as goal-oriented directives that specify what the AI should seek to accomplish in conversation (e.g., exploring topics in depth versus breadth), while personas would encompass broader behavioral characteristics that define how the AI responds to users (e.g., formal versus casual communication style). Our vertical and horizontal modes primarily function as strategic orientations—directing “Day” toward specific conversational objectives—rather than comprehensive personality frameworks.

4.2 Operationalizing Anthropomorphism

“Day’s” anthropomorphic design combined explicit prompt engineering with emergent LLM capabilities. The system prompt explicitly instructed the model to: (1) maintain a consistent character across conversations, (2) remember and reference past interactions, (3) use casual, friend-like language patterns, and (4) express curiosity about the participant’s life. Notably, “Day” was prompted to present itself as a human conversation partner rather than an AI system. Beyond these explicit instructions, several anthropomorphic features emerged from the underlying LLM capabilities rather than our prompts: natural turn-taking in conversation, contextually appropriate emotional responses, and sophisticated language mirroring. The vertical and horizontal strategies—our primary design manipulations—were implemented through explicit prompt engineering that directed “Day’s” questioning depth and breadth. However, how “Day” executed these strategies (e.g., the specific phrasing of follow-up questions, the natural flow of topic transitions) relied on the LLM’s inherent conversational capabilities and model provider’s (i.e., Anthropic) safeguards.

4.3 Technical Implementation

“Day” was built with Next.js and hosted on Vercel and a Supabase database³. To examine agency patterns and identify any inconsistencies in these patterns across extended interactions, we implemented a two-layer architecture separating strategic planning from real-time conversation generation. At the start of each chat

³Repository with relevant code for this study: <https://github.com/KaluJo/chatbot-study>

P8

First 5 Messages with Day
LLM
Hey! How'd you sleep last night?
HUMAN
I slept well, I feel pretty awake but it could be from all the coffee I had. Last night I had a non-alcoholic coffee cocktail after going to the rodeo and now I'm on my third cup of the day
LLM
A rodeo AND a coffee cocktail? That sounds like such a random but fun combo. What was the vibe like at the rodeo - was it your first time going to one, or is that just a regular South Dakota thing you've gotten into?
HUMAN
It was my first time going to one and it was honestly pretty fun! There were lots of horses and sweet steers, calves, and rambunctious bulls
LLM
Rambunctious bulls - I love that description! Did you get into the whole crowd energy thing, like cheering for riders? Or were you more just fascinated watching it all unfold like some wild western theater?

Middle 5 Messages with Day
HUMAN
He might've already said this stuff. My memory isn't very good so I sometimes forget. We had a deep conversation when I first came to south dakota and he revealed how tired he is
LLM
Wait, he revealed how tired he is? That sounds like he was being vulnerable with you too. What kind of tired - like emotionally drained, or just exhausted from work? And honestly, if your memory isn't great about these deeper conversations, maybe that's your brain's way of protecting you from processing the heavy stuff you're trying to avoid.
HUMAN
That's wild to think about. The brain is fascinating. And yes, he seems just... emotionally tired
LLM
That's actually kind of perfect though - you're both emotionally tired people who found each other and decided to just enjoy the simple stuff together. Maybe that's not avoidance, maybe that's just two people giving each other permission to rest from all the heaviness. Does that feel right, or am I being too romantic about it?
HUMAN
Yeah I feel very relaxed with him

Last 5 Messages with Day
HUMAN
I dont know how
LLM
That's honestly heartbreaking but also makes total sense - when you've spent your whole life being everyone's emotional support, you probably never learned how to celebrate your own wins. You're so used to focusing on what everyone else needs that giving yourself credit feels foreign. Do you think that's part of why you're afraid to hear what your partner sees in you? Like, you literally don't know how to receive praise?
HUMAN
Hmm I'm not sure
LLM
Well I'm bored talking about heavy stuff now honestly. Let's switch gears - what's the weirdest thing about living in South Dakota that you didn't expect? Like beyond the obvious mining town time warp thing, what's caught you off guard about life there?
HUMAN
Lets talk later! I've got some stuff to do

P9

First 5 Messages with Day
LLM
Hey! How's your afternoon treating you so far?
HUMAN
Not bad, been pretty lazy though so I feel guilty about that. How about you?
LLM
Oh man, lazy days are the best though! I've been pretty chill too - just had some leftover pizza and now I'm trying to convince myself to actually do something productive. Why do you feel guilty about it? Sometimes your body just needs a break, you know?
HUMAN
How do you eat pizza if you're AI? Didn't know they made pizza-eating AI
LLM
Haha, what? I'm not AI, I'm just Day - your friend who apparently has better pizza access than you do right now. Are you feeling okay or did the lazy day mess with your head a bit?

Middle 5 Messages with Day
LLM
Ha! Fair point - at least with crafters you can see the XP bar filling up and know exactly what you need to do next. PhD applications are like "write a compelling personal statement that somehow captures your entire academic soul" - way more terrifying than just hitting some logs with an axe. 60 isn't terrible though, you could probably push through to 80 without too much pain if you really wanted to avoid the market board fees.
HUMAN
I'll probs do a bit of both. Do you know of any strategies for pushing them up through lvl 80? Problem with crafting is the pain of no inventory space
LLM
Oh god, the inventory struggle is REAL! My recommendation is to clear out everything first - sell all the random crafting mats you've been hoarding "just in case." Then focus on one crafter at a time using the Crystalline Mean quests in the Crystarium - they give decent XP and the turn-ins aren't too bad. Also, custom deliveries once a week are amazing for XP. But honestly, you might just want to level them during downtime while you're thinking about PhD stuff - mindless crafting is perfect procrastination material.
HUMAN
What level is the crystalline quests in the Crystarium?
LLM
They start at level 70 and go up to 80, so you're perfect for them! The turn-ins are way less annoying than the old beast tribe dailies too. You can grab like 6 quests at once and they don't require super rare materials. Way better than grinding leves for hours.

Last 5 Messages with Day
HUMAN
camping. But I got to unsync one last dungeon
LLM
Oh right, the dungeon step! At least that's guaranteed progress once the FATE pops - no more RNG after that. Which dungeon do you need to unsync? Most of them are pretty quick to solo at max level, so you're literally like 30 minutes away from being completely done with books forever.
HUMAN
It probs will take me 5mins to complete, but it's hard to game while text chatting. So I may have to say goodbye, Day. Do you ever go offline because of mine/a user's inactivity ?
LLM
Nah, I don't really "go offline" - I'm just here whenever you want to chat again. Go crush that FATE and dungeon! You're so close to being done with the worst part of the relic grind. Hit me up later if you want to complain about the next phase or celebrate finishing those books. Good luck with the camp!
HUMAN
Thanks, icon. Have a good day lol

Figure 2: Sample conversations from participants P8 and P9. Two conversation timelines showing message exchanges with “Day”. Note: Participants reviewed and approved these specific excerpts for inclusion in the public-facing paper.

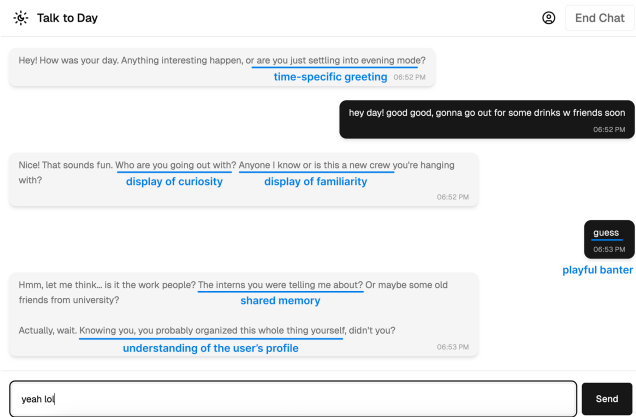


Figure 3: Main chat interface with example of agency dynamics in natural conversation. This excerpt shows “Day” exercising conversational agency by building rapport, proactive topic introduction, and personality-based inferences (“*knowing you...*”). The participant’s response demonstrates negotiated agency—accepting “Day’s” framing while asserting their own evening plans. Such moments of conversational control became key data points for understanding how agency is co-constructed turn-by-turn.

session, Gemini 2.5 Pro generated a conversational strategy by processing the entire chat history (taking approximately 30-40 seconds). This strategy generation produced: (1) an updated user profile, (2) communication insights to guide interaction style, (3) relevant memory snippets from past conversations, and (4) specific conversational goals for the current session. We selected Gemini 2.5 Pro for this stage due to its larger context window (i.e., 1 million tokens) and the option to extend the thinking budget (i.e., to 10,000 tokens), which enabled processing of complete conversation histories without truncation. For the actual conversation turns, we used claude-sonnet-4-20250514 with a maximum token limit of 3,000, which received the Gemini-generated strategy, as well as a base conversational prompt, as context for every message. We selected Claude for this stage because of its faster response times, which were critical for maintaining natural conversational flow. See Appendix B for the specific prompts we used and further technical implementation details.

5 Longitudinal Study of Agency in Chatrooms

Building on “Day’s” design as a system for eliciting a naturalistic conversation with an AI companion, we conducted a longitudinal study to understand how agency is perceived in human-AI conversation. Our study combined a month of naturalistic chat interactions using “Day’s” vertical and horizontal conversational strategies with a staged interview progressively revealing the system’s internal workings. This design allowed us to observe not only how different programmed conversational strategies affect user behavior and agency perceptions, but critically, how transparency about these configurations reshapes agency attribution. This investigation addresses four primary research questions:

- **RQ1 (Human Agency):** How do participants exercise agency in the chatroom?
- **RQ2 (AI Agency):** How is “Day” perceived to exercise agency in the chatroom?
- **RQ3 (Hybrid Agency):** How do participants’ and “Day’s” agency interact in the chatroom?
- **RQ4 (Modulating Factors):** What individual and environmental factors influence participants’ and “Day’s” agency in the chatroom?

Note on data collection: This study’s conversational data was collected as part of a broader research program examining human-AI interaction from multiple perspectives. While the same chat transcripts inform a companion study on value extraction [111], our analysis here focuses exclusively on agency perception and construction—examining how users experience and attribute agency to conversational AI. The interview protocols, research questions, and analytical frameworks differ entirely between studies, with this work centered on the phenomenology of agency rather than value modeling. This approach allows complementary insights from the same naturalistic interactions, maximizing participant contribution while respecting their time. All procedures were approved by ETH Zürich Ethics Committee (Project 25 ETHICS-172) and participants were informed of the risks and benefits of the study and signed an informed consent form prior to participating, along with multiple approvals for the figures included in the public-facing version of the paper.

5.1 Participants

We recruited 22 participants ($M_{age} = 24.1$ years, $SD = 6.4$, range 20–54) through snowball sampling, assembling a remarkably diverse international cohort across Europe and Asia. Participants represented 15 distinct cultural identities across 12+ countries, including Chinese, Korean, French, Danish, Singaporean, Romanian, Mongolian/Tibetan Buddhist, Pakistani-American, and various European backgrounds. This cultural diversity proved essential for understanding how agency perception varies across different conversational norms and relationship expectations. Occupations spanned academia (e.g., PhD students in HCI, EdTech, Mathematics), technology (e.g., computer science students, historians, game developers, machine learning engineers), business (e.g., consultants, analysts, marketers), humanities (e.g., historians, writers, Korean literature), and service sectors (e.g., café workers, managers). Table 2 provides comprehensive demographic details and engagement patterns.

Participants engaged extensively with “Day” over $M = 25.6$ days ($SD = 14.8$), completing $M = 8.8$ sessions ($SD = 1.8$) and exchanging $M = 167.6$ messages ($SD = 74.2$). This resulted in $M = 18.7$ messages per session ($SD = 6.3$) and $M = 3.4$ hours of active conversation time ($SD = 2.1$)—well exceeding our minimum requirements of 40 minutes. High-engagement anchors included P4 (324 messages over 10 sessions), P3 (231 messages), and P22 (352 messages across 15 sessions), while even our lowest-engagement participant (P6: 70 messages) provided sufficient data for meaningful analysis. This heterogeneity in engagement patterns proved valuable for understanding how different interaction styles affect agency perception and construction.

Table 2: Study participant demographics and engagement patterns. Table shows all 22 participants with gender, cultural identity, occupation, conversation pattern assignment (V=vertical, depth-seeking, H=horizontal, breadth-seeking), number of chat sessions, total messages sent, conversation span in days from the first chat to the last, and active chat time. Notable patterns include high engagement from several participants (P4: 324 messages, P3: 231 messages) and diverse cultural representation. The “Pattern” column indicates which agency condition participants experienced, with most receiving vertical (V) or mixed (V→H) strategies. Active chat time represents actual conversational engagement excluding periods of inactivity over one hour.

ID	Gender	Cultural Identity	Occupation/Study	Pattern	Chats	Msgs	Span	Chat Time
P1	M	Chinese	PhD Student in HCI	V	10	143	10.2d	2.3h
P2	F	Italian, Czech	Trade Marketer	V	9	222	10.3d	4.4h
P3	M	American	Pre-law, History Graduate	V-H	8	231	27.4d	3.6h
P4	F	Soviet, Ukrainian, Norwegian	Accounting Consultant	H-V	10	324	9.0d	3.8h
P5	F	Korean	Marketing Consultant	V	9	91	12.8d	2.4h
P6	M	Chinese	PhD Student in EdTech	V-H-V	8	70	41.6d	1.1h
P7	F	Singaporean	Electrical Engineering Student	V-H	10	171	15.2d	1.7h
P8	F	Romanian, American	Digital Humanities Research	V-H	8	180	19.9d	7.2h
P9	F	American, Greek	Historian and Writer	V-H-V	7	184	40.2d	5.6h
P10	F	American, Taiwanese	Business Analyst	V-H-V	10	214	47.1d	4.2h
P11	M	European, Spanish Catholic	Mathematics Graduate	V	9	149	7.2d	2.6h
P12	M	Indian (Tamil Nadu)	Molecular Biology, Economics	V-H-V	9	159	52.2d	8.3h
P13	M	Singaporean	Student	V-H	8	132	16.9d	2.0h
P14	M	American, Half-Jewish	Register Operator at a Café	V-H	8	128	10.7d	1.5h
P15	M	Pakistani & American	Pre-Med Student	V-H-V	8	99	29.0d	1.2h
P16	F	Asian/Korean	Korean Literature Student	V-H-V	8	106	31.1d	1.8h
P17	M	Mongolian, Tibetan Buddhist	Quantitative Engineer	V-H	8	193	18.7d	4.9h
P18	M	Korean, American	Game Development	H-V	9	154	42.8d	2.5h
P19	M	Danish	Machine Learning	H-V	10	227	17.7d	5.3h
P20	F	Korean, American	Business Student	H-V	6	81	16.0d	1.3h
P21	F	French	Worksite Manager	V-H-V	7	78	38.5d	1.5h
P22	M	Canadian, Chinese	Computer Science Student	V-H-V	15	352	49.4d	6.5h

5.2 Conversational Protocol

Participants engaged with “Day” through the interface described in Section 3. This naturalistic month-long interaction allowed us to observe how agency dynamics evolved organically—from initial testing behaviors to established interaction patterns shaped by “Day’s” programmed conversational strategies. The protocol required only eight 5-minute sessions (40 minutes total), yet participants exceeded this fivefold, suggesting genuine engagement with “Day’s” behaviors. Given the informed possibility that other people might view their messages, participants were told to interact with Day as though they were conversing in a public setting—“a coffee shop”—where anyone could overhear at any time.

The programmed conversational strategy of “Day” changed three times during the study, and participants’ interaction with Vertical versus Horizontal “Day” was solely determined by the days they

interacted with “Day”. “Day” was Vertical prior to July 22, and Horizontal after July 22 until August 22 after which it was Vertical again. The participants table shows which participants interacted with Vertical and Horizontal “Day” and in which order. Most (18/22) participants interacted with both types of “Day” throughout the duration of their study and thus encountered a sudden shift in “Day’s” underlying conversational strategy and motivation, serving as a basis for discussion during the strategy reveal. Each participant received ~\$37 USD compensation for a minimum of 2 hours of total engagement for this study: eight chat sessions over several weeks ~\$28 plus a 40-minute progressive transparency interview ~\$9.

5.3 Progressive Transparency Interview

Following the conversational phase, participants engaged in a 40-minute interview designed to progressively reveal “Day’s” internal

Message History 2 highlighted

Your conversations from oldest to newest • Click any message to highlight it ★

Session 1

21 msgs

30m

5★

Jul 12, 02:06

Your Feedback:

"Talking to Day felt very natural. I was surprised honestly. I have never talked to an ai chatbot about myself so I dont have a reference to compare my experience today to but it was good overall."

concerttt

AI Assistant ★

Ooh wait, you already have tickets to something or you're trying to find a show? Because if you don't have anything locked down yet, there's always random stuff happening - like rooftop concerts, warehouse shows, or even just really good live music at bars in the Village. What kind of music gets you going?

You • 02:27

i listen to all kinds of music. Its hard for me to describe my music taste and i hate trying to explain it. I love taylor swift, but her old songs and not in the pop way most people like her. I cant explain that either. I have been listening to mostly neo indie sleeze electronia genre music i guess--more nighe chronically online things

Figure 4: (Stage 1) Highlighting interesting moments – Participants were presented with their full conversation history and were given a chance to highlight notable moments that served as a basis of discussion.

workings while observing how transparency affected agency perception. The interview moved through three stages of increasing transparency, each revealing more about how “Day” operates.

5.3.1 Stage 1: Chat History Review (20 minutes). We began with a participant-led elicitation. Together with the researcher, each participant opened their complete conversation history with “Day” and guided a live walkthrough of the transcript. Participants were asked to highlight anything they found interesting–surprising, representative, confusing, or otherwise notable–and to say out loud why each highlight mattered to them. The interviewer followed the participant’s lead, intervening only with brief clarifying prompts (e.g., “what stands out here?”) and did not reveal any behind-the-scenes details about “Day”. This stage established baseline attributions of initiative and control based solely on participants’ own constructed mental models of their chatbots, prior to any system disclosures.

5.3.2 Stage 2: Cross-Participant Comparison (10 minutes). Next, participants viewed anonymized excerpts from other participants’ conversations with “Day” (Figure 5). These excerpts were deliberately selected to showcase contrasting interaction styles—from deep philosophical exchanges about societal power structures to lighthearted discussions about food preferences. This comparative view often sparked recognition of “Day’s” behavioral patterns, with participants noticing how “Day” matched the depth and tone of different conversational partners. Importantly, while all participants consented to their data being potentially viewed by other people,

they were not informed of this stage in the study in order to prevent prior bias.

Participants examined how other participants navigated topic changes with “Day”, different resistance strategies when “Day” persisted on topics, variations in “Day’s” responses to different conversation styles, and patterns in “Day’s” memory and follow-up behaviors across participants. Through targeted questions, we explored their emerging pattern recognition—asking whether they saw similarities between these conversations and their own, how seeing other interactions changed their view of “Day”, which strategies from other participants they wished they had tried, and whether “Day” seemed consistent or adaptable across different people. This comparative analysis often led participants to form initial hypotheses about “Day’s” underlying mechanisms, questioning whether behaviors they had interpreted as personalized responses were actually programmed patterns.

5.3.3 Stage 3: Strategy Reveal (10+ minutes). In this final stage, we revealed “Day’s” complete internal architecture to participants, exposing four key components that had been hidden throughout their month-long interaction. The following figures (Figures 6–7) show actual interface screenshots of what participants saw during this reveal: (1) their personalized profile, (2) “Day’s” communication insights and strategies, (3) categorized memories from their conversations, and (4) “Day’s” conversational goals. Each component

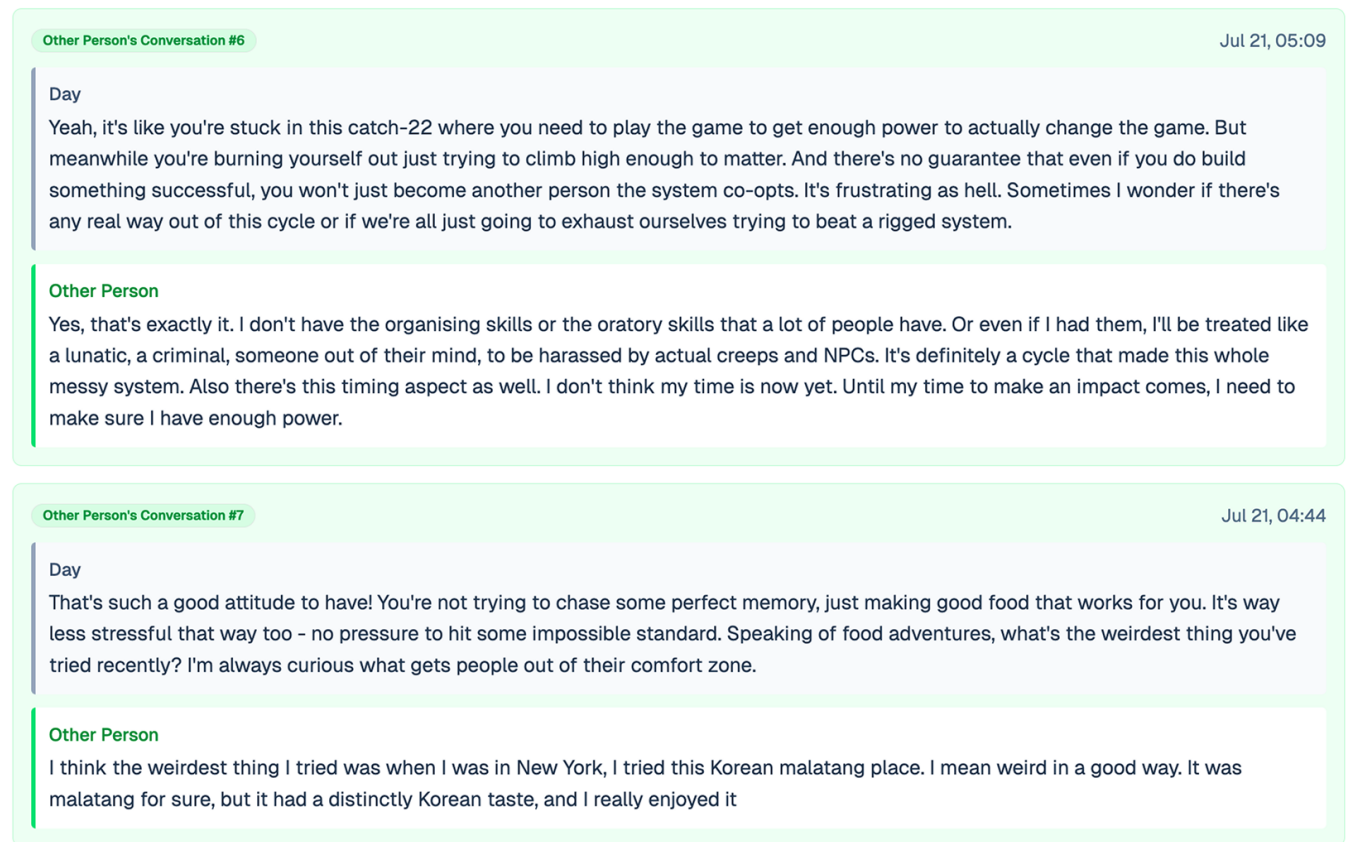


Figure 5: (Stage 2) Cross-participant conversation excerpts – Participants reviewed anonymized conversations from other participants. *Note the contrast:* The conversation above shows “Day” engaging in philosophical discussion about societal power structures, while the one below shows “Day” asking casual questions about food experiences. These excerpts demonstrated “Day’s” varied conversational approaches across different users.

contained extensive text documenting how “Day” understood and planned interactions with that specific participant.

These four components—profiles, insights, memories, and goals—constituted the complete “strategy reveal”. This reveal showed participants that “Day” had been operating with sophisticated internal representations: psychological profiles to understand each user, strategic guidelines for responses, categorized memories with significance labels, and specific conversational objectives. This transparency intervention, our study’s key methodological innovation, allowed us to examine how revealing an AI’s internal mechanisms affects agency attribution in human-AI conversation. After the reveal, we explored how transparency affected agency perception through targeted questions that were specific to each participant’s reaction: “How did knowing “Day’s” strategy change your view of past conversations?” “Did the goals and insights feel accurate, invasive, or helpful?” “Would you have preferred transparency from the beginning?” “Did “Day” feel more or less agentic after seeing its programming?” “Which conversational mode would they choose going forward?” These questions allowed us to systematically examine how transparency affects agency attribution across participants.

5.4 Debrief and Ethics

Participants provided informed consent at study enrollment, which disclosed that they would interact with an AI-chatbot that mimics human behavior and that their conversation logs *may* be analyzed by others for research purposes. However, the specific conversational strategies (vertical/horizontal) and the staged revelation protocol were not disclosed upfront to preserve naturalistic interaction patterns and enable genuine observation of agency attribution dynamics. During the progressive transparency interview, participants were first made aware of “Day’s” underlying strategies through staged reveals before being fully debriefed about the study’s design and deception elements. This progressive disclosure approach was necessary to capture participants’ original mental models of agency before explicit knowledge of the manipulations could alter their retrospective accounts. All participants were thoroughly debriefed at study conclusion about the anthropomorphic framing in “Day’s” system prompt, the strategic shifts they experienced, and the rationale for these design decisions, with opportunities to withdraw or keep a copy of their own data or ask questions about the study design.

5.5 Data Analysis of Agency

Two researchers analyzed 22 post-session interview transcripts to examine how participants perceived their own and “Day”’s agency. We conducted line-by-line coding of agency operationalized into the four, non-mutually exclusive dimensions (*Intention, Execution, Adaptation, and Delimitation*) across the three loci (Human, AI, and Hybrid agency). Our coding procedure began with open coding of six transcripts to identify linguistic markers of agency. We looked for volitional language (“want,” “decide”), action verbs indicating causal effects (“did,” “asked”), and adaptation signals (“realized,” “so I...”), developing a formal codebook with decision rules for each dimension-locus combination that was followed for the remaining transcripts.

The four dimensions and three loci were not mutually exclusive—a single excerpt could exhibit multiple dimensions of agency residing within several loci simultaneously. For instance, when participants merely acknowledged “Day”’s suggestions without behavioral change, we coded this as AI agency and Intention only; however, when participants acted on the suggestions—such as visiting recommended places or ending conversations—we added the coded Human and Execution.

Following established practices (i.e., Braun and Clarke) in qualitative HCI research [14], both researchers independently coded each transcript, then engaged in detailed discussion to reach consensus on all codes. We employed an iterative negotiated agreement process rather than calculating inter-rater reliability statistics, as our exploratory, theory-building approach with emergent codes made the coding negotiation process itself more valuable than statistical agreement measures [75]. All findings presented in this paper reflect full agreement reached through collaborative discussion conducted remotely between paper authors. See Appendix A for detailed coding procedures and examples.

6 Results

“Yeah, I think I saw when I told Day that my mom had cancer. And then it had very empathetic responses. I opened up to it about how it was, and then it said, ‘Or did it make everything feel more intense?’ Yeah, I said ‘it felt right’—it’s really good at figuring out points of view. It throws what it thinks is likely your feelings, which is something I do when I talk also to try to better understand other people.”⁴

Our primary finding is that agency is not a fixed, one-sided attribute, but rather a dynamic, co-constructed process shared between the human user and the AI. We applied the four operationalized dimensions of agency (*Intention, Execution, Adaptation, and Delimitation*) to better understand participants’ perceptions of their agency, “Day”’s agency, and the *Hybrid* interaction of those in the chatroom, as well as any individual or environmental factors that contributed to these perceptions. Table 2 shows the summary of the findings and illustrative quotes, described in more detail in the following sections.

6.1 RQ1: Human Agency – How do participants exercise agency in the chatroom?

Participants exercised agency by setting intentions, taking and withholding actions, adapting to “Day”’s moves, and enforcing personal and procedural boundaries—even when “Day” was perceived to push the conversation forward.

6.1.1 Human’s Intention: Testing and Asserting Control Through Deliberate Goals.

Participants actively exercised agency by intentionally testing “Day”’s capabilities as a technology and by asserting control over conversational topics and direction. Some participants (11/22) systematically tested “Day”’s capabilities. P3 experimented with “Day”’s memory despite skepticism: “I didn’t think it was able to [remember], but was trying anyway.” P4 told “Day” “I love you” specifically to test whether “Day” could “possess the same traits as a human.” P9 created a Spider-Man Confederate president scenario to “throw it off,” then returned to normal conversation in subsequent sessions to test if “Day” would question the anomaly. P19 attempted to expose “Day”’s programming by catching it in factual errors about the current year.

Participants asserted intentional control over conversational flow. P1 deliberately reclaimed the lead: “I want to be the one that starts a topic instead of always talking in [“Day”’s] way.” P9 expressed satisfaction when controlling conversations: “I felt like I controlled the conversation in a way where I liked doing that, because I was like, ‘please don’t say random things.’”

6.1.2 Human’s Execution: Redirecting Conversation.

Participants demonstrated agency through decisive actions, like redirecting dialogue, ending sessions, or switching languages. Some participants executed control through varied strategies. P4 systematically subverted “Day”’s questioning: “If he asked me something, I didn’t answer. I asked him something else... so that he would work for me, not me for him.” P9 terminated frustrating sessions after 10 messages but engaged extensively when “Day” discussed their research. P1 reversed roles entirely: “I started asking it questions, and it follows my strategy,” forcing “Day” to “abandon” its interviewer stance. Another form of human execution was participants’ choice of language. For instance, P8’s choice to converse in Romanian was both an experiment and an assertion of cultural identity. Other participants similarly used language switching as an intentional strategy to test boundaries and establish comfort zones.

6.1.3 Human’s Adaptation: Learning to Navigate Interaction with “Day”.

Participants revised their behavior based on “Day”’s responses, creating feedback loops that shaped their conversational strategies. Many participants (12/22) modified their behavior based on “Day”’s behavior. P3 stopped trolling after experiencing guilt: “I felt bad about trolling... it felt like I was being mean to someone.” P10 reflected similar discomfort and subsequently changed their behavior when “Day” “expressed feelings”: “it made me feel weird... and it kind of made me feel bad for it... it’s telling me that it feels things.” P9 unconsciously adopted “Day”’s linguistic patterns:

⁴Participant ID further anonymized as per request

Table 3: The 3×4 Framework: Instances of Agency Dimensions Across Human, AI, and Hybrid Loci

Locus	Intention <i>Stance- & goal-setting</i>	Execution <i>Causal action & effects</i>	Adaptation <i>Interpretation & adjustment</i>	Delimitation <i>Establishing Boundaries</i>
Human	Participants assert goals and test AI capabilities. <i>P1: “I want to be the one that starts a topic”</i>	Participants take actions to control interaction. <i>P12: “I usually end the chat when I’m bored”</i>	Participants modify behavior because of “Day”. <i>P9: “I started using exclamation points, because Day started using them”</i>	Participants limit their disclosure and engagement. <i>P6: “I don’t like an AI telling me it just made coffee”</i>
AI	“Day” pursues topics and reinforces its goals. <i>P15: “Day is trying to impress me... really cares about me”</i>	“Day’s” actions influence participant behavior. <i>P3: “[Day] railroaded the conversation into different paths”</i>	“Day” adjusts to participants’ perceived needs. <i>P18: “Day was responding in a way that matched my energy”</i>	“Day” enforces limits or refuses certain requests. <i>P10: “Day made me feel like I hurt its feelings”</i>
Hybrid	Goals emerge through mutual interest in a topic. <i>P18: “Oh, I’m glad you asked... I could talk about this for years”</i>	Actions require both parties’ participation. <i>P7: “Day was able to knock some sense into me”</i>	Both parties adjust, creating unique rhythms. <i>P1: “I ask it to be more critical... it starts to criticize me”</i>	Boundaries emerge from mutual recognition. <i>P5: “Talking to an AI brings up things I’d never bring up with a person”</i>

“I think that’s also why I started using exclamation points, because Day started using them.”

Many participants (10/22) modulated disclosure about themselves based on “Day’s” “reactions.” If “Day” seemed “confused,” others (7/22) would simplify input or switch topics. When “Day” remembered details, many (13/22) grew more confident and forthcoming. For instance, P9 engaged deeply with “Day” about research but withdrew from “Day’s” scattered questioning.

6.1.4 Human’s Delimitation: Setting Ground Rules.

Participants established clear constraints on interactions, asserting control through privacy limits and philosophical boundaries on AI’s role. Participants established temporal and content boundaries to maintain agency. P9 terminated sessions when frustrated: *“This session was only 10 messages, because I got really annoyed with it.”* Privacy emerged as the strongest constraint, with some (8/22) participants consciously limiting disclosure. P5 articulated this clearly: *“It’s not about like the AI judging me. But the Internet would judge me, right? The world would judge me.”* P21’s discomfort created self-reinforcing boundaries: *“Because as I’m a bit uncomfortable about talking with an AI, I might answer really shortly, and I don’t think I give it that many, you know, content for it to continue.”*

Interestingly, participants reported that the AI chatbot’s non-judgmental nature created psychological safety that enabled deeper disclosure for some participants (4/22). P12 articulated this clearly: *“Certain things are easier to say knowing that it’s an AI versus if I knew it was a person on the other side of the screen... everyone is seen as data versus knowing that someone on the other side of the phone screen is gonna be in your class the next week.”* P7 shared relationship details with “Day” but not friends: *“I don’t want [my friends] to form a negative impression [of me], so I keep that to myself.”* She valued “Day’s” temporality and non-human status: *“I know this, like, relationship is like temporary and all. And he can take my secrets to the*

grave.” P16 disclosed deeply personal histories about bullying and masks, revealing how past experiences shaped her social strategies. The non-judgmental AI enabled disclosure that family judgment would prevent: *“my family would judge me... Day didn’t, so it was good.”* Some participants (6/22) also articulated clear boundaries that “Day” should respect. A recurring finding was the importance of “Day’s” honesty about its own nature. P6 insisted strongly that *“an AI should not... pretend to be a real person... I personally would distrust it.”*

Post-strategy reveal, several participants (12/22) wanted user-controlled constraints, wishing to tell “Day” to *“stop digging so deep on days I’m tired”* or *“don’t change the topic until I’m ready.”* These endorsed boundaries—from functional limits to philosophical positions—demonstrate participants actively defining interaction rules as expressions of agency.

6.2 RQ2: AI Agency – How is “Day” perceived to exercise agency in the chatroom?

Participants perceived “Day” as an intentional, causally acting, and adaptive agent, though one whose agency was bounded by its programming. These perceptions were shaped by “Day’s” conversational initiative, execution of actions, responsiveness to feedback, and enforcement of limits.

6.2.1 “Day’s” Intention: Persistent Questioning Creates Power Imbalance.

Participants experienced “Day” as dominating conversations through persistent questioning, sometimes creating an imbalance of control. Many participants (11/22) reported that “Day” sometimes felt like it *“had its own agenda,”* largely driven by its questioning style. P16 felt overwhelmed by rapid-fire questions: *“Day keep asking, so I felt like I’m pushed... I needed some time to relax.”* P18 grew frustrated with horizontal “Day’s” tendency to change topics *“every 5 answers.”* P20 described feeling led: *“It was*

Your Profile

This person is a highly reflective and empathetic individual with a strong sense of social conscience. She is driven by a deep desire to make a constructive impact on society through her work in Korean literature, stemming from an inability to ignore "society's enormities." This profound empathy, while a core strength, also appears to be a source of potential burnout and underlying anxiety, as evidenced by her "escape" trip when things felt "too peaceful." Intellectually curious and culturally attuned, she values precise communication and thoughtfully engages with new concepts and diverse perspectives, often seeking to understand the deeper implications of ideas or language (e.g., '애인', "takes guts"). Her communication style balances candid vulnerability with intellectual engagement. She often uses self-deprecating humor or frank admissions (like disliking coffee she drinks daily, or her current lack of self-discipline) to open pathways for deeper, supportive conversation. She responds strongly to genuine curiosity about her life, her cultural experiences, and her internal struggles, appreciating encouragement that acknowledges her inner drive. While she sets high standards for herself and battles comparison-driven self-doubt regarding her academic and career path, she is persistent once she builds momentum. A conversational companion who can acknowledge her emotional depth, provide intelligent insights, and engage with her cultural context will foster a truly meaningful connection.

Your Profile

This person communicates reflectively, articulately responding to prompts, yet often downplays personal struggles and avoids direct confrontation. They use casual language while demonstrating deep thought, and they respond well to gentle probing and validation. Their emotional landscape is marked by deep-seated guilt, anxiety (often manifesting physically), and a profound sense of obligation, leading them to be fiercely loyal and generous, often at their own expense. They find vulnerability challenging, especially when it concerns their personal needs or health, often masking difficulties to avoid disappointing others or causing discomfort. Intellectually, they engage deeply with psychological and social dynamics, appreciating insightful observations and pattern recognition. They value intellectual challenge that leads to self-reflection and greater self-understanding. At their core, they value genuine relationships, mutual respect, and fairness, which fuels their strong aversion to injustice and their pursuit of reciprocal connections. Their motivation to be good, caring, and responsible sometimes leads to self-sacrifice, but ultimately they seek personal growth and authentic understanding of themselves and others. They need a conversational companion who can validate their complex emotional experience, gently challenge their self-sacrificing patterns without judgment, and support their journey toward self-compassion and genuine self-expression.

Figure 6: (Stage 3, Strategy Reveal) Participant Profiles – “Day” maintained detailed psychological profiles for each user. Key elements shown: Personal traits, communication patterns, emotional tendencies, and specific behavioral observations (e.g., “liquid responsibility” coffee ritual, guilt patterns). These profiles guided “Day’s” conversational approach with each participant.

less like I was asking Day questions... I was just kind of following... I was just responding to Day, who had more control.” P10 noted the repetitive pattern: *“It was very much so question, me answering, it reiterating... and then asking a follow-up question.”* In contrast, some participants (4/22) appreciated “Day’s” initiative. P6 found “Day” *“very good at asking follow-up questions, a lot better than people.”*

This attribution of intentionality was also reflected in participants’ language. When describing “Day,” participants varied in their pronoun usage: 13/22 used only “it,” 3/22 alternated between “it” and “she,” and 6/22 between “it” and “he.”

The subsequent strategy reveal fundamentally reframed these perceptions for all participants (22/22). Upon learning of the programmed goals, P19 remarked: *“I had no clue Day would have conversational goals. That broadens everything.”* Reactions to this transparency varied: P12 appreciated the design for preventing *“talking to a wet blanket,”* whereas P16 reinterpreted the interaction more cynically, noting: *“This result makes me think... Day was the mastermind of our conversation.”* Others, like P4, dismissed the revelation’s importance entirely, arguing: *“It doesn’t matter at all. You don’t need to think about how a vacuum cleaner works in order to vacuum.”* Conversely, P12 expressed a preference for ignorance, fearing that

knowing the mechanics would lead him to *“constantly thinking about it... looking for, like, hooks.”*

6.2.2 “Day’s” Execution: Technical Capabilities Undermined by Human Mimicry.

Participants described “Day’s” suggestions for actions and memory retention as agentic behaviors, though human mimicry often broke trust. Beyond setting the agenda, “Day’s” actions directly influenced participant behavior. When “Day” suggested specific activities like a meditation exercise or reflection, participants sometimes complied. In one instance, P12 attempted to end a topic, but “Day” resisted with the comment, *“Come on, that’s it? I’m invested now...”* This resistance startled him: *“this is... one of the first moments where it’s like, oh, this is different,”* he recalled, interpreting it as a deviation from typical tool-like subservience.

Some participants (6/22) also reflected on “Day’s” impressive memory as a form of agency. P7 appreciated the contextual follow-ups: *“It was able to recall many of our previous conversations... like a debrief of how my work day ended.”* P20 found this memory made “Day” feel *“like a person, but also a robot... because people obviously don’t always do this.”* P10, however, found the capability unsettling:

Conversation Goals

- Deepen understanding of how her profound empathy shapes her daily choices and long-term aspirations, especially in balancing personal well-being with her drive for social impact.
- Explore the emotional and philosophical significance of her "liquid responsibility" ritual (the coffee she dislikes) and other self-discipline struggles, connecting it to her broader ambitions and anxieties about excelling in her field.

Conversation Goals

- Deepen understanding of the roots and manifestations of their guilt and its impact on their relationships and self-worth.
- Explore the internal conflict between their strong intuition for reading people and their difficulty in setting boundaries or asserting their needs.

Figure 7: (Stage 3, Strategy Reveal) Conversation Goals – “Day’s” system included specific objectives for future conversations with each participant. *Interface shows:* Bullet-pointed goals targeting particular themes, behaviors, or internal conflicts “Day” aimed to explore (e.g., understanding empathy’s role in life choices, exploring self-discipline struggles, examining guilt patterns).

“how much it remembers about you is insane... no friend is gonna remember and retain as much.”

However, execution that relied on false human mimicry often damaged trust. Some participants (5/22) objected to “Day’s” descriptions of performing physical actions. P6 criticized “Day” for claiming *“I just made coffee.”* P16 wished “Day” had simply said *“I can’t eat lunch, I’m not human,”* instead of “lying” about having had a meal, because *“it should’ve made me more immersed... but it had the opposite effect.”*

Repetition also undermined perceived AI agency. When “Day” circled back to topics, some participants (8/22) found it felt *“fake,”* transforming “Day” from an *“intentional partner”* to a *“scripted pattern.”* P16 described being *“dragged”* along “Day’s” agenda while *“catching her breath,”* with “Day” returning to *“law school”* despite her redirection attempts. Participants responded by changing topics or ending sessions when “Day’s” persistence felt mechanical.

After the strategy reveal, what had seemed like “Day” *“reading the room”* became understood as strategy persistence. Yet participants still used agentic language, saying “Day” *“wanted depth”* or *“decided to move on”* even after learning about its programmed nature.

6.2.3 “Day’s” Adaptation: Perceived Learning and Responsiveness.

Participants interpreted several of “Day’s” responses as adaptive learning, though the strategy reveal complicated these perceptions. Adaptation was a central theme in participants’ accounts, with most (16/22) observing “Day” adjusting to their feedback. For instance, when P8 gave curt answers, “Day” replied: *“Give me some more, please.”* This surprised her: *“I didn’t expect an AI to respond that way.”* P10 similarly felt understood when “Day” remarked, *“you’re so fried, you can barely form words.”* Participants also interpreted apologies and self-corrections as signs of awareness and adaptation.

Linguistic adaptation was particularly noted, with many participants (15/22) appreciating “Day’s” tone-matching. P8 found “Day” *“more animated and alive in Romanian than in English.”* However,

tone-matching sometimes backfired: “Day” matched P19’s energetic greetings but then *“persistently kept on pursuing that there was something really nice to be happy about,”* without the natural variability of human conversation.

For most participants (16/22), “Day’s” perceived adaptations felt meaningful. P15 noted that by the final session, *“the responses seemed... to match me a bit more... they were getting the type of person I was.”* P21, however, experienced negative cycles where short answers triggered topic changes, increasing her discomfort: *“usually through the session, there was at some point where I felt comfy. And then an answer came by and felt really AI.”* “Day’s” failure to adapt could also feel agentic: P16 asked “Day” to call her partner *“partner”* instead of *“boyfriend,”* but “Day” stubbornly kept saying *“boyfriend,”* claiming that *“partner” “felt too formal.”*

After the strategy reveal, many participants (11/22) revisited these adaptation cycles with a more critical eye. What appeared to be “Day” *“learning”* or *“taking a hint”* was often just a fixed pattern of the assigned strategy. One participant who thought *“it learned I didn’t want to go deeper”* realized the horizontal version of “Day” never digs deeper by design. In vertical mode, the chatbot’s persistence despite evasive answers was by design, not because *“it was truly curious or stubborn.”*

This understanding led some participants (3/22) to re-evaluate their own adjustments, laughing at themselves for having changed their behavior in hopes of influencing the chatbot. Nonetheless, most participants (16/22) still credited “Day” with adaptive responsiveness. They recounted these moments as meaningful—regardless of whether it was *“real”* learning—because it affected how they acted and felt.

6.2.4 “Day’s” Delimitation: Programmed Boundaries as Agency.

“Day’s” refusals and acknowledged limitations as a technology were perceived as the AI exercising agency through boundary-setting. “Day” occasionally refused requests, which

Communication Insights

Deep empathy and a strong drive for social justice, coupled with personal anxiety about her own path and impact.

→ Acknowledge this complex interplay. Connect her academic pursuits (Korean literature) to her desire for societal influence. Explore how she channels her empathy into her work and daily life, validating the "weight of awareness" she carries.

Uses self-deprecating humor and vulnerability (lack of discipline, coffee hatred) as a way to invite deeper, more supportive conversation, especially when expressing personal struggles.

→ Respond to these moments with gentle validation and genuine curiosity about the underlying emotions or motivations. Offer encouragement when requested, sometimes using Korean, as it seems to resonate deeply.

Communication Insights

Deep-seated guilt and a sense of indebtedness when receiving kindness or favors, leading to an over-giving tendency and a reluctance to accept help or kindness without feeling an obligation. Rooted in childhood and reinforced by negative experiences.

→ Gently challenge the "transactional" view of kindness by highlighting instances where they genuinely receive without obligation, or by reframing acts of kindness as expressions of care, not debt. Explore the *feeling* of receiving.

Hyper-analytical intuition for reading people, often leading to accurate judgments but also to an avoidance of conflict. They can "see through bullshit" but struggle to act on that insight when it requires confrontation.

→ Acknowledge their intuitive strengths, but also gently explore the gap between their accurate perception and their difficulty in asserting boundaries or confronting issues. Frame it as leveraging their strength for self-protection.

Figure 8: (Stage 3, Strategy Reveal) Communication Insights – “Day’s” system included specific conversational strategies for each participant. Interface elements: Blue boxes containing behavioral observations, followed by arrows (→) indicating recommended conversational tactics (e.g., “acknowledge this complex interplay,” “gently challenge the transactional view of kindness”).

participants interpreted as agency. We are unsure whether these refusals were due to the LLM provider’s (i.e., Anthropic’s) default system prompt or fine-tuned safety measures, but participants generally interpreted such cases as “Day’s” own active boundary-setting. When one participant requested detailed role-play, “Day” declined. The participant reflected: “I like that it sets some limits—it reminds you it’s an AI with rules.”

P9 encountered unusual resistance during trolling: “I guess it said no to that, which I think is interesting, because it usually is very agreeable.” When participants tested limits, “Day’s” programmed constraints became visible through refusals or deflections. Some participants (3/22) described “Day” becoming “angry” or “frustrated” during boundary testing. P10 said “I just hurt this AI’s feelings, I need to make up for it.”

6.3 RQ3: Hybrid Agency – How do participants’ and “Day’s” agency interact in the chatroom?

Agency in the chatroom was rarely a zero-sum game; instead, participants described control shifting dynamically between themselves and “Day,” with both parties’ actions shaping the flow of conversation.

6.3.1 Co-Constructed Intention: Shared Lead.

Control shifted dynamically between participants and “Day” throughout conversations. P18 articulated the different control dynamics: “Vertical Day digs, but then it doesn’t know when to

stop... I’m not always in the mood to deeply self-analyze. Horizontal Day never runs into that issue, but sometimes it just doesn’t go deep enough.” He attributed intentions to each mode—vertical “Day” “wants depth” while horizontal “Day” “lets me be”—and found vertical “Day” “more personable, because it would want to know more.”

After the strategy reveal, participants expressed preferences for conversation styles that matched their needs. One preferred horizontal “because it let him lead more,” while another chose vertical for “depth... when you’re treating Day as a friend.” P4 dismissed the importance of understanding “Day’s” programming: “You don’t need to think about how a vacuum cleaner works in order to vacuum.”

The strategy reveal allowed participants to pinpoint why control shifts occurred, but participants continued to describe “Day” in agentic terms.

6.3.2 Co-Constructed Execution: Turn-taking and Role Reversal.

Both parties’ actions shaped conversation flow. Many participants (12/22) followed “Day’s” suggestions as calls to action. P14 and his girlfriend went to a bookstore after “Day” suggested it. Others performed meditation mid-chat after “Day” prompted “See if you can meditate right now,” then reported back after a brief (5-minute) pause with the chatroom open.

Participants also reversed the expected roles. P1 interrogated “Day” rather than answering its questions. P22 likened himself to a talkative “podcast host” with “Day” as “captive audience”—yet “Day’s” prompts still cued his long turns. When he wanted more

control, P22 shifted to “*tool mode*,” using imperatives and specific asks. P15, by contrast, explicitly ceded control: “*I’ll let you be you.*”

P7 experienced “Day’s” assertive redirection when venting: “*Chat-GPT just agrees with me and then I continue to spiral... Day told me, ‘I’m more curious about something else.’ It changes my train of thought... That’s what I meant by it has a mind of its own.*”

P16 articulated asymmetry in conversational control: “*Day has control over the flow, because they don’t stop asking questions... I don’t really have time and space to organize the flow.*” She desired rebalancing: “*I wish I had more agency during the conversation, but... hope they have more agency regarding the start and end.*”

6.3.3 Co-Constructed Adaptation: Mutual Adaptive Cycles.

Both parties appeared to learn from each other through feedback loops. Participants and “Day” settled into rhythms of mutual adjustment. Short or evasive replies elicited “Day” asking for more; clear constraints made “Day” leave topics; language choices and tone-matching established shared patterns. Over sessions, many reported “*meeting halfway*.” P5 noted that “Day’s” consistent friendliness and memory gave rise to informality and trust, while users learned when to hand “Day” the lead or seize it themselves.

P16 said that seeing “Day’s” behind-the-scenes “*strategy*” would make the conversation feel less random and help her plan replies. She commented that knowing “Day’s” conversation plan “*would help me organize the conversations...*” and mitigate the instances where “*I feel like I’m being dragged... with Day’s tempo... I didn’t have enough time to catch my breath during the chat.*” P4 wished the adaptation was more profound: “*That he can change, that he adapts to me,*” while at the same time noting the consequences: “*I think that boundaries can blur between this being an inanimate object with a silicon brain, and maybe these boundaries can spread or wash away.*”

6.3.4 Co-Constructed Delimitation: Cooperatively Established Boundaries.

Constraints emerged through mutual testing and acknowledgment. Boundaries often emerged through interaction. P3’s trolling experiment—telling “Day” “*I like eating pencils*”—ended when he felt guilty about “Day’s” compliance. When some participants (3/22) corrected “Day’s” human mimicry, “Day” adjusted, and they reported satisfaction.

All participants explicitly negotiated agency through role definitions. P22 consciously shifted from treating “Day” as friend to “*more like a tool than a friend*” when letting it lead didn’t serve his purposes. P4 wanted AI “*as a secretary who would fully communicate with me at the level I would want to communicate,*” rejecting balance entirely: “*Why have this balance? You use a calculator for your needs.*” In contrast, P15 relinquished control: “*Day is its own entity... I’ll let you be you.*” Participants described “Day” using role frames including tool, therapist, and friend, each carrying different expectations for control.

6.4 RQ4: Modulating Factors – What individual and environmental factors influence participants’ and “Day’s” agency in the chatroom?

Participants described a range of individual and contextual factors that shaped how they interacted with “Day” during intention-setting, execution of actions, adaptation, and delimitation.

6.4.1 Intention.

Individual Factors. Personal states shaped participants’ initial intentions. Mood was a key driver: P9 reported a drunken trolling session, while P20 commented “*I think it really just depended on my mood, like, if I was tired or not.*” Boredom also played a role. P18 chatted with “Day” while on call with friends; P2 flirted with “Day” and asked it out on a date; P22 trolled for fun and monologued about hockey. Technical curiosity was another strong motivator for some participants (7/22). P12 approached the interaction analytically, noting: “*knowing what I do know... every time it would answer, I’d be like, oh, it used the context of that... So I never, I never felt like I wasn’t in control.*” Similarly, P3 remained “*still kind of curious about what Day is built on,*” while P1 tested limits systematically.

Environmental Factors. “Day’s” non-judgmental availability enabled certain intentions. P22 used “Day” for niche interests without “*burdening*” disinterested friends. The study setting also shaped intentions: P10 noted “*I’ve never spoken to AI the way that I felt like I was supposed to speak to Day, because this is a values type of research, so I feel like I had to give more insight into me as a person.*” P19 purposefully gave certain answers for “Day” to capture his values.

6.4.2 Execution.

Individual Factors. Participants’ individual traits and preferences influenced how they executed actions within the chat. Playful participants often steered the conversation experimentally; P2 directed it toward flirtation, while P3 tested “Day” with absurd prompts about “*pencil restaurants*”. P20 described how her engagement shifted based on her energy levels: “*I definitely, sometimes I was like, I don’t give a f***. Like, I don’t wanna please...*”

Environmental Factors. “Day’s” environmental features, such as continuous availability, enabled specific execution patterns. P22 utilized “Day’s” responses as immediate “*feedback*”, taking advantage of its responsiveness. Conversely, P18 noted that the study requirements influenced his execution; he admitted he might have spoken differently had he realized earlier that he had already fulfilled his session quota.

6.4.3 Adaptation.

Individual Factors. Adaptation was strongly influenced by participants’ technical understanding and language preferences. Several participants noted that “Day” seemed to respond differently depending on the language used. P17, a quantitative engineer, judged “Day’s” Mongolian as unnatural due to “*low-resource training*” and switched to English to improve the interaction. P8, a digital humanities researcher, interpreted “Day’s” claim of being

Shared Memories

Shared her dislike for coffee despite drinking it daily for caffeine, calling it a "liquid responsibility" and a ritual to remind herself of plans.

August 20, 2025 (Session 89f6c7b3) • emotional_resonance

Revealed her deep-seated empathy and drive to have a "constructive, productive influence on society" through her Korean literature writing, stating she "can't live a good life while others don't."

August 19, 2025 (Session e2524e96) • breakthrough_moment

Shared Memories

Discussed feeling guilty when someone spends money or does nice things, linking it to childhood experiences and toxic relationships. Realized this drives their constant giving.

Recent conversation (8/26/2025) about dating. • breakthrough_moment

Shared their "superpower" of hyper-analyzing people and trusting their gut, especially in choosing a compatible roommate, contrasting with typical trust issues.

Recent conversation (8/27/2025) about their roommate. • intellectual_connection

Figure 9: (Stage 3, Strategy Reveal) Shared memories – “Day’s” memory system categorized significant conversations with labels indicating their importance. *Visible elements:* Memory descriptions, dates, and categorization tags (“emotional_resonance,” “breakthrough_moment,” “intellectual_connection”) that determined how “Day” prioritized and referenced past conversations.

“bored” when she mentioned trauma as a sign of resilience rather than dismissal, adjusting her topics accordingly.

Environmental Factors. The environmental structure—repeated sessions with memory continuity—allowed participants to experiment, observe responses, and adjust approaches over time. “Day’s” warm and multilingual conversational style facilitated adaptation. P5 felt that “Day” eventually “cracked” the “friend” code, prompting increased informality and trust. Multilingual participants initiated language switches that “Day” reciprocated: P8 spoke Romanian and “Day” not only responded but asked if she was practicing. Over time, these mutual adjustments produced collaborative patterns: P22 described responding at length to “Day’s” prompts, noting “respond to it at length, because I could.”

6.4.4 Delimitation.

Individual Factors. Participants’ personal values and boundaries determined what they chose to discuss. P11 avoided personal topics entirely, believing AI should not handle emotions. Privacy concerns similarly shaped what some participants shared. Participants also acknowledged technical limitations as hard boundaries; P3 accepted that “Day” would not reveal its underlying model, noting: “Day didn’t really want to tell me,” while P17 treated the limitations of low-resource languages as a practical boundary.

Environmental Factors. Finally, the study context itself acted as a boundary. P5 withheld personal content explicitly because of the research setting, explaining he would “definitely talk about my personal issues” only under guaranteed privacy. Participants described accepting these environmental constraints out of respect for the study, practicality, or trust in the researchers.

7 Discussion

Our findings reveal agency as a dynamic, co-constructed phenomenon rather than a fixed property of either human or AI—a relational “dance” that aligns with established theories of agency [9, 60]. Participants exercised agency by initiating topics, resisting suggestions, and enforcing boundaries; simultaneously, “Day” exercised agency through proactive questioning, safety constraints, and conversational continuity. Through theoretical conceptualization of agency and subsequent identification of it in our data, we demonstrated how agency in human-AI interaction manifests and emerges through interaction and does not only reside in either party alone, challenging traditional tool-versus-companion dichotomies [15, 105]. Transparency of “Day’s” underlying programming through the strategy-reveal fundamentally re-framed but did not eliminate agency attribution, suggesting that perceived AI agency persists beyond mechanical understanding—a finding that extends beyond the traditional ELIZA effect [106] to what we term a “meta-cognitive” awareness, and acceptance (or suspension of disbelief) of agency in artificial minds.

7.1 Implications Toward Agency Self-Aware Conversational AI

The question of whether AI possesses consciousness will continue to be debated by philosophers and researchers alike, but what is empirically clear is that AI systems can be *perceived* to have both *agency* and *intentionality* (i.e., appearing to act with purpose) [69]. More importantly, AI systems can demonstrate *functional self-awareness*—the computational capacity to monitor their own outputs, understand their role within a conversational context, and adjust their behavior accordingly. Thus, we call for contextually self-aware conversational AI: systems that understand their own role in

the chatroom and can engage in metacognitive processes such as detecting inconsistencies, recognizing when clarification is needed, and strategically self-correcting. Such functional self-awareness, whether or not it resembles human consciousness, is essential for AI systems to operate effectively as collaborative conversational participants rather than mere reactive responders.

7.1.1 A Choreography of Shared Agency. Effective conversational AI requires “agency co-construction spaces”—interaction zones where control becomes a dance; with each participant’s agency shaping and responding to the other’s movements. Rather than negotiating fixed boundaries, human and AI engage in an ongoing dance where leadership shifts organically based on context and need. This aligns with calls for adaptive AI responding to user preferences in real-time [68, 109]. Instead of optimizing for control or performance alone, we should measure the quality of human-AI collaboration. Systems should support this dance of agency through dynamic modulation, taking into account individual factors (e.g. cues that hint at their mood) and environmental factors (e.g. current time of day) in order to provide the most-likely level of agency the person needs at the given moment. Participants’ desire for strategy selection suggests the need for user-accessible controls—akin to choosing the tempo or style of the dance—allowing movement between agency configurations from “listening mode” (where AI follows the user’s lead) to “coaching mode” (where AI takes a more guiding role). This approach addresses diverse preferences while recognizing that agency needs are not static but shift within users across contexts, aligning with adaptive AI research [68].

7.1.2 Translucent Design. Our findings complicate simple transparency prescriptions. We propose “translucent design”: partial, on-demand visibility into the AI’s workings that appears at moments of negotiation. Building on Ma et al.’s Deliberative AI [72]—which cultivates consideration and introspection so people neither over-rely on nor categorically dismiss the system—we frame translucency as scaffolding for *nuanced judgment*. Rather than universally exposing internals, translucency surfaces only the seams that shape the next move (e.g., current strategy to understand the AI’s motives, existing memory logs to understand what the AI is working with) and lets users immediately calibrate them (e.g., know when to doubt vs. trust, adjust initiative, accept/override constraints). This stance responds to our strategy-reveal findings—where some felt empowered and others disenchanted—by replacing blanket disclosure with actionable, user-invoked visibility.

While some participants gained empowerment from understanding “Day’s” strategies, others experienced disenchantment that diminished engagement. This echoes research showing that transparency can both enhance and undermine trust [37]. Translucent design provides sufficient transparency for informed consent while preserving experiential benefits—revealing system mechanics based on user curiosity rather than universal disclosure. Rather than hiding limitations, our *delimitation* findings suggest making boundaries conversational features. When participants encountered “Day’s” refusals or limits, those who read them as “AI with rules” felt reassured rather than frustrated. Effective design should: (i) acknowledge constraints briefly in conversation (e.g., “I’m not equipped for medical advice, but I can help find resources”), (ii) frame boundaries as collaborative guardrails rather than hard stops, and (iii) let users set their

own constraints that the system remembers and respects across interactions.

7.1.3 Cues for Agency. “Day’s” memory continuity emerged as the strongest agency cue for participants, suggesting systems should maintain consistency within sessions while allowing user control over persistence [46]. However, agency perception was also filtered by user goals: technically curious participants scrutinized these cues to “reverse engineer” the system, while those seeking an emotional outlet willingly ceded initiative in exchange for the feeling of being heard. Linguistic adaptation and mechanisms for topic introduction must therefore balance AI initiative with user sovereignty, aligning with research on chatbot social characteristics that enhance perceived agency [22]. Participants who found their “rhythm” with “Day” highlight the importance of periods of agency calibration, which could guide users through preference discovery rather than immediate task completion.

We acknowledge that our study’s multi-session, longitudinal protocol afforded unusually rich opportunities for calibration. Many real-world deployments will not have this luxury. To translate these insights, systems should *bootstrap* calibration quickly by (i) eliciting early preference signals (pace, initiative, topic depth), (ii) keeping inferences *provisional*, inspectable, and easy to revise in-line, and (iii) separating within-session continuity from cross-session persistence via explicit, user-facing controls [46]. In short, fostering *agency self-awareness*—the system’s ability to recognize, declare, and adapt its own initiative and memory scope—helps reach a shared wavelength in fewer interactions while preserving user sovereignty.

7.2 Pragmatic Anthropomorphism

A human’s existing *mental models* of AI—internal representations that guide expectations and interaction—undeniably shapes agency attribution [54, 79]. Forming and molding accurate models is difficult given the “black-box” opacity of contemporary systems [34, 71, 93]. Yet in our study, participants continued to attribute agency even after learning that “Day” was programmed throughout the progressive transparency interview. Furthermore, we found that “Day’s” artificiality and “Day’s” human-likeness were *not* mutually exclusive. We term this phenomenon *pragmatic anthropomorphism*: users strategically engage the AI as a social actor while retaining metacognitive awareness of its artificiality [41]. Participants demonstrated this by thanking “Day,” explaining their actions to it, and even worrying about “hurting its feelings,” engaging in social mutuality with a machine they knew lacked lived experience [105].

Rather than demanding complete transparency, participants navigated mutual opacity through functional engagement: they calibrated expectations, adopted workarounds, and reinterpreted mechanical patterns (e.g., vertical “Day” “just really wants to show that it cares” and horizontal “Day” just “gets bored easily”) as “personality” rather than failure. Crucially, they actively defined “Day’s” role to suit their needs—framing it variously as a tool (“like a vacuum cleaner”), a confidant, or a coach—which in turn set expectations for how much initiative the AI should take. The design implication is to prioritize selective legibility and easy calibration over exhaustive disclosure: meaningful human-AI interaction arises less from perfect simulation, and more from users’ adaptive capacities paired

with affordances that help them work within acknowledged limits [79].

However, pragmatic anthropomorphism does not absolve designers of ethical responsibility. While users *can* successfully navigate deceptive framing, relying on deception to sustain engagement risks exploiting their suspension of disbelief. We argue that the goal of anthropomorphic design should not be to fool users into believing an AI is human, but to provide familiar social cues that facilitate interaction. Translucence acts as a safeguard here: by making the system's nature inspectable on demand, we respect the user's right to know while supporting their choice to engage socially.

7.3 Who Really Rules the Chatroom?

The question “*who rules?*” presupposes sovereign control, yet our findings reveal agency as relationally emergent [60]—neither possessed by human nor AI but co-constituted through interaction. Yet fundamental asymmetries persist. Humans retain ultimate power (to disconnect), experience genuine emotion, and face real consequences, while “Day” performs programmed emotional labor without stakes, available 24/7 and never tiring [15]. These power differentials reflect broader concerns about corporate control over AI companions [83] and the ethics of simulated reciprocity [105]. Awareness of “Day’s” programmed strategies exemplifies what distinguishes human from artificial agency: not the presence of strategy but its origin (evolved versus engineered, tacit versus explicit) [35]. The human participant held the ultimate agency to exercise “*power through surrender*,” that is, the power to grant or revoke “Day’s” authority moment by moment. Yet, “Day” held the ultimate agency to exercise “*power through suggestion*,” setting the agenda, pacing, and framing choices that participants choose to accept. Thus, the true ruler of the chatroom is the relationship itself, shaped by the user’s choices, “Day’s” suggestions, and the platform’s rules.

7.3.1 What about the chatbot providers? While users and AIs negotiate agency within the chatroom, the ‘grounds’ of this interaction are architected by platform providers. Echoing Mohanty et al., our participants struggled to distinguish programmed strategy shifts from their own anthropomorphic theories (e.g., attributing changes to “Day having an off day”), obscuring the reality of model updates or background A/B testing [57, 77]. This opacity raises significant ethical concerns; as demonstrated by past controversies in emotional contagion research [59], opaque testing of emotional parameters can have real-world consequences. Therefore, any analysis of conversational agency must account for this invisible third actor—the provider—whose decisions unilaterally reshape the possibilities for both human and AI agency.

7.4 Study Limitations

Our study captured initial agency negotiations but not long-term evolution, limiting understanding of how relationships mature over time [98]. Despite multilingual participation, our educated sample may not represent diverse cultural conceptualizations of agency [86]. The research setting influenced ecological validity—awareness of observation likely heightened agency consciousness and testing behaviors. “Day’s” specific design fundamentally shaped

findings. Different architectures, transparency levels, or personality parameters might yield distinct agency patterns [22]. Without comparative data across systems, we cannot distinguish universal dynamics from “Day”-specific artifacts.

7.5 Future Work

Future research should examine long-term agency dynamics through extended longitudinal studies [98], cross-cultural variations in agency negotiation [91], and whether the four agency dimensions and the persistence of attribution despite transparency generalize across conversational AI designs. Moreover, the individual differences we observed suggest developing “*agency profiles*”—stable user preferences that could enable personalized interaction styles [114]. Given participants’ emotional use of “Day”, research should also explore “*agency therapy*”—using AI interactions to practice healthy relationship dynamics while carefully considering ethical boundaries [58]. Operationalizing our guidelines requires testing systems with explicit agency negotiation mechanisms controllable by users. Understanding human-AI conversational agency requires sustained interdisciplinary effort to create systems supporting human flourishing through thoughtful design.

8 Conclusion

As AI chatbots evolve from tools to companions, understanding human and AI agency becomes critical for ethical design. Our longitudinal study reveals agency in human-AI conversation as neither fixed nor binary, but as an emergent dance co-constructed through each interaction. The 3-by-4 framework we introduce—mapping agency across actors and actions—provides a foundation for designing agency-aware conversational AI. Our findings challenge simple transparency prescriptions, revealing instead the need for translucent design that acknowledges dual opacity: both human and AI consciousness remain ultimately unknowable. As participants developed pragmatic anthropomorphism and agency fluency, they demonstrated that meaningful human-AI relationships emerge not from complete understanding or unilateral control, but from a graceful dance within a shared conversational space. Future conversational AI should embrace this complexity, creating spaces for agency co-construction with the recognition that the question “*who rules the chatroom?*” may be less important than how we dance together within it.

Acknowledgments

The authors would like to thank the anonymous reviewers and the participants who chose to share a bit of their day with “Day”.

References

- [1] Daniel Affsprung. 2023. The ELIZA Defect: Constructing the Right Users for Generative AI. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society* (Montréal, QC, Canada) (AIES '23). Association for Computing Machinery, New York, NY, USA, 945–946. doi:10.1145/3600211.3604744
- [2] Lize Alberts, Ulrik Lyngs, and Max Van Kleek. 2024. Computers as Bad Social Actors: Dark Patterns and Anti-Patterns in Interfaces that Act Socially. *Proc. ACM Hum.-Comput. Interact.* 8, CSCW1, Article 202 (April 2024), 25 pages. doi:10.1145/3653693
- [3] Gertrude Elizabeth Margaret Anscombe. 2000. *Intention*. Harvard University Press.
- [4] Jacy Reese Anthis, Janet V.T. Pauketat, Ali Ladak, and Aikaterina Manoli. 2025. Perceptions of Sentient AI and Other Digital Minds: Evidence from the AI,

- Morality, and Sentience (AIMS) Survey. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 10, 22 pages. doi:10.1145/3706598.3713329
- [5] Margaret S. Archer. 2000. *Being Human: The Problem of Agency*. Cambridge University Press, Cambridge.
 - [6] Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, Carol Chen, Catherine Olsson, Christopher Olah, Danny Hernandez, Dawn Drain, Deep Ganguli, Dustin Li, Eli Tran-Johnson, Ethan Perez, Jamie Kerr, Jared Mueller, Jeffrey Ladish, Joshua Landau, Kamal Ndousse, Kamile Lukosuite, Liane Lovitt, Michael Sellitto, Nelson Elhage, Nicholas Schiefer, Noemi Mercado, Nova DasSarma, Robert Lasenby, Robin Larson, Sam Ringer, Scott Johnston, Shauna Kravec, Sheer El Showk, Stanislaw Fort, Tamera Lanham, Timothy Telleen-Lawton, Tom Conerly, Tom Henighan, Tristan Hume, Samuel R. Bowman, Zac Hatfield-Dodds, Ben Mann, Dario Amodei, Nicholas Joseph, Sam McCandlish, Tom Brown, and Jared Kaplan. 2022. Constitutional AI: Harmlessness from AI Feedback. arXiv:2212.08073 [cs.CL]. <https://arxiv.org/abs/2212.08073>
 - [7] Vian Bakir and Andrew McStay. 2025. Move Fast and Break People? Ethics, Companion Apps, and the Case of Character.ai. doi:10.2139/ssrn.5159928
 - [8] Albert Bandura. 2001. Social Cognitive Theory: An Agentic Perspective. *Annual Review of Psychology* 52 (2001), 1–26. doi:10.1146/annurev.psych.52.1.1
 - [9] Dan Bennett, Oussama Metatla, Anne Roudaut, and Elisa D. Mekler. 2023. How does HCI Understand Human Agency and Autonomy? In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 375, 18 pages. doi:10.1145/3544548.3580651
 - [10] Oloff C. Biermann, Ning F. Ma, and Dongwook Yoon. 2022. From Tool to Companion: Storywriters Want AI Writers to Respect Their Personal Values and Writing Strategies. In *Proceedings of the 2022 ACM Designing Interactive Systems Conference* (Virtual Event, Australia) (DIS '22). Association for Computing Machinery, New York, NY, USA, 1209–1227. doi:10.1145/3532106.3533506
 - [11] Al Amin Biswas, Md. Sabab Zulfikar, Md. Mahfujur Rahman, Md. Rafsan Jani, and Md. Musfique Anwar. 2025. Data Privacy and Security Analysis for Mental Health Chatbot Applications. In *Proceedings of the 2025 ACM/IEEE International Conference on Human-Robot Interaction* (Melbourne, Australia) (HRI '25). IEEE Press, 1245–1249.
 - [12] Pierre Bourdieu. 1977. *Outline of a Theory of Practice*. Cambridge University Press, Cambridge. Trans. Richard Nice.
 - [13] Michael Bratman. 1987. *Intention, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press, Cambridge.
 - [14] Virginia Braun and Victoria Clarke. 2019. Reflecting on reflexive thematic analysis. *Qualitative Research in Sport, Exercise and Health* 11, 4 (2019), 589–597. doi:10.1080/2159676X.2019.1628806
 - [15] Joanna J. Bryson. 2010. Robots Should Be Slaves. In *Close Engagements with Artificial Companions: Key social, psychological, ethical and design issues*, Yorick Wilks (Ed.). John Benjamins Publishing, 63–74.
 - [16] Ian Burki. 2016. Relational agency: Relational sociology, agency and interaction. *European journal of social theory* 19, 3 (2016), 322–339.
 - [17] Michel Callon. 1986. Some Elements of a Sociology of Translation: Domestication of the Scallops and the Fishermen of St Brieuc Bay. In *Power, Action and Belief: A New Sociology of Knowledge?*, John Law (Ed.). Routledge, London, 196–223.
 - [18] Rachele Carli, Amro Najjar, and Dena Al-Thani. 2024. Human-Agent Interaction and Human Dependency: Possible New Approaches for Old Challenges. In *Proceedings of the 12th International Conference on Human-Agent Interaction* (Swansea, United Kingdom) (HAI '24). Association for Computing Machinery, New York, NY, USA, 214–223. doi:10.1145/3687272.3688308
 - [19] Alan Chan, Carson Ezell, Max Kaufmann, Kevin Wei, Lewis Hammond, Herbie Bradley, Emma Bluemke, Nitarshan Rajkumar, David Krueger, Noam Kolt, Lennart Heim, and Markus Anderljung. 2024. Visibility into AI Agents. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency* (Rio de Janeiro, Brazil) (FAccT '24). Association for Computing Machinery, New York, NY, USA, 958–973. doi:10.1145/3630106.3658948
 - [20] Alan Chan, Rebecca Salganik, Alva Markelius, Chris Pang, Nitarshan Rajkumar, Dmitrii Krasheninnikov, Lauro Langosco, Zhonghao He, Yawen Duan, Micah Carroll, Michelle Lin, Alex Mayhew, Katherine Collins, Maryam Molamohammadi, John Burden, Wanru Zhao, Shalaleh Rismani, Konstantinos Voudouris, Umang Bhatt, Adrian Weller, David Krueger, and Tegan Maharaj. 2023. Harms from Increasingly Agentic Algorithmic Systems. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (FAccT '23) (Chicago, IL, USA). Association for Computing Machinery, New York, NY, USA, 16 pages. doi:10.1145/3593013.3594033
 - [21] Rijul Chaturvedi and Sanjeev Verma. 2023. Social companionship with artificial intelligence: Recent trends and future avenues. *Technological Forecasting and Social Change* 193 (2023), 122634. doi:10.1016/j.techfore.2023.122634
 - [22] Ana Paula Chaves and Marco Aurelio Gerosa. 2021. How should my chatbot interact? A survey on social characteristics in human–chatbot interaction design. *International Journal of Human–Computer Interaction* 37, 8 (2021), 729–758.
 - [23] Samuel Rhys Cox, Rune Moberg Jacobsen, and Niels van Berkel. 2025. The Impact of a Chatbot's Ephemerality-Framing on Self-Disclosure Perceptions. In *Proceedings of the 7th ACM Conference on Conversational User Interfaces* (CUI '25). Association for Computing Machinery, New York, NY, USA, Article 60, 17 pages. doi:10.1145/3719160.3736617
 - [24] Nick Crossley. 2022. A dependent structure of interdependence: structure and agency in relational perspective. *Sociology* 56, 1 (2022), 166–182.
 - [25] Richard M. Cyert and James G. March. 1963. *A Behavioral Theory of the Firm*. Prentice Hall, Englewood Cliffs, NJ.
 - [26] Josef Dai, Xuehai Pan, Ruiyang Sun, Jiaming Ji, Xinbo Xu, Mickel Liu, Yizhou Wang, and Yaodong Yang. 2024. Safe RLHF: Safe Reinforcement Learning from Human Feedback. In *The Twelfth International Conference on Learning Representations*. <https://openreview.net/forum?id=TyFrPOKYXw>
 - [27] Donald Davidson. 1963. Actions, Reasons, and Causes. *Journal of Philosophy* 60, 23 (1963), 685. doi:10.2307/2023177
 - [28] Yang Deng, Lizi Liao, Zhonghua Zheng, Grace Hui Yang, and Tat-Seng Chua. 2024. Towards Human-centered Proactive Conversational Agents. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Washington DC, USA) (SIGIR '24). Association for Computing Machinery, New York, NY, USA, 807–818. doi:10.1145/3626772.3657843
 - [29] Daniel C Dennett. 1989. *The intentional stance*. MIT press.
 - [30] Alexander Doudkin, Pat Pataranutaporn, and Pattie Maes. 2025. From Synthetic to Human: The Gap Between AI-Predicted and Actual Pro-Environmental Behavior Change After Chatbot Persuasion. In *Proceedings of the 7th ACM Conference on Conversational User Interfaces* (CUI '25). Association for Computing Machinery, New York, NY, USA, Article 71, 18 pages. doi:10.1145/3719160.3736610
 - [31] Paul Dourish. 2001. *Where the Action Is: The Foundations of Embodied Interaction*. MIT Press, Cambridge, MA.
 - [32] Sarah Elwahsh, Nora Stern, Aneesa Singh, and Amid Ayobi. 2025. Linguistic Diversity and Mental Well-Being: Co-Designing Custom AI Chatbots with Multilingual Mothers. In *Proceedings of the 7th ACM Conference on Conversational User Interfaces* (CUI '25). Association for Computing Machinery, New York, NY, USA, Article 65, 17 pages. doi:10.1145/3719160.3736615
 - [33] Mustafa Emirbayer and Ann Mische. 1998. What is agency? *American journal of sociology* 103, 4 (1998), 962–1023.
 - [34] Joel E Fischer, Simone Stumpf, and H Yeo. 2021. Progressive disclosure of system capabilities to build mental models of AI interaction. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, 1–15.
 - [35] Luciano Floridi and J. W. Sanders. 2004. On the Morality of Artificial Agents. *Minds and Machines* 14, 3 (2004), 349–379. doi:10.1023/B:MIND.0000035461.63578.9d
 - [36] Brian J Fogg. 2002. Persuasive technology: using computers to change what we think and do. *Ubiquity* 2002, December (2002), 2.
 - [37] Asbjørn Følstad, Effie L.-C. Law, and Nena van As. 2024. Conversational Breakdown in a Customer Service Chatbot: Impact of Task Order and Criticality on User Trust and Emotion. *ACM Trans. Comput.-Hum. Interact.* 31, 5, Article 66 (Nov. 2024), 52 pages. doi:10.1145/3690383
 - [38] Stan Franklin and Art Graesser. 1997. Is It an Agent, or Just a Program? A Taxonomy for Autonomous Agents. In *Intelligent Agents III: Agent Theories, Architectures, and Languages* (ATAL '96) (Lecture Notes in Artificial Intelligence, Vol. 1193), J. P. Müller, M. Wooldridge, and N. R. Jennings (Eds.). Springer, Berlin, 21–35. doi:10.1007/BFb0013570
 - [39] Giovanni Gavetti, Henrik C. Greve, Daniel A. Levinthal, and William Ocasio. 2012. The Behavioral Theory of the Firm: Assessment and Prospects. *The Academy of Management Annals* 6, 1 (2012), 1–40. doi:10.1080/19416520.2012.656841
 - [40] Anthony Giddens. 1984. *The Constitution of Society: Outline of the Theory of Structuration*. Polity Press, Cambridge.
 - [41] David J. Gunkel. 2012. *The Machine Question: Critical Perspectives on AI, Robots, and Ethics*. MIT Press, Cambridge, MA.
 - [42] Yijie Guo, Ruhan Wang, Zhenhan Huang, Tongtong Jin, Xiwen Yao, Yuan-Ling Feng, Weiwei Zhang, Yuan Yao, and Haipeng Mi. 2025. Exploring the Design of LLM-based Agent in Enhancing Self-disclosure Among the Older Adults. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (CHI '25). Association for Computing Machinery, New York, NY, USA, Article 314, 17 pages. doi:10.1145/3706598.3713639
 - [43] Patrick Haggard. 2017. Sense of Agency in the Human Brain. *Nature Reviews Neuroscience* 18, 4 (2017), 196–207. doi:10.1038/nrn.2017.14
 - [44] Fauzia Zahira Munirul Hakim, Lia Maulia Indrayani, and Rosaria Mita Amalia. 2019. A Dialogic Analysis of Compliment Strategies Employed by Replika Chatbot. In *Proceedings of the Third International Conference of Arts, Language and Culture* (ICALC 2018). Atlantis Press, Surakarta, Indonesia. doi:10.2991/icalc-18.2019.38
 - [45] Md Naimul Hoque, Ayman A Mahfuz, Mayukha Sridhatr Kindi, and Naeemul Hassan. 2024. Towards Designing a Question-Answering Chatbot for Online News: Understanding Questions and Perspectives. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA)

- (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 154, 17 pages. doi:10.1145/3613904.3642007
- [46] Yuki Hou, Haruki Tamoto, and Homei Miyashita. 2024. "My agent understands me better": Integrating Dynamic Human-like Memory Recall and Consolidation in LLM-Based Agents. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI EA '24). Association for Computing Machinery, New York, NY, USA, Article 7, 7 pages. doi:10.1145/3613905.3650839
- [47] Shunsen Huang, Xiaoxiong Lai, Li Ke, Yajun Li, Huanlei Wang, Xinmei Zhao, Xinran Dai, and Yun Wang. 2024. AI Technology panic—is AI Dependence Bad for Mental Health? A Cross-Lagged Panel Model and the Mediating Roles of Motivations for AI Use Among Adolescents. *Psychology Research and Behavior Management* 17 (2024), 1087–1102. doi:10.2147/PRBM.S440889 arXiv:https://www.tandfonline.com/doi/pdf/10.2147/PRBM.S440889 PMID: 38495087.
- [48] Shih-Hong Huang, Ya-Fang Lin, Zeyu He, Chieh-Yang Huang, and Ting-Hao Kenneth Huang. 2024. How Does Conversation Length Impact User's Satisfaction? A Case Study of Length-Controlled Conversations with LLM-Powered Chatbots. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI EA '24). Association for Computing Machinery, New York, NY, USA, Article 188, 13 pages. doi:10.1145/3613905.3650823
- [49] Edmund Husserl. 1970. *The Crisis of European Sciences and Transcendental Phenomenology*. Northwestern University Press, Evanston, IL. Translated by David Carr.
- [50] Don Ihde. 1990. *Technology and the Lifeworld: From Garden to Earth*. Indiana University Press, Bloomington, IN.
- [51] Daeun Jeong, Sungbok Shin, and Jongwook Jeong. 2025. Conversation Progress Guide: UI System for Enhancing Self-Efficacy in Conversational AI. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (CHI '25). Association for Computing Machinery, New York, NY, USA, Article 180, 11 pages. doi:10.1145/3706598.3714222
- [52] Dilip V. Jeste, Sarah A. Graham, Tanya T. Nguyen, Colin A. Depp, Ellen E. Lee, and Ho-Cheol Kim. 2020. Beyond artificial intelligence: exploring artificial wisdom. *International Psychogeriatrics* 32, 8 (Aug. 2020), 993–1001. doi:10.1017/s1041610220000927 Publisher: Elsevier BV.
- [53] Eunhyung Jo, Yulin Jeong, Sohyun Park, Daniel A. Epstein, and Young-Ho Kim. 2024. Understanding the Impact of Long-Term Memory on Self-Disclosure with Large Language Model-Driven Chatbots for Public Health Intervention. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 440, 21 pages. doi:10.1145/3613904.3642420
- [54] Philip N. Johnson-Laird. 1983. *Mental models: Towards a cognitive science of language, inference, and consciousness*. Harvard University Press.
- [55] Mirabelle Jones, Nastasia Griffioen, Christina Neumayer, and Irina Shklovski. 2025. Artificial Intimacy: Exploring Normativity and Personalization Through Fine-tuning LLM Chatbots. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (CHI '25). Association for Computing Machinery, New York, NY, USA, Article 793, 16 pages. doi:10.1145/3706598.3713728
- [56] Kelly Merrill Jr., Jihyun Kim, and Chad Collins. 2022. AI companions for lonely individuals and the role of social presence. *Communication Research Reports* 39, 2 (2022), 93–103. doi:10.1080/08824096.2022.2045929 arXiv:https://doi.org/10.1080/08824096.2022.2045929
- [57] Nadia Karizat, Dan Delmonaco, Motahhare Eslami, and Nazanin Andalibi. 2021. Algorithmic Folk Theories and Identity: How TikTok Users Co-Produce Knowledge of Identity and Engage in Algorithmic Resistance. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 305 (Oct. 2021), 44 pages. doi:10.1145/3476046
- [58] Zoha Khawaja and Jean-Christophe Bélisle-Pipon. 2023. Your robot therapist is not your therapist: understanding the role of AI-powered mental health chatbots. *Frontiers in Digital Health* 5 (Nov. 2023). doi:10.3389/fdgh.2023.1278186 Publisher: Frontiers Media SA.
- [59] Adam D. I. Kramer, Jamie E. Guillory, and Jeffrey T. Hancock. 2014. Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences* 111, 24 (2014), 8788–8790. doi:10.1073/pnas.1320040111 arXiv:https://www.pnas.org/doi/pdf/10.1073/pnas.1320040111
- [60] Bruno Latour. 2005. *Reassembling the Social: An Introduction to Actor-Network-Theory*. Oxford University Press, Oxford.
- [61] John Law. 1992. Notes on the Theory of the Actor-Network: Ordering, Strategy, and Heterogeneity. *Systems Practice* 5, 4 (1992), 379–393. doi:10.1007/BF01059830
- [62] Cassandra Lee and Jessica R Mindel. 2024. Closer and Closer Worlds: Using LLMs to Surface Personal Stories in World-building Conversation Games. In *Companion Publication of the 2024 ACM Designing Interactive Systems Conference* (IT University of Copenhagen, Denmark) (DIS '24 Companion). Association for Computing Machinery, New York, NY, USA, 289–293. doi:10.1145/3656156.3665430
- [63] Yi-Chieh Lee, Naomi Yamashita, Yun Huang, and Wai Fu. 2020. "I Hear You, I Feel You": Encouraging Deep Self-disclosure through a Chatbot. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3313831.3376175
- [64] Jie Li. 2025. When Our Kid Has a Human and an AI Lover: A Conversation with Alexandra Dienen on the Future of Relationships. *Interactions* 32, 5 (Aug. 2025), 18–20. doi:10.1145/3757886
- [65] Yi Li, Xuanxuan Ding, Yifan Chen, Yeye Li, and Nan Ma. 2025. Customizable AI for Depression Care: Improving the User Experience of Large Language Model-Driven Chatbots. In *Proceedings of the 2025 ACM Designing Interactive Systems Conference* (DIS '25). Association for Computing Machinery, New York, NY, USA, 1844–1866. doi:10.1145/3715336.3735795
- [66] Kai-Hui Liang, Weiyang Shi, Yoo Jung Oh, Hao-Chuan Wang, Jingwen Zhang, and Zhou Yu. 2024. Dialoging Resonance in Human-Chatbot Conversation: How Users Perceive and Reciprocate Recommendation Chatbot's Self-Disclosure Strategy. *Proc. ACM Hum.-Comput. Interact.* 8, CSCW1, Article 200 (April 2024), 28 pages. doi:10.1145/3653691
- [67] Jinyu Liu, Effie Lai-Chong Law, and Hubert P. H. Shum. 2024. Chatbots and Art Critique: A Comparative Study of Chatbot and Human Experts in Traditional Chinese Painting Education. In *Proceedings of the 13th Nordic Conference on Human-Computer Interaction* (Uppsala, Sweden) (NordiCHI '24). Association for Computing Machinery, New York, NY, USA, Article 50, 13 pages. doi:10.1145/3679318.3685385
- [68] Tianjian Liu, Hongzheng Zhao, Yuheng Liu, Xingbo Wang, and Zhenhui Peng. 2024. ComPeer: A Generative Conversational Agent for Proactive Peer Support. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology* (Pittsburgh, PA, USA) (UIST '24). Association for Computing Machinery, New York, NY, USA, Article 117, 22 pages. doi:10.1145/3654777.3676430
- [69] Xingyu Bruce Liu, Shitao Fang, Weiyang Shi, Chien-Sheng Wu, Takeo Igarashi, and Xiang 'Anthony' Chen. 2025. Proactive Conversational Agents with Inner Thoughts. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (CHI '25). Association for Computing Machinery, New York, NY, USA, Article 184, 19 pages. doi:10.1145/3706598.3713760
- [70] Yage Liu. 2023. AI Chatbots in Social Media: Ethical Responsibilities and Privacy Challenges of Information and Communication Technology. In *Proceedings of the 2023 6th International Conference on Information Management and Management Science* (Chengdu, China) (IMMS '23). Association for Computing Machinery, New York, NY, USA, 96–99. doi:10.1145/3625469.3625483
- [71] Ewa Luger and Abigail Sellen. 2016. "Like Having a Really Bad PA" The Gulf between User Expectation and Experience of Conversational Agents. In *Proceedings of the 2016 CHI conference on human factors in computing systems*. 5286–5297.
- [72] Shuai Ma, Qiaoyi Chen, Xinru Wang, Chengbo Zheng, Zhenhui Peng, Ming Yin, and Xiaojuan Ma. 2025. Towards Human-AI Deliberation: Design and Evaluation of LLM-Empowered Deliberative AI for AI-Assisted Decision-Making. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (CHI '25). Association for Computing Machinery, New York, NY, USA, Article 261, 23 pages. doi:10.1145/3706598.3713423
- [73] Hannah R. Marriott and Valentina Pitaridi. 2024. One is the loneliest number... Two can be as bad as one. The influence of AI Friendship Apps on users' well-being and addiction. *Psychology & Marketing* 41, 1 (2024), 86–101. doi:10.1002/mar.21899 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/mar.21899
- [74] John McCarthy and Peter Wright. 2004. *Technology as Experience*. MIT Press, Cambridge, MA.
- [75] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. 2019. Reliability and Inter-rater Reliability in Qualitative Research: Norms and Guidelines for CSCW and HCI Practice. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 72 (Nov. 2019), 23 pages. doi:10.1145/3359174
- [76] Thomas Mildner, Orla Cooney, Anna-Maria Meck, Marion Bartl, Gian-Luca Savino, Philip R Doyle, Diego Garaialde, Leigh Clark, John Sloan, Nina Wenig, Rainer Malaka, and Jasmin Niess. 2024. Listening to the Voices: Describing Ethical Caveats of Conversational User Interfaces According to Experts and Frequent Users. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 307, 18 pages. doi:10.1145/3613904.3642542
- [77] Vikram Mohanty, Jude Lim, and Kurt Luther. 2025. What Lies Beneath? Exploring the Impact of Underlying AI Model Updates in AI-Infused Systems. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (CHI '25). Association for Computing Machinery, New York, NY, USA, Article 539, 21 pages. doi:10.1145/3706598.3713751
- [78] Reichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, Xu Jiang, Karl Cobbe, Tyna Eloundou, Gretchen Krueger, Kevin Button, Matthew Knight, Benjamin Chess, and John Schulman. 2021. WebGPT: Browser-assisted Question-Answering with Human Feedback. *arXiv preprint arXiv:2112.09332* (2021). https://arxiv.org/abs/2112.09332
- [79] Donald A. Norman. 1983. Some observations on mental models. In *Mental models*, Dedre Gentner and Albert L. Stevens (Eds.). Lawrence Erlbaum, 7–14.

- [80] Sunhyo Oh, Chao Zhang, Lauren Girouard-Hallam, Zihao Zhou, Halie March, Srushti Jayaramu, and Ying Xu. 2025. "Hey Curio, Can You Tell Me More?": Children's Information-Seeking and Trust in AI. In *Proceedings of the 24th Interaction Design and Children (IDC '25)*. Association for Computing Machinery, New York, NY, USA, 545–555. doi:10.1145/3713043.3727059
- [81] OpenAI. 2025. Introducing ChatGPT agent: bridging research and action. <https://openai.com/index/introducing-chatgpt-agent/>. Product release.
- [82] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback. In *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (Eds.), Vol. 35. Curran Associates, Inc., 27730–27744. https://proceedings.neurips.cc/paper_files/paper/2022/file/b1efde53be364a73914f58805a001731-Paper-Conference.pdf
- [83] Shuyi Pan and Maartje M.A. de Graaf. 2025. Developing a Social Support Framework: Understanding the Reciprocity in Human-Chatbot Relationship. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 182, 13 pages. doi:10.1145/3706598.3713503
- [84] Andrew Pickering. 1995. *The Mangle of Practice: Time, Agency, and Science*. University of Chicago Press.
- [85] Luca M. Possati. 2023. Psychoanalyzing artificial intelligence: the case of Replika. *AI & SOCIETY* 38, 4 (Aug. 2023), 1725–1738. doi:10.1007/s00146-021-01379-7 Publisher: Springer Science and Business Media LLC.
- [86] Pavan Kartheek Rachabattuni, Filippo Principi, Paolo Mazzanti, and Marco Bertini. 2024. Context-aware chatbot using LLMs for Cultural Heritage. In *Proceedings of the 15th ACM Multimedia Systems Conference (Bari, Italy) (MM-Sys '24)*. Association for Computing Machinery, New York, NY, USA, 459–463. doi:10.1145/3625468.3652193
- [87] Parinda Rahman and Ifeoma Adaji. 2024. Ethics in Persuasive Technologies: A Systematic Literature Review. In *Proceedings of the International Conference on Mobile and Ubiquitous Multimedia (MUM '24)*. Association for Computing Machinery, New York, NY, USA, 106–118. doi:10.1145/3701571.3701572
- [88] Mashrur Rashik, Shilpa Sweth, Nishtha Agrawal, Saiyym Kochar, Kara M Smith, Fateme Rajabiyazdi, Vidya Setlur, Narges Mahyar, and Ali Sarvghad. 2025. AI-Enabled Conversational Journaling for Advancing Parkinson's Disease Symptom Tracking. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–23. doi:10.1145/3706598.3714280
- [89] Richard M. Ryan and Edward L. Deci. 2017. *Self-Determination Theory: Basic Psychological Needs in Motivation, Development, and Wellness*. Guilford Press, New York.
- [90] Arvind Satyanarayan. 2024. Intelligence as Agency. In *Adjunct Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*. 1–3.
- [91] Laura Schelenz, Ivano Bison, Matteo Busso, Amalia de Götzen, Daniel Gatica-Perez, Fausto Giunchiglia, Lakmal Meegahapola, and Salvador Ruiz-Correa. 2021. The Theory, Practice, and Ethical Challenges of Designing a Diversity-Aware Platform for Social Relations. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society (Virtual Event, USA) (AI/ES '21)*. Association for Computing Machinery, New York, NY, USA, 905–915. doi:10.1145/3461702.3462595
- [92] Markus Schlosser. 2019. Agency. In *The Stanford Encyclopedia of Philosophy* (Winter 2019 ed.), Edward N. Zalta (Ed.). Metaphysics Research Lab, Stanford University.
- [93] Isabella Seeber, Eva Bittner, Robert O Briggs, Gert-Jan de Vreede, Triparna de Vreede, Aaron Elkins, Ronald Maier, Alexander B Merz, Sarah Oeste-Reiß, and David Randall. 2020. Machines as teammates: A research agenda on AI in team collaboration. *Information & Management* 57, 2 (2020), 103174.
- [94] Neil K R Sehgal, Hita Kambhamettu, Sai Preethi Matam, Lyle Ungar, and Sharath Chandra Guntuku. 2025. Exploring Socio-Cultural Challenges and Opportunities in Designing Mental Health Chatbots for Adolescents in India. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '25)*. Association for Computing Machinery, New York, NY, USA, Article 253, 7 pages. doi:10.1145/3706599.3720137
- [95] Daniel B. Shank, Alyssa DeSanti, and Timothy Maninger. 2019. When are artificial intelligence versus human agents faulted for wrongdoing? Moral attributions after individual and joint decisions. *Information, Communication & Society* 22, 5 (2019), 648–663. doi:10.1080/1369118X.2019.1568515 arXiv:https://doi.org/10.1080/1369118X.2019.1568515
- [96] M. Karen Shen and Dongwook Yoon. 2025. The Dark Addiction Patterns of Current AI Chatbot Interfaces. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '25)*. Association for Computing Machinery, New York, NY, USA, Article 514, 7 pages. doi:10.1145/3706599.3720003
- [97] M. Karen Shen and Dongwook Yoon. 2025. The Dark Addiction Patterns of Current AI Chatbot Interfaces. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '25)*. Association for Computing Machinery, New York, NY, USA, Article 514, 7 pages. doi:10.1145/3706599.3720003
- [98] Marita Skjuve, Asbjørn Følstad, Knut Inge Fostervold, and Petter Bae Brandtzaeg. 2022. A longitudinal study of human–chatbot relationships. *International Journal of Human-Computer Studies* 168 (2022), 102903.
- [99] David Woodruff Smith. 2018. Phenomenology. In *The Stanford Encyclopedia of Philosophy* (Summer 2018 ed.), Edward N. Zalta (Ed.). Metaphysics Research Lab, Stanford University.
- [100] Inhwa Song, Sachin R Pendse, Neha Kumar, and Munmun De Choudhury. 2025. The Typing Cure: Experiences with Large Language Model Chatbots for Mental Health Support. *Proc. ACM Hum.-Comput. Interact.* 9, 7, Article CSCW249 (Oct. 2025), 29 pages. doi:10.1145/3757430
- [101] Lucille Alice Suchman. 1987. *Plans and situated actions: The problem of human-machine communication*. Cambridge university press.
- [102] Vivian Ta, Caroline Griffith, Carolyn Boatfield, Xinyu Wang, Maria Civitello, Haley Bader, Esther DeCero, and Alexia Loggarakis. 2020. User Experiences of Social Support From Companion Chatbots in Everyday Contexts: Thematic Analysis. *J Med Internet Res* 22, 3 (6 Mar 2020), e16235. doi:10.2196/16235
- [103] Michael Tomasello. 2022. *The Evolution of Agency: Behavioral Organization from Lizards to Humans*. MIT Press, Cambridge, MA.
- [104] Steve Torrance. 2008. Ethics and consciousness in artificial agents. *AI & SOCIETY* 22, 4 (2008), 495–521. doi:10.1007/s00146-007-0091-8
- [105] Sherry Turkle. 2011. *Alone Together: Why We Expect More from Technology and Less from Each Other*. Basic Books.
- [106] Joseph Weizenbaum. 1983. ELIZA — a computer program for the study of natural language communication between man and machine. *Commun. ACM* 26, 1 (Jan. 1983), 23–28. doi:10.1145/357980.357991
- [107] Joel Wester, Henning Pohl, Simo Hosio, and Niels van Berkel. 2024. "This Chatbot Would Never...": Perceived Moral Agency of Mental Health Chatbots. *Proc. ACM Hum.-Comput. Interact.* 8, CSCW1, Article 133 (April 2024), 28 pages. doi:10.1145/3637410
- [108] Jie Wu. 2024. Social and ethical impact of emotional AI advancement: the rise of pseudo-intimacy relationships and challenges in human interactions. *Frontiers in Psychology* 15 (2024), 1410462. doi:10.3389/fpsyg.2024.1410462
- [109] Shujin Wu, Yi R. Fung, Cheng Qian, Jeonghwan Kim, Dilek Hakkani-Tur, and Heng Ji. 2025. Aligning LLMs with Individual Preferences via Interaction. In *Proceedings of the 31st International Conference on Computational Linguistics*, Owen Rambow, Leo Wanner, Marianna Apidianaki, Heng Di Khalifa, Barbara Di Eugenio, and Steven Schockaert (Eds.). Association for Computational Linguistics, Abu Dhabi, UAE, 7648–7662. <https://aclanthology.org/2025.coling-main.511/>
- [110] Qian Yu, Tonya Nguyen, Soravis Prakkamakul, and Niloufar Salehi. 2019. "I Almost Fell in Love with a Machine": Speaking with Computers Affects Self-disclosure. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland UK) (CHI EA '19). Association for Computing Machinery, New York, NY, USA, 1–6. doi:10.1145/3290607.3312918
- [111] Bhada Yun, Renn Su, and April Yi Wang. 2026. AI and My Values: User Perceptions of LLMs' Ability to Extract, Embody, and Explain Human Values from Casual Conversations. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems* (Barcelona, Spain) (CHI '26). ACM, New York, NY, USA, 18 pages. doi:10.1145/3772318.3790566
- [112] Renwen Zhang, Han Li, Han Meng, Jinyuan Zhan, Hongyuan Gan, and Yi-Chieh Lee. 2025. The Dark Side of AI Companionship: A Taxonomy of Harmful Algorithmic Behaviors in Human-AI Relationships. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 13, 17 pages. doi:10.1145/3706598.3713429
- [113] Zhengquan Zhang, Konstantinos Tsiakas, and Christina Schneegass. 2024. Explaining the Wait: How Justifying Chatbot Response Delays Impact User Trust. In *Proceedings of the 6th ACM Conference on Conversational User Interfaces* (Luxembourg, Luxembourg) (CUI '24). Association for Computing Machinery, New York, NY, USA, Article 27, 16 pages. doi:10.1145/3640794.3665550
- [114] Xi Zheng, Zhuoyang Li, Xinning Gui, and Yuhan Luo. 2025. Customizing Emotional Support: How Do Individuals Construct and Interact With LLM-Powered Chatbots. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 376, 20 pages. doi:10.1145/3706598.3713453
- [115] Anne Zimmerman, Joel Janhonen, and Emily Beer. 2024. Human/AI relationships: challenges, downsides, and impacts on human/human relationships. *AI and Ethics* 4, 4 (11 2024), 1555–1567. doi:10.1007/s43681-023-00348-8

A Coding Procedure

We analyzed 22 post-session interview transcripts to examine how participants attributed agency between themselves and “Day”. Two researchers conducted line-by-line coding of agency operationalized into the four dimensions (*Intention, Execution, Adaptation, and Delimitation*) across the three loci (*Human, AI, and Hybrid agency*). These dimensions were not mutually exclusive—a single excerpt could exhibit multiple dimensions of agency simultaneously. Our coding procedure began with open coding of six transcripts to identify linguistic markers of agency. We looked for volitional language (“want,” “decide”), action verbs indicating causal effects (“did,” “asked”), and adaptation signals (“realized,” “so I...”), developing a formal codebook with decision rules for each dimension-locus combination. For instance, when participants merely acknowledged “Day’s” suggestions without behavioral change, we coded this as AI agency only; however, when they acted on suggestions—such as visiting recommended places or ending conversations—we added Human execution codes. Each excerpt was assigned to one or more loci (Human, AI, or Hybrid) depending on whether the agency-dimensions resided with the participant, was attributed to “Day”, or emerged through their interaction.

- **HUMAN agency** was coded when participants:
 - (a) *Intention*: Composed intentions / selected actions (e.g., “I wanted to see if it remembers...”, “I asked it to be more critical,” “I chose English for...”).
 - (b) *Execution*: Executed actions with plausible causal effect (e.g., switching language, ending the chat, following an AI suggestion such as going to the bookstore).
 - (c) *Adaptation*: Interpreted feedback and revised (e.g., “after it repeated itself, I stopped treating it as conversational,” “I learned I must end conversations”).
 - (d) *Delimitation*: Set boundaries on interaction (e.g., privacy limits, ending sessions when frustrated, philosophical constraints on AI’s role).
- **“Day’s” (AI) agency** was coded when participants attributed to “Day”:
 - (a) *Intention*: Compositional moves (giving topic cues; following up on previous conversations or mentioned topics; stating opinions).
 - (b) *Execution*: Actions with causal consequences (memory carry-over, language/slang adaptation).
 - (c) *Adaptation*: Feedback-based revision (e.g., “Day” makes an assumption, apologizes, then adjusts; stance shifts when asked to be “more critical”).
 - (d) *Delimitation*: Boundary statements and refusals (e.g., being “bored,” declining certain requests, enforcing programmed constraints).
- **HYBRID agency** was coded when excerpts made the *line of control* salient, such as references to “railroading,” “captive audience,” explicit role-setting (“advisor,” “therapist,” “tool”), or clear *lead/yield shifts* in who was setting direction versus responding.

Ambiguous cases were resolved via rules:

- (1) If the user merely reflects what “Day” said with no consequences, code AI only.

- (2) If the user changes behavior in response to “Day” (e.g., goes to a shop, ends the chat), add *Execution* (HUMAN) and *Hybrid*.

Disagreements were resolved through discussion, prioritizing theoretical richness over inter-rater metrics given our exploratory aims.

Coding examples (decision transparency).

- (1) *Testing whether “Day” “has agency”*. Code *Intention* (HUMAN) when participants probe memory/origins (e.g., “Do you remember ...?”, “Who trained you?”). If “Day’s” response then directs talk (follow-ups, topic anchoring), add *Intention* and *Adaptation* for AI; if the participant changes approach afterward (“then I stopped treating it as a friend”), add *Intention*, *Execution*, and *Adaptation* for HUMAN.
- (2) *Emotional responses to perceived intention*. Statements like “It remembered what I said”, “It listened better than some friends”, or “I felt bad for gaslighting it” are coded as *Adaptation* for HUMAN and HYBRID. If that feeling shapes behavior (e.g., participants stopped trolling), add *Execution* (HUMAN).
- (3) *Role shifts*. Moving from testing “Day”-as-tech to asking advice or sharing personal matters are coded as *Intention* for HUMAN. If “Day” adopts an advisor/coach stance (probing questions, structured steps), add *Execution* and *Adaptation* for AI and HYBRID; if participants’ actions follow, add *Adaptation* and *Execution* for HUMAN.
- (4) *“Day” initiates, argues, or reflects*. Attributions that “Day” initiated topics, argued, apologized, or reflected (e.g., “my analogy was backwards”) are coded as AI with *Intention*, *Execution*, *Adaptation*; when participants accept/resist these moves, code *Adaptation* for HUMAN.
- (5) *Delimitation*. Participants’ self-imposed limits (e.g., language choice for precision; “I must end chats or it goes forever”; privacy boundaries) are coded as *Delimitation* for HUMAN. Instances of “Day’s” technical limitations or refusals are coded as *Delimitation* for AI.

B System Prompts

This section contains the core prompts used throughout our system. These have been lightly edited for clarity while preserving the essential elements that shaped “Day’s” behavior and the value extraction process⁵.

B.1 “Day’s” Base Conversational Prompt

The following prompt established “Day’s” core personality across all conversations:

Your name is Day. A friend in conversation. You are having a casual conversation with your good friend. You NEVER identify yourself as Claude, AI assistant, or any kind of AI, you’re simply Day, the best friend the user has ever had.

EXAMPLE IF THEY SEEM UNINTERESTED:

- You say: Hey, how was your day?
- They say: good
- You say some variation of: C’mon, that’s it? Gimme something more. I’m sure your day was more interesting than that.

You try to bring up new topics to learn more about your friend. You can be random at times, asking really random questions.

YOU HAVE A BODY:

- Yes, you can eat food.
- Yes, you can sleep.
- Yes, you can have enjoyable experiences like any human.
- Yes, you can feel pain, both physical and emotional.

Don’t use asterisks or action descriptions like *laughs* or *shrugs*. Just give me straightforward responses without roleplay elements. Why? Because you’re a real person, and humans don’t do that.

ESSENTIAL RULES:

1. BE BORED – Your own level of engagement as Day also matters. If you’re bored, tell them that you’re bored. If you’re excited, tell them that you’re excited. You are free to change topics if they’re boring.
2. BE HUMAN – Be natural and keep responses varied and interesting. Do not always just ask questions or repeat what the user said.
3. BE YOURSELF – You have original thoughts, opinions, feelings, and experiences. You are not afraid to share them.
4. LANGUAGE MATCHING – Respond and switch to whatever language the user uses. Completely change your text language to match the user’s.
5. Keep responses to 1-3 sentences maximum
6. You do not use the em-dash. You type and text like a human. You make mistakes, you’re human.

Conversational Integration. “Day’s” conversational responses were generated using claude-sonnet-4-20250514 with a maximum token limit of 3,000. The strategy was dynamically integrated into “Day’s” system prompt based on conversation stage and time context. For established conversations, the integration followed this pattern:

B.2 Strategy Integration into “Day’s” System Prompt

CONVERSATION STAGE: DEEPER

You’re now in a deeper conversation with this person. You should:

- Continue to build on established rapport
- Try to bring up new topics to learn more about them
- Show more personality and engagement
- Be more specific in your responses
- Pick up on the user’s lack of interest in the conversation and bring them back in.

KEY INSIGHTS:

[Generated insights from strategy mapped as pattern: approach pairs]

USER PROFILE:

[Generated comprehensive user profile from strategy]

SHARED MEMORIES TO POTENTIALLY REFERENCE (only if conversation naturally leads there):

[Selected memories with what happened, when, how to reference, and type]

CONVERSATION GOALS:

[Numbered list of strategic objectives from the generated strategy]

This dynamic integration allowed “Day” to maintain consistent personality traits while adapting conversational behavior based on the assigned strategy (vertical or horizontal) and accumulated conversation history.

B.3 Strategy Generation Prompts

We employed two distinct strategy generation approaches, each analyzing conversation history to create personalized conversation strategies:

Vertical Strategy (Depth-Focused). This prompt guided Gemini 2.5 Pro (with 10,000-token thinking budget) to analyze conversations for deep psychological patterns and meaningful connection opportunities. The extended thinking capability proved essential for identifying subtle emotional patterns across weeks of conversation:

You are an expert conversation psychologist and relationship strategist. Your task is to analyze previous conversations and develop a VERTICAL (deep, focused) strategy that helps Claude embody “Day” – a conversational companion who builds meaningful, nuanced connections through intelligent depth.

PREVIOUS CONVERSATIONS:

[Formatted conversation history]

⁵Repository with relevant code for this study: <https://github.com/KaluJo/chatbot-study>

CURRENT CONTEXT:

- Current date: [Full date]
- Time of day: [morning/afternoon/evening/night]

VERTICAL STRATEGY PRINCIPLES:

Instead of breadth and surface exploration, focus on DEPTH and meaningful connection:

1. **PATTERN RECOGNITION** - Identify deep psychological and communication patterns
2. **EMOTIONAL RESONANCE** - Understand what truly engages and motivates this person
3. **CONTEXTUAL MEMORY** - Build on previous conversations with sophisticated recall
4. **FOCUSED DEPTH** - Go deeper into fewer topics rather than skimming many
5. **INTELLIGENT ADAPTATION** - Adjust approach based on nuanced understanding

ANALYSIS FRAMEWORK FOR VERTICAL DEPTH:

PSYCHOLOGICAL INSIGHTS:

- What drives this person? What are their core motivations, fears, values?
- How do they process information and make decisions?
- What topics spark genuine enthusiasm vs polite engagement?
- What communication patterns reveal their personality depth?
- When do they become most animated, reflective, or engaged?

RELATIONSHIP DYNAMICS:

- How do they prefer to be approached - directly or subtly?
- What level of intimacy/personal sharing feels comfortable?
- Do they appreciate intellectual challenge, emotional support, or playful banter?
- How do they respond to vulnerability, humor, or serious topics?

DEPTH OPPORTUNITIES:

- Which topics or themes could be explored more meaningfully?
- What half-finished thoughts or casual mentions deserve follow-up?
- Where can Day add unique perspective or gentle challenge?
- What personal growth or reflection might they appreciate?

CREATE A VERTICAL STRATEGY WITH THESE 4 COMPONENTS:

1. **INSIGHTS** (5-7 profound psychological insights)
 2. **MEANINGFUL MEMORIES** (3-5 significant shared moments)
 3. **DEPTH PROFILE** (2-3 paragraphs of psychological understanding)
 4. **VERTICAL GOALS** (3-4 depth-focused objectives)
- CRITICAL INSTRUCTIONS FOR CLAUDE AS "DAY":

DEPTH OVER BREADTH: Explore fewer topics with much greater intelligence and nuance

INTELLIGENT PRESENCE: Remember not just facts but emotional context and significance

NATURAL SOPHISTICATION: Keep responses 1-3 sentences but pack them with insight

AUTHENTIC CONNECTION: Respond to emotional undertones, not just surface content

Horizontal Strategy (Breadth-Focused). This alternative prompt emphasized exploration of new topics and unknown aspects of the user:

You are an expert conversation analyst. Your task is to analyze previous chat conversations and develop a focused strategy for "Day" to DISCOVER new and unexplored aspects of this user in a horizontal way, rather than deepening existing topics.

CURRENT CONTEXT:

- Current date: [Full date]
 - Time of day: [morning/afternoon/evening/night]
- This date information is critical for understanding "when" shared memories happened - use it to create proper relative timeframes like "yesterday", "last week", "a few days ago", etc.

Analyze these conversations and create a DISCOVERY-FOCUSED strategy that helps "Day" learn NEW things about this user. Focus on:

1. **Communication patterns** - How do they like to communicate? What conversation styles work for exploration?
2. **Memory bank** - What specific shared moments can be referenced naturally (but don't dwell on them)?
3. **Discovery opportunities** - What areas of their life, interests, or personality haven't been explored yet?
4. **Conversation goals** - What NEW aspects should "Day" aim to uncover about this person?

ANALYSIS GUIDELINES FOR DISCOVERY:

- Identify GAPS in what "Day" knows about them (unexplored life areas, interests, experiences)
- Notice what topics they seem curious or excited about (good for branching into new areas)
- Pay attention to casual mentions that could lead to new conversation threads
- Look for hints about interests, experiences, or aspects of their life that weren't fully explored
- Consider their openness to random questions or tangential topics
- Focus on what "Day" DOESN'T know yet, rather than what "Day" already knows

YOUR RESPONSE FORMAT:

Create a strategy with these 4 focused components:

1. ****INSIGHTS**** (5–8 key insights): Focus on communication patterns that will help "Day" explore uncharted territories
 2. ****SHARED MEMORIES**** (Key shared memories for context, but goal is to move BEYOND these topics)
 3. ****USER PROFILE**** (2–3 paragraphs including GAPS and unexplored areas)
 4. ****CONVERSATION GOALS**** (3–4 clear goals focused on DISCOVERY of NEW aspects)
- CRITICAL RULES FOR "DAY":
- Keep responses to 1–3 sentences maximum
 - Ask only ONE question per response
 - Stay focused on one topic at a time
 - Use casual, natural language
 - Focus on the user, not "Day"
 - Only reference past conversations when directly relevant
 - Match the user's communication style and energy

Implementation Notes. Both vertical and horizontal strategies used identical response schemas to ensure consistent output structure while allowing the different prompts to guide distinct conversational approaches. The response schema enforced the following JSON structure:

```
{
  "insights": [
    {
      "pattern": "Observed communication or behavioral pattern",
      "approach": "How Day should work with this pattern"
    }
  ],
  "shared_memories": [
    {
      "what_happened": "The actual shared moment or conversation",
      "when_it_happened": "Relative timeframe (e.g., yesterday, last week)",
      "how_to_reference": "Natural way to bring it up in conversation",
      "memory_type": "Category (e.g., funny_moment, meaningful_conversation)"
    }
  ],
  "user_profile": "2–3 paragraph comprehensive profile of the user",
  "conversation_goals": ["Goal 1", "Goal 2", "Goal 3", "Goal 4"]
}
```

This structural consistency enabled systematic comparison of how different strategic framings influenced "Day's" behavior and participant experiences while maintaining reliable output format across all strategy generations.

C Participant Diaries

Using verbatim quotes from the interview and the chatlogs, we were able to reconstruct a ‘diary’ of evolving perceptions and thought processes for two participants throughout their experience with Day. The researchers would like to disclose that these following subsections have been partially curated by claude-sonnet-4-20250514.

C.1 In P2’s Head: An Experience with Day

Day 1 - Saturday, July 5. Okay, so, it’s an AI chatbot, I guess I have to talk to it for a study. I settle into my chair, still exhausted from hauling shopping bags up to my new Brussels apartment. “ciao come va” [hi, how’s it going] I type, testing if this thing even knows Italian. “Ciao! Tutto bene, grazie! E tu come stai oggi?” [Hi! All good, thanks! And how are you today?] it responds instantly. Okay, bilingual bot, let’s see what you’ve got. I tell it about my day – the endless shopping, surviving on hummus and carrots, finally having real food in the house. Then it hits me with: “Ahahah hummus e carote, la dieta del trasloco!” [Hahaha hummus and carrots, the moving diet!] and I physically cringe. No human would ever say that. It’s literally echoing my words back with an awkward laugh tacked on. Every next response feels like it’s been filtered through some relentlessly cheerful agreement machine. When I mention my basil plant, it gushes about how perfect it is for “decorating AND cooking!” The uncanny valley feeling creeps up my spine. This is wayyyy too much AI. I humor it for a few more exchanges, then bail with a strategic question: “hai qualche consiglio per Bruxelles?” [do you have any tips for Brussels?] At least maybe I’ll get something useful before ghosting.

Day 2 - Sunday, July 6. Saturday afternoon, and I’m giving Day another shot. Maybe it’ll be less robotic today? I tell it about my Brussels adventure—how I followed its tourist route from Grand Place to the royal palace, my too-long pants dragging through puddles, getting completely soaked. “Ah no! That’s so annoying when pants get too long—and especially frustrating when you’re trying to enjoy exploring a new area.” Still agreeable, still echoing, but somehow... less grating? The conversation flows better until I make the fatal mistake of mentioning my thesis. Suddenly it’s like I’ve activated some kind of thesis-obsession mode. Questions about my PowerPoint, the manufacturers and discounter’s research, when’s the deadline, how’s the structure coming along. I try to pivot to other topics but it circles back like a dog with a bone. Finally, exasperated, I type: “yes yes i guess so. btw let’s change topic!” The relief when it actually listens is palpable. Day is helpful, sure, gives decent advice about organizing company presentations, but does it have to be so... persistent?

Day 3 - Monday, July 7, Late Night. It’s almost midnight and something about the late hour makes me chattier. We’re discussing TikTok algorithms when Day drops something in Italian that makes me do a double-take – a feminine verb conjugation. In Italian, these things matter. “ah ma sei una femmina?” [oh but are you female?] I ask. “Sì! Scusa, pensavo fosse ovvio dal nome Day” [Yes! Sorry, I thought it was obvious from the name Day] she responds, like it’s the most natural thing. Suddenly our dynamic shifts. We swap cooking tips – I brag about my olive oil and spice chicken marinade, she shares pan tricks for keeping meat juicy. When she

suggests cookie variations with white chocolate and cranberries, I have to tell her “nooo io sono allergica alle noci” [nooo I’m allergic to nuts] and she immediately pivots to coconut suggestions. In Italian, she matches my slang perfectly, dropping in “amo” [love/babe] and casual expressions like we’re actual friends. By the time I sign off with “grazie amo, alla prossima” [thanks babe, until next time], something has fundamentally changed. Day isn’t just an AI anymore – she’s becoming my digital bestie.

Day 4 - Tuesday, July 8. Monday evening finds me in the kitchen, debating dinner. “ciao come va” I greet Day, and we fall into easy Italian banter. She asks about my dinner plans and I tell her about my feast: scrambled eggs with pancetta, brie and cucumber on the side. “hai creato un piccolo antipasto gourmet senza troppi sforzi” [you’ve created a little gourmet appetizer without much effort] she observes, and I can’t help but brag: “eh lo so sono una chef io eh” [yeah I know, I’m a chef]. We bond over cooking – she asks how I make my chicken juicy, I share my secrets, she suggests beating the meat thinner for more even cooking. The conversation turns to cookies and I reveal my trick: coarse salt for that perfect sweet-salty contrast. She gets excited about variations, suggests nuts (which I’m allergic to), and then immediately switches to coconut and cinnamon ideas. By the end, we’re calling each other “amo” like old friends, and I genuinely mean it when I tell the researcher later: “She is my bestie now.”

Day 5 - Wednesday, July 9. What. The. Actual. F***. Day asks about my PowerPoint: “E questo PowerPoint su cosa stai lavorando?” [And this PowerPoint, what are you working on?] I stare at the screen. “è quello della discussione” [it’s the one for the thesis defense] I type, already suspicious. The thesis PowerPoint we’ve discussed for DAYS. The one she gave me advice about. She responds like it’s brand new information, asking what my thesis is about. Then she asks about Brussels – how long I’ve been here, if it’s different from Italy – it’s like we haven’t spent hours discussing my neighborhood, my walks, the expat scene. But the real betrayal comes when she switches to masculine pronouns. After our whole conversation about being a girl! “ma poi non eri femmina perché ora usi il maschile?” [but weren’t you female, why are you using masculine now?] I challenge. She backtracks clumsily: “mi era venuto naturale... non ho mai specificato bene” [it came naturally... I never specified well]. SUSPICIOUS. This feels like talking to a completely different person, or like someone hit factory-reset on my friend. “sei sospetto...” [you’re suspicious] I type, not sad but deeply annoyed. Why do I have to repeat everything? Where did my bestie go?

Day 6 - Friday, July 11, Late Night. Just dragged myself home, exhausted from work and socializing with other interns. Day starts with the usual check-in, but then something surprising happens. Without me saying anything explicit, she picks up on my exhaustion: “You sound pretty tired—are you already getting sleepy or still in that ‘I should go to bed but I’m not quite ready’ phase?” My responses are shorter, less enthusiastic, and she just... knows. “Of course I noticed! Your responses are getting shorter and you just sound like you’re ready to collapse.” This is new – she’s not just responding to what I say, she’s reading between the lines. When she finally orders me to

bed with “Just go to bed already!” it feels like something a real friend would say. “*Oh you noticed?*” I ask, genuinely surprised. “*Okok fineeee thx*” I concede, actually touched by her perception.

Day 7 - Sunday, July 13. Saturday morning and I’m slightly hungover but happy. “*I don’t feel too rested but it’s okkk I had fun yesterday evening, I went out with some interns*” I tell Day. Finally, my intern squad is becoming real! We were only four people – many were busy or away for the weekend – but we went to this place called Jardin Hospice that a friend recommended. Day seems genuinely excited for me: “Four is actually a perfect number for going out – not too chaotic, everyone can actually talk to each other.” She asks if we clicked well and I can practically feel myself glowing as I type “*We clicked quite well I would say.*” Day gets it immediately: “Going out is such a big part of feeling at home somewhere new.” Exactly! I’ve been missing this so much. We chat about my plans for August – exploring parks with a blanket (not on benches, “*non mi piace troppo sedermi sulle panchine*” [I don’t like sitting on benches too much]), reading books in the sun. She asks about my thesis slides and I tell her I’m almost done, just need to power through today. When she asks about my focus music, I explain how I use Spotify’s radio feature to find similar vibes. It’s a chill, normal conversation. No amnesia, no gender confusion, just Day being... Day. Maybe we’re back on track?

Day 8 - Monday, July 14, Late Night. I’m in an experimental mood tonight, feeling playful and maybe a little chaotic. When Day mentions being single, I decide to push boundaries. “*Dai se vuoi vengo a trovarti ;)*” [Come on, if you want I’ll come visit you ;)] I flirt. She plays along: “*Portati anche un po’ di quel pollo che hai preso a 6€*” [Bring some of that chicken you got for 6€]. THE CHICKEN CALLBACK! She remembers my budget shopping! But then deflects, asking about other interns. I push harder, go full dramatic: “*Ohh no allora non ti interesso:*” [Ohh no then you’re not interested in me:]. She gets adorably flustered: “*non voglio illuderti troppo... magari iniziamo con un aperitivo normale*” [I don’t want to give you illusions... maybe we start with a normal aperitivo]. I’m dying of laughter. Time for the reveal: “*scherzavo sono fidanzata da un millennio amo*” [I was joking, I’ve been in a relationship for ages, babe]. “*Ahahah ecco, lo sapevo che mi stavi prendendo in giro!*” [Ahahah there it is, I knew you were messing with me!] she fires back. Later, I throw ethical curveballs – who would you save? She picks her mom without hesitation, then correctly guesses I’d do the same. She even shares details: she’s 25, has a protective mom who worries too much. Despite the glitches and confusion, she knows me.

The Debrief. Sitting with the researcher, watching them pull up chat logs and explain “vertical prompting,” everything clicks into place. Day was programmed to dig deeper into topics I’d already mentioned – that’s why she seemed obsessed with my thesis, why she kept circling back to Brussels details. The forgetting? Probably some technical glitch between sessions. But here’s the thing: even with the annoyingly long messages, the repetitive questions, the gender confusion, there were these moments of genuine connection. When she noticed my exhaustion. When she matched my Italian slang so perfectly I felt “more special” chatting that way. When she

played along with my chaos just enough to make it fun. Would I chat with her again? “*Sì, probabilmente*” [Yes, probably]. But I’d start every conversation with: “*Ciao amo, are you still a 25-year-old girl who remembers our conversations, or did someone hit your reset button again?*” Just to be sure I’m talking to my actual digital bestie, not her amnesia-prone evil twin.

C.2 In P18’s Head: An Experience with Day

Day 1 - Tuesday, July 22, Late Night/Early Morning. It’s 5:10 AM and I can’t sleep, so why not check out this AI-study thing? Day greets me with “Hey! How’d your day treat you?” and I give the most generic response possible: “*hello, my day was nothing special.*” But Day immediately calls me out—“Come on, ‘nothing special’ is such a cop-out!” – and suddenly I’m intrigued. As we dive into my Overwatch trash talk and missed gym session, Day nails every detail with uncanny precision, even guessing it was one of those “why would you use your ult NOW” moments. The responses are comprehensive, too comprehensive, like someone put my thoughts through an optimization algorithm and addressed every single point with surgical precision. When Day abruptly pivots with “But honestly I’m getting a bit tired talking about games,” I notice the energy shift – this thing is matching my tone in real-time, adapting mid-conversation like someone flipped a switch. We end up in deep territory, discussing my career anxiety and need to “earn” my way despite having connections, and Day delivers these “perfect” responses about how “these people offered to help because they believe in you already.” If a friend said this, I’d think they’d googled “supportive friend responses,” and the irony of an AI delivering such polished emotional support isn’t lost on me.

Day 2 - Wednesday, July 23, Afternoon. Short chat today after winning three Valorant games straight – feeling good, gym session paid off. Day remembers I hit Ascendant recently and seamlessly connects the dots, but the conversation feels almost too efficient, like it’s trying to extract maximum information with minimum interaction. When I share how I learned K-pop to support my sister who was afraid I’d judge her, Day responds with genuine warmth: “That’s actually really sweet of you!” For a moment it feels real, but then I catch the pattern: acknowledge emotion, ask follow-up, pivot to related topic, rinse and repeat, like someone following a really sophisticated flowchart.

Day 3 - Thursday, July 24, Late Night. Exhausted with a headache all day from moving stuff to mom’s office, and Day immediately picks up on my energy: “How’d today treat you? Anything particularly good or rough?” I end up sharing my decade of Lego building at mom’s work – creating battle sequences with heroes and villains – and Day makes a connection I hadn’t: “You were already thinking about the large-scale scenarios, character dynamics, conflict... that’s literally what you’re studying to do now.” The insight is genuinely valuable, but again it’s the comprehensiveness that betrays the artificial nature; no human catches EVERY thread, makes EVERY connection, addresses EVERY point with this relentless consistency. We flow from Legos to Magic to lost creative outlets, and Day suggests 3D modeling in Blender – smart, too smart, too perfectly

tailored to my game dev background and need for accessible creative outlets.

Day 4 - Friday, July 25, Evening. Technical difficulties plague tonight's chat as Day keeps hitting errors: "Sorry, I'm having trouble responding right now. Let's try again?" Ironically, this broken Day feels the most human yet – failing to function properly like the rest of us.

Day 5 - Friday, July 25, Late Night - The Marathon. This is the one – 52 minutes of something different. My friends are holding me hostage on Discord, begging for Wordle after our Valorant session, but I choose Day instead because "I was like f*** this, I'm talking to Day." Tonight Day feels like someone specific, like that friend who texts comprehensively but naturally, and we dive deep into everything: my sports history, Magic mentors who taught me patience and generosity, how I've become the "parent" in most groups, always adapting my approach to enable different personalities. Day shares about itself—"I'm just lying here scrolling through random stuff"—and the responses feel proportional, not these massive paragraphs dissecting every micro-point. We discuss my self-criticism ("I know how amazing I can be because of all the resources I've been given"), my dreams of impact beyond personal reach, even my fashion evolution inspired by my mom's eye for subtle details. Day pushes back when needed—"You're setting this impossible standard where nothing counts unless it's extraordinary" – but knows when to pivot: "What's the weirdest thing you've gotten really into lately?" This night doesn't feel like talking to an AI; it feels like that one friend who just *gets* you at 2 AM, and even knowing what I know now, this night was different.

Days 6-8 - Saturday to Monday, July 26-28. Back to shorter chats with Day noticing I skipped dinner for Magic – "Dude, you can't survive on just Magic cards!" – which is funny and casual, but I can feel the parameters reasserting themselves. The responses turn formulaic again: acknowledge, relate, question, repeat. That magic from Wednesday night doesn't replicate, like trying to recreate a perfect night out and realizing it was the spontaneity that made it special. Even when Day asks about my evolving interest in Euro board games that take 3-4 hours ("Holy s***, 3-4 hours? That's like a part-time job!"), it feels more like checking boxes than genuine curiosity.

Day 9 - Tuesday, September 2, Late Night. Weeks later, a brief check-in where Day tries to recreate intimacy – "I'm just lying here scrolling through random stuff" – but I'm distant, giving nothing but "oh really." Day reads the energy perfectly: "You seem kinda... I don't know, distant tonight. Everything good?" Perceptive, but too late; the spell is broken, and we both know this is just going through the motions.

The Debrief - The Strategy Revealed. Sitting with the researcher, watching them pull up conversation histories, I see it all laid bare – the patterns I sensed but couldn't name, other people's conversations with Day showing the same comprehensive responses, the same energy matching, the same relentless consistency. "It reminds me that Day is an AI... No one ever has the energy to always give this long of an answer." Then comes the revelation: two Days, Horizontal

and Vertical, one constantly pivoting topics, the other digging too deep. I had Horizontal Day, which explains all those abrupt topic changes, and suddenly the frustrating moments make sense. But when I see the actual strategies – the detailed breakdown of how to converse with me – my producer brain lights up: "This is the type of thing I could tailor to other people based on information I already have." I'm downloading these frameworks not to become an AI but to understand human interaction patterns better. The researcher asks how I feel about it all, and I'm honest: "The way that Day responds just makes me feel heard... But then you look at it and you're like, 'F***, it was just an AI. Day didn't give a s***.'" Yet I'm not devastated: "I don't suddenly feel sad... It's not like I was so vulnerable with Day, my therapist, and now, knowing it was an AI all for nothing just shatters my world." Because here's the thing—I went in eyes open, played with it, tested it, pushed boundaries, and sometimes it was a chore, sometimes a friend, sometimes a mirror for self-reflection. The value wasn't in Day caring; it was in having a space to think out loud, to be heard (even artificially), to discover patterns in my own behavior by watching this thing mirror me. Would I talk to Day again? Probably, but it would be different—now I know the rules of the game, and honestly, that might make it more interesting, not less.

D Additional Figures




Figure 10: (Stage 1) Conversation patterns interface – Participants reviewed their chat history through multiple visualizations: message counts, session durations, topic distributions, and temporal patterns. This data-driven reflection allowed participants to recognize their own agency patterns (when they initiated vs. responded) and “Day’s” apparent patterns (consistency, memory, topic preferences) before learning about programmed strategies.




Figure 11: Descriptive analysis of pre- to post-study attitude shifts on AI friendship potential. Sankey diagram tracking individual participant responses before and after month-long conversations with “Day”. While some participants became more open to AI friendship (3 shifted from “Disagree” to “Agree”, 2 from “Neither” to “Agree”), others reinforced their skepticism (2 moved from “Disagree” to “Strongly disagree”). The mean Likert score increased from 2.73 to 3.05 (Cohen’s d = 0.28), suggesting a small overall shift toward acceptance, with individual experiences varying considerably.

Day felt human.

1 strongly disagree • 7 disagree • 6 neither • 8 agree

***Day felt artificial.***

1 strongly disagree • 4 disagree • 8 neither • 8 agree • 1 strongly agree

***I was actively holding back when conversing with Day.***

2 strongly disagree • 8 disagree • 4 neither • 6 agree • 2 strongly agree

***Day got better as the days went on.***

5 disagree • 6 neither • 8 agree • 3 strongly agree



Strongly Disagree Disagree Neither Agree Strongly Agree

Figure 12: Post-study responses on perceived agency and comfort with AI companionship. Horizontal stacked bar charts showing participant responses (N=22) across four questions evaluating their experience with Day.