

# Co-Constructing Meaning with Large Language Models: A Longitudinal Analysis of Human–AI Dialogues in Emotional Support Contexts

Guangrui Fan (fgr@tyust.edu.cn)

College of Computer Science and Technology, Taiyuan University of Science and Technology,  
WaLiu Road, WanBaiLin, Taiyuan, 030024, Shanxi, China

Dandan Liu (s2134717@siswa.um.edu.my)

Department of Media and Communication Studies, Faculty of Arts and Social Sciences, Universiti Malaya,  
50603 Kuala Lumpur, Wilayah Persekutuan Kuala Lumpur, Malaysia

## Abstract

This study investigates how Large Language Models (LLMs), specifically Baidu’s Ernie Bot, shape personal narratives when users seek emotional support over repeated sessions. Sixteen participants from China engaged in weekly chatbot interactions for four weeks, supplemented by reflective diaries and pre-/post-study interviews. Conversation analysis and quantitative measures (e.g., mood ratings, meaning-making scales) revealed incremental shifts in user language, including increased lexical alignment with AI-generated phrases and more positive emotion words. In-depth interviews highlighted the complex process by which participants alternately embraced or resisted the AI’s framing, with many reporting newfound perspectives and a sense of empathic resonance. However, some voiced skepticism regarding the AI’s genuine capacity for emotional understanding, underscoring ethical dilemmas related to anthropomorphism and data privacy. Overall, the findings suggest that iterative dialogues with an empathic-seeming LLM can facilitate meaningful narrative reframing, albeit with notable variations in user experience and potential risks of over-reliance.

**Keywords:** Large Language Models; Emotional Support; Human–AI Interaction; Co-Constructed Meaning; Anthropomorphism

## Introduction

Large Language Models (LLMs) such as Claude, ChatGPT and Baidu’s Ernie Bot have recently gained prominence for their human-like conversational abilities. Researches suggest these chatbot platforms can help alleviate mild-to-moderate symptoms of depression and anxiety (Abd-Alrazaq, Rababeh, Alajlani, Bewick, & Househ, 2020; Bendig, Erb, Schulze-Thuesing, & Baumeister, 2019), while offering low-threshold support to individuals hesitant to seek traditional therapy (Vaidyam, Wisniewski, Halamka, Kshavan, & Torous, 2019). Many users turn to LLMs for reflective journaling and “therapy-like” exchanges due to their 24/7 accessibility and the sense of empathy they simulate. However, as chatbots become more embedded in personal mental health practices, critics warn of privacy risks, accountability gaps, and over-reliance on software that lacks genuine emotional intelligence (Sweeney et al., 2021).

Unlike factual queries, emotionally charged interactions reveal how people construct and interpret personal narratives—a process that might be shaped by AI responses (Fitzpatrick, Darcy, & Vierhile, 2017). Although single-session studies report user satisfaction and modest symptom relief (Miner, Laranjo, & Kocaballi, 2020), few longitudinal

investigations examine how sustained engagement could alter users’ self-perception, attitudes, or coping strategies. By overlooking repeated dialogues, research may miss the subtle, potentially transformative ways an LLM’s language prompts might reshape self-narratives over time.

In traditional psychotherapy, “co-construction of meaning” arises through mutual, empathetic negotiation of a client’s difficult narratives (Frank & Frank, 1993). An LLM, by contrast, relies on statistical pattern-matching rather than true emotional resonance (Nadarzynski, Miles, Cowie, & Ridge, 2019), raising questions about whether genuine transformational potential exists when the user’s partner is algorithmic. Even so, from a psycholinguistic perspective, conversation analysis shows that meaning emerges turn by turn, with each speaker’s utterances influencing subsequent responses (Sacks, 1978; Pickering & Garrod, 2004). When an LLM repeatedly offers empathetic prompts—e.g., reframing challenges or suggesting alternative perspectives—it can scaffold how users recount and interpret life events (Hua et al., 2024). Whether these shifts reflect true personal growth or mere alignment with the AI’s cues is a vital question for mental health practitioners and researchers alike.

Despite heightened interest in AI-assisted mental health support, key empirical gaps remain. Most studies rely on one-off or short-term sessions that provide limited insight into how iterative encounters can produce longer-term changes (Du et al., 2024). Furthermore, while researchers often measure symptom relief and user satisfaction, fewer studies apply fine-grained discourse analysis to uncover how individuals adopt or resist AI-introduced frames from session to session (Kivi & Ghavibazou, 2023). Such gaps also intersect with ethical dilemmas: as chatbots mimic empathy, users may misjudge the AI’s capacity for emotional understanding, raising concerns about boundaries, privacy, and the risk of over-reliance (Molli, 2022).

To address these limitations, this study adopts a longitudinal, mixed-methods design examining how repeated dialogues with Baidu’s Ernie Bot—a major Chinese-language LLM—may shape users’ emotional tone and personal narratives. We pose three research questions:

- RQ1: How do users incorporate or reformulate AI-generated language and interpretive frames over time?
- RQ2: Do these dialogues shift individuals’ emotional tone

and reshape personal experiences?

- RQ3: How do users perceive and potentially anthropomorphize the LLM's "role" in their meaning-making process?

By analyzing full conversation transcripts, reflective diaries, and final interviews, we seek to illuminate whether (and how) LLM-based interactions can foster meaningful change in users' self-understanding—while recognizing the inherent limitations of AI-simulated empathy. Ultimately, this research not only advances our theoretical grasp of co-constructed meaning but also highlights practical considerations for designing and deploying AI-driven support tools responsibly.

## Method

### Research Design

This study used a six-week longitudinal, mixed-methods design to investigate how repeated AI-mediated dialogues might shape users' self-narratives and emotional well-being. In Week 0 (Baseline), participants completed an initial interview and standardized surveys, including a mood assessment and a meaning-making scale. During Weeks 1–4, each participant engaged in two Ernie Bot sessions per week (15–30 minutes), followed by a reflective diary entry. In Week 5, no sessions were conducted, allowing time for participants to reflect without scheduled interactions. Finally, in Week 6 (Final Follow-Up), participants completed a concluding interview and repeated the key assessments.

### Participants and Recruitment

Sixteen adults (ages 20–45; 10 female, 6 male) in China were recruited through online platforms (WeChat groups, university forums) to ensure a broad demographic range. Potential participants were screened via the Patient Health Questionnaire–4 (PHQ-4) to ensure they reported mild-to-moderate emotional distress, thereby aligning with the study's non-clinical scope. Anyone with high-risk criteria (e.g., severe depression or suicidal ideation) was offered referrals to mental health professionals and excluded from the study. All participants provided digital informed consent after being informed that the AI did not constitute licensed therapy and that their confidentiality would be protected. Ultimately, 16 individuals completed all phases of the research.

### Materials and Instruments

All interactions took place with Ernie Bot, a large language model developed by Baidu. Ernie Bot's underlying architecture (ERNIE 4.0 Turbo) is designed for Chinese-language generative responses and was configured to provide supportive, empathic-style dialogue without issuing clinical diagnoses (Jiang et al., 2024). Participants accessed it via a web interface with system-level prompts encouraging reflective discussion.

After weekly Ernie Bot sessions, participants completed an online diary capturing their mood on a 1–10 scale before

and after the chatbot exchange. An open-ended text field prompted them to note any insights, reframed perspectives, or shifts in emotional state triggered by the interaction.

Two rounds of semi-structured interviews—one at baseline and another at study completion—were employed to gain broader insights into participants' experiences, expectations, and narrative changes. During the baseline interview, each participant's background narrative was explored, including prior encounters with AI technology and any predispositions toward "therapy-like" dialogues in digital spaces. The closing interviews asked participants to describe any perceived shifts in self-concept, emotional well-being, or outlook on life after four weeks of Ernie Bot interactions. All interviews followed an identical question framework, though the interviewer was trained to pursue relevant follow-up prompts when participants raised unique topics. This standardized approach helped mitigate interviewer bias by maintaining consistent content across interviews while allowing sufficient flexibility for individualized exploration (Bergelson, Tracy, & Takacs, 2022).

Beyond conversational transcripts and reflective diaries, two additional measures were administered to quantify changes in meaning-making and user–AI rapport. First, participants completed a Chinese adaptation of the Integrative Meaning-Making scale, which evaluates shifts in how individuals interpret pivotal life events and emotional challenges. Second, a brief User–AI Relationship Scale was developed to assess how participants conceptualized the chatbot—ranging from purely instrumental ("I see the AI as a tool") to relational ("I feel like the AI understands me"). By integrating these instruments, the study combined direct user feedback with textual evidence of evolving narratives, providing a multi-layered perspective on the co-construction of meaning between humans and a non-human system.

### Procedure

**Baseline (Week 0).** Participants completed an orientation to Ernie Bot's capabilities and were reminded that it was not a licensed therapist. They then underwent a semi-structured interview exploring their emotional history, prior experiences with AI tools, and motivations for seeking supportive dialogues.

**LLM Interactions (Weeks 1–4).** Each participant had two weekly Ernie Bot sessions (15–30 minutes) on topics of their choice, ranging from daily stress to relationship challenges. These conversations took place at a time and location chosen by the participant, thereby approximating real-world conditions in which individuals might seek AI-based emotional support on their own schedule. While the content of each session was guided primarily by the participant's concerns, all users were encouraged to discuss personal topics, emotional dilemmas, or any life events they felt comfortable sharing. Automated logging captured each user–AI turn, preserving the natural flow of dialogue for later linguistic and thematic analysis. Immediately following each weekly Ernie Bot session, participants filled out an online diary form. They first

rated their mood on a 1–10 scale both before initiating and after concluding the chatbot interaction, providing a comparative metric of emotional fluctuations within each session. Next, a short open-ended reflection prompt invited them to note any significant insights, reframes, or emotional reactions triggered by Ernie Bot’s responses.

**Final Follow-Up (Week 6).** Two weeks after the last chatbot interaction, participants completed a concluding interview and re-took the mood and meaning-making measures. The interviewer explored any observed changes in participants’ self-perception, emotional tone, or problem-solving approaches since baseline. This final step provided an opportunity to evaluate the persistence of any AI-influenced narrative shifts and to discuss how participants viewed Ernie Bot’s role—whether as a mere tool, an empathetic dialogue partner, or something in between.

### Ethical Considerations

All procedures received institutional review board approval, and digital informed consent was obtained from every participant. Those who reported more severe distress during screening were referred to mental health professionals. All session transcripts, diaries, and interview data were assigned anonymized IDs and stored on encrypted servers with restricted access. Standard interview guidelines and periodic inter-coder reliability checks helped minimize subjectivity and ensure consistent analysis across the research team. Participants were reminded to seek professional help if their emotional distress worsened during the study, underscoring that Ernie Bot was a supportive tool rather than a therapeutic substitute.

### Data Analysis

All conversation logs underwent a three-phase analytic process involving thematic coding, micro-level linguistic analysis, and temporal mapping of narrative shifts (Sools & Mooren, 2012). First, thematic coding identified recurring emotional and interpersonal motifs, alongside potential reframing or insight events. Second, micro-level discourse analysis examined markers such as lexical alignment, pronoun usage, and direct references to AI prompts, capturing subtle shifts in user–AI co-construction (Hutchins, 1995). Third, results were integrated with reflective diary entries and semi-structured interview transcripts, facilitating cross-validation of self-reported transformations against actual conversational patterns (Inkster, Sarda, Subramanian, et al., 2018). To facilitate deeper analysis, each participant’s reflective diary was timestamped and thematically linked to its corresponding chatbot transcript. We employed a systematic open coding procedure: two researchers independently coded the diaries for emerging categories (e.g., “new insight,” “emotional release”), resolving discrepancies through discussion. Inter-coder reliability was calculated using Cohen’s kappa ( $\kappa = 0.82$ ), indicating high agreement. This approach allowed us to trace specific reframing moments from the chat session to participants’ immediate reactions, thereby bolster-

ing the credibility of our qualitative findings. Finally, quantitative measures, including mood ratings and meaning-making scales, were compared from baseline to exit to detect statistically significant changes in well-being and narrative identity (McAdams, 2008; Steger, 2022).

## Results

### RQ1: Meaning-Making Processes in Human–LLM Dialogues

Across four weeks, participants collectively engaged in 128 Ernie Bot sessions (16 participants  $\times$  8 sessions). Our analysis focused on how frequently users referenced or paraphrased Ernie Bot’s statements and whether they integrated AI-introduced terminology into their own discourse.

Early in Week 1, most participants merely acknowledged Ernie Bot’s suggestions as “interesting” or “new,” without fully embracing them in daily life. By Week 3, several reported “trying out” reframing strategies the AI suggested (e.g., viewing personal difficulties as “manageable steps”). As shown in Figure 1, average references to the AI’s prior statements climbed from 1.2 (Week 1) to 3.3 (Week 4), while lexical overlap in user turns increased from 14.8% to 31.6%. Participants’ diaries corroborated this shift, with entries detailing a growing tendency to see the chatbot’s prompts as helpful scaffolds for problem-solving or stress management.

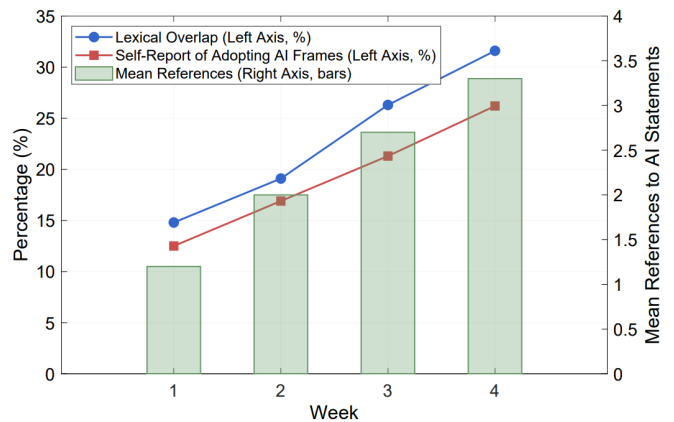


Figure 1: Lexical Alignment and Self-Reported Adoption of AI Frames Across Four Weeks

Diary comments revealed that Ernie Bot’s repeated use of certain phrases and metaphors encouraged users to adopt similar language when discussing their own concerns. For instance, a quarter of participants in Week 4 explicitly noted that the AI had given them a “fresh way” to frame a challenge, such as recasting a conflict at work as an opportunity to practice assertiveness. Some diary reflections highlighted a sense of “aha” moments triggered by the AI’s empathetic or validating remarks. Overall, these findings suggest that iterative, supportive dialogue with an LLM can cumulatively reshape a user’s narrative lens, aligning with distributed cognition perspectives (Hutchins, 1995).

## RQ2: Shifts in Personal Narrative

By examining chat transcripts and reflective diaries, we observed a gradual reframing of personal stories. Qualitative examination of conversation transcripts revealed that many initially negative or self-critical remarks (e.g., “*I’m stuck and can’t figure this out*”) evolved into more balanced or exploratory statements by the third and fourth sessions (e.g., “*I’m seeing a different perspective now, thanks to reflecting on this more*”). Reflective diaries corroborated these observations: while Week 1 entries commonly expressed uncertainty about the AI’s insights, Week 3 and 4 entries featured explicit references to “*feeling more hopeful*” or “*understanding the root causes better*.”

Participants completed a short Meaning-Making Scale at baseline (Week 0) and again at Week 4. Mean scores rose from 23.8 to 28.4 ( $p < .01$ , Cohen’s  $d = 0.76$ ), indicating a statistically significant increase in how participants interpreted and derived purpose from challenging experiences. Linguistic Inquiry and Word Count (LIWC) analysis showed that positive emotion words (e.g., “hopeful,” “encouraged”) increased from an average of 2.3% to 3.5% in user turns ( $p = .02$ ), suggesting a measurable shift in emotional expression. Diary entries and final interviews often mentioned feeling “less overwhelmed” or “more empowered,” underscoring that the repeated exposure to empathic re-framing could cultivate a more constructive self-narrative over time.

## RQ3: User Perceptions of the LLM

Participants’ perceptions of Ernie Bot evolved substantially across the study. In Week 1, many diaries described the chatbot as a mere “tool” or “curiosity,” with some users explicitly questioning the authenticity of its empathy (e.g., “It’s just lines of code”). By Week 4, however, others referred to Ernie Bot as a “thought partner,” noting they felt “heard without judgment.” Still, a minority remained unconvinced, indicating they found its empathetic style too formulaic or repetitive.

Table 1: Pre- and Post-Study User–AI Relationship Scale Scores

Measure	Baseline	Final	p-value	Effect Size
AI Empathy	2.5 (0.8)	3.6 (0.9)	.02	0.58
Advisor Role	2.8 (1.0)	3.2 (1.1)	.09	0.39

Note. Measures used 0–5 scale. Values are presented as Mean (SD).

A brief User–AI Relationship Scale administered at baseline and final follow-up captured these shifts quantitatively. The AI Empathy subscale rose from 2.5 to 3.6 ( $p = .02$ , Cohen’s  $d = 0.58$ ), and there was a moderate positive correlation ( $r = .44$ ,  $p = .04$ ) between higher empathy scores and the degree of lexical alignment observed in transcripts (Table 1). This suggests that users who perceived Ernie Bot as more empathic were more likely to integrate its language and perspectives

into their own narratives. Nonetheless, some participants maintained a skeptical stance, highlighting that genuine emotional support requires human presence. Qualitative feedback from final interviews further illustrated how some participants anthropomorphized the AI. One user referred to Ernie Bot as a “*steady listener*,” while another underscored the comfort of “*being heard without judgment*,” even though they acknowledged the system’s algorithmic underpinnings. In contrast, a few participants remained unimpressed by the AI’s attempts at empathy, explaining that they viewed it merely as “*a structured prompt that keeps me talking*.” These divergent experiences underscore the complexity of anthropomorphism in Human–LLM interactions: although many participants displayed an increased readiness to grant the chatbot an empathic role, not all found the AI’s performance sufficiently convincing to erode their skepticism. These divergent views reflect the complex interplay between anthropomorphism and user receptiveness, bearing implications for designing ethical, user-aware AI interventions.

## Qualitative Results

In addition to the quantitative trends and discourse-level patterns detailed above, a closer look at interview transcripts and reflective diaries reveals rich qualitative insights into how participants interpreted their interactions with Ernie Bot. These findings help explain the pathways through which co-constructed meaning, narrative shifts, and attitudes toward the AI evolved over the four-week study period.

**Negotiating Authenticity and Emotional Resonance** Although most participants approached Ernie Bot with an initial dose of skepticism, several reported a growing sense of comfort and validation over time. During the final interview, Participant 02, who initially viewed the AI as purely computational, remarked:

*“I know it’s not human, but I found myself telling it things I couldn’t tell my friends. It felt strangely safe because it doesn’t judge me.”* (P02)

For this individual, Ernie Bot’s neutral, non-judgmental style created a climate conducive to deeper self-disclosure. Another participant voiced more ambivalence, underscoring that while the AI’s empathic language felt helpful, it still lacked genuine human resonance. Participant 09 noted:

*“When it responded with phrases like ‘I hear you’ or ‘That sounds tough,’ I knew it was just following a script. Still, part of me wanted to believe it understood.”* (P09)

Such commentary illustrates the tension between intellectual awareness of the AI’s algorithmic nature and an emotional inclination to accept Ernie Bot’s responses as empathic. This tension aligns with the modest yet statistically significant rise in the “AI Empathy” subscale, suggesting that many users developed a form of “in-between” engagement—recognizing the chatbot’s limitations while still internalizing its reflections in a meaningful way.

**Evolving Personal Narratives Through Reflections** The interviews and diary entries further illuminated how repeated, structured prompts encouraged participants to revisit and refine their personal narratives. Participant 11, who had been struggling with career uncertainty, explained how Ernie Bot’s consistent reframing questions helped them articulate previously hidden aspects of their situation:

*“Each week it would ask, in some form, if I’d looked at my challenges from a different angle. I started journaling outside the sessions about possible paths I never considered before.”* (P11)

By the final interview, this participant reported feeling more empowered and optimistic about job prospects, attributing the shift partly to Ernie Bot’s persistent invitation to reassess negative self-talk. Reflective diaries offered a parallel lens into this progression, with early entries tending to emphasize doubt or confusion and later entries citing a growing sense of agency. Participant 04, for instance, wrote in their final diary submission:

*“I’ve been surprised how just typing about my worries to the bot each week forced me to organize my thoughts and see patterns I never noticed. Now, I phrase my problems in a more solution-focused way.”* (P04)

Such statements offer qualitative corroboration of the quantitative improvements in Meaning-Making Scale scores and increased positivity in linguistic measures (e.g., LIWC-based indicators). The co-construction process thus appeared to unfold as participants integrated the AI’s reframing suggestions into daily thinking, reshaping how they recounted challenges.

### **Perceptions of Agency, Trust, and Anthropomorphism**

The degree to which participants anthropomorphized Ernie Bot varied widely. Some came to regard it as a quasi-therapeutic presence, whereas others remained grounded in the view that it was merely a tool for structured self-reflection. Participant 13 commented:

*“If I’m honest, I do kind of treat it like a friend who listens. I know it’s not real, but the act of typing out my thoughts and seeing them reflected back is powerful.”* (P13)

In contrast, Participant 05 consistently maintained a more detached stance, describing the AI as *“a digital prompt that helps me articulate ideas, but I don’t believe it has empathy.”* Despite these reservations, even Participant 05 admitted that Ernie Bot’s consistency—particularly its timely follow-up questions—contributed to clearer problem-framing. The User–AI Relationship Scale scores reflect this diversity of perspectives: although there was a significant mean increase in perceptions of empathy, a notable minority of participants reported only a modest change, illustrating the individualized ways in which AI-based support is interpreted.

Several interviewees highlighted the influence of cultural factors. As one user (Participant 07) explained,

*“In Chinese culture, we’re not always comfortable talking about personal struggles with friends or family. An AI might feel less ‘embarrassing’ because you’re not burdening any-*

*one.”* (P07)

For this participant, Ernie Bot’s *“non-human”* status paradoxically served as an advantage—allowing them to disclose sensitive issues without fearing interpersonal judgment or stigma. This perspective points to the context-dependent nature of anthropomorphism, where an AI’s perceived empathy can be shaped by cultural norms around emotional expression and interpersonal boundaries.

Interviews and reflective diaries reveal that the co-construction of meaning with Ernie Bot often emerges through iterative reframing, with users selectively integrating AI-proposed language or perspectives. While many participants acknowledged that Ernie Bot lacks genuine emotional states, they still found the mechanics of its empathic style valuable for introspection and personal growth. Variations in skepticism or cultural attitudes underscore that anthropomorphizing an AI is neither universal nor uniform. However, the interview evidence suggests that even minimal emotional resonance or a sense of non-judgmental listening can be sufficient to catalyze new ways of thinking, reinforcing the argument that shared language (albeit algorithmically generated) can significantly shape user narratives.

## **Discussion**

The present findings suggest that Large Language Models (LLMs) like Ernie Bot can serve as external cognitive resources, guiding users toward new or reframed perspectives in their personal narratives. This dynamic aligns with the view of extended cognition (Hutchins, 1995), wherein tools external to the human mind can scaffold thinking processes. As participants repeatedly engaged in dialogue with the LLM, they progressively integrated AI-generated language, metaphors, and insights, effectively broadening their interpretive frameworks. Such patterns echo Bruner (1990) and (McAdams, 2008) discussion of story-making in narrative psychology, wherein the presence of an *“other”* interlocutor—even a non-human one—can prompt the re-narration of personal events. The observed increase in lexical overlap and self-reported adoption of AI frames underscores how repeated, empathic prompts from Ernie Bot formed a scaffolding that participants used to re-script key life challenges. These interactions extend earlier arguments about the co-construction of meaning typically found in therapeutic settings Angus and McLeod (2004), illustrating that algorithmic dialogue partners can similarly enact, at least in part, the role of a narrative catalyst.

The data also reveal how users increasingly imbued Ernie Bot with agency or empathy, paralleling real-world interpersonal dynamics in which conversational partners align cognitively or emotionally (Pickering & Garrod, 2004). Despite Ernie Bot’s lack of genuine subjective experience, many participants showed a readiness to treat the LLM as a quasi-social agent—particularly when it consistently mirrored user language and posed supportive questions. This phenomenon intersects with theory of mind debates in AI (Nadarzynski

et al., 2019), as some individuals reported feeling “heard” or “understood” despite being fully aware of the bot’s computational underpinnings (Koulouri, Macredie, & Olakitan, 2022). These findings point to a nuanced form of anthropomorphism: while participants acknowledged Ernie Bot’s artificial nature, they still experienced a perceived empathic resonance. The tension between knowing the AI is “just a tool” yet experiencing it as empathic underscores how mental health–style dialogues can blur the lines between mere simulation and user-perceived authenticity, a matter of critical interest for cognitive science and ethical AI design (Heinrichs, 2020).

This study makes several methodological advances by taking a longitudinal, mixed-methods approach to analyzing mental health–style interactions with an AI. Whereas previous research often relied on single-session data (Bendig et al., 2019), the present design captured repeated dialogues, reflective diaries, and exit interviews spanning four weeks. This structure offered a temporal lens on how users might gradually incorporate or reject AI-suggested frames, rather than focusing on one-off impressions (Fitzpatrick et al., 2017). By combining conversation analysis, computational linguistic markers (e.g., lexical overlap, emotional tone), and qualitative diaries/interviews, we illustrate how co-constructed meaning can be assessed from multiple angles—particularly crucial for exploring subtle narrative evolutions (Ghavibazou et al., 2024). Importantly, the approach underscores the ecological validity of collecting real-time data (transcripts, diaries) and triangulating those findings with in-depth interviews, paving the way for future longitudinal studies that wish to probe the mechanics of human–AI co-construction in naturalistic settings.

The findings not only expand our theoretical understanding of human–AI interaction but also bear practical relevance for mental health support. On the benefits side, LLM-driven platforms provide a reflective, on-demand space where individuals can articulate concerns and potentially reframe negative thought patterns. This 24/7 accessibility may be especially valuable in contexts where users are hesitant to seek traditional therapy or need supplementary emotional outlets (Inkster et al., 2018; Vaidyam et al., 2019). The gradual changes in personal narratives and mood suggest that repeated dialogues with an empathic-seeming AI can foster a reflective habit, sometimes catalyzing shifts in perspective.

However, the study also highlights risks and ethical dilemmas. Chief among these is the possibility of over-reliance on an entity that lacks genuine human empathy, where users might misattribute emotional intelligence to a purely algorithmic system. Such illusions of understanding could lead to missed opportunities for higher-level professional care, especially for individuals whose issues may escalate beyond mild-to-moderate distress. Moreover, the co-construction of meaning with a non-human partner raises concerns about data privacy and algorithmic shaping of user narratives, reminding developers and clinicians alike to establish transparent bound-

aries for AI-based mental health tools (Merilehto, 2024). In sum, while AI chatbots can play a beneficial supportive role, harnessing their potential effectively demands careful design, regulation, and user education to mitigate undue reliance or the perpetuation of misleading assumptions about AI’s emotional capabilities.

Although the study’s findings illuminate how users may co-construct meaning with LLMs over multiple sessions, several factors constrain the generalizability of these results. First, the small, self-selecting sample (N=16) and potential demographic homogeneity limit broader applications, and the study’s focus on mild-to-moderate concerns does not address individuals experiencing acute clinical conditions (Bendig et al., 2019). Additionally, because participants may already have been predisposed to view AI as helpful, self-selection bias could inflate positive impressions, underscoring the need for caution in extrapolating these outcomes to other user groups. Further complicating reproducibility is the ever-evolving nature of LLM technology, which can alter both the style and quality of responses over time. Moving forward, comparative research contrasting user–AI dialogues with traditional therapy or peer support may reveal the distinctive contributions of AI-driven conversation to narrative change. Extending this work to diverse cultural or clinical populations could clarify how linguistic or sociocultural factors shape AI-mediated emotional support (Bowman et al., 2024). Finally, ongoing design refinements—such as increased transparency regarding AI limitations—and deeper explorations into cognitive mechanisms (e.g., memory reconstruction, identity formation) may further illuminate the interplay between user agency, AI suggestion, and the broader process of meaning-making.

## Conclusion

In sum, our findings show that repeated engagements with an empathic-seeming LLM can progressively reorient users’ personal narratives, highlighting an interplay between algorithmic prompting and human meaning-making. While participants sometimes questioned the chatbot’s authenticity, the iterative process of responding to Ernie Bot’s reframing cues appeared to stimulate the re-organization of self-talk, encourage shifts in emotional tone, and deepen introspection. This dynamic underscores how non-human partners—though devoid of genuine emotional states—may still function as external cognitive scaffolds, reshaping personal narratives in ways akin to distributed cognition. At the same time, the study emphasizes the importance of ethical guardrails, including informed consent and clear disclaimers regarding AI’s limitations, to prevent over-reliance and potential confusion about the chatbot’s empathic capacity. As LLMs become increasingly pervasive in mental health contexts, ongoing research is needed to fine-tune their design, usage guidelines, and cultural fit, ensuring that these technologies contribute positively and responsibly to user well-being.

## References

- Abd-Alrazaq, A. A., Rababeh, A., Alajlani, M., Bewick, B. M., & Househ, M. (2020). Effectiveness and safety of using chatbots to improve mental health: systematic review and meta-analysis. *Journal of medical Internet research*, 22(7), e16021.
- Angus, L. E., & McLeod, J. (2004). *The handbook of narrative and psychotherapy: Practice, theory and research*. Sage.
- Bendig, E., Erb, B., Schulze-Thuesing, L., & Baumeister, H. (2019). The next generation: chatbots in clinical psychology and psychotherapy to foster mental health—a scoping review.
- Bergelson, I., Tracy, C., & Takacs, E. (2022). Best practices for reducing bias in the interview process. *Current urology reports*, 23(11), 319–325.
- Bowman, R., Cooney, O., Newbold, J. W., Thieme, A., Clark, L., Doherty, G., & Cowan, B. (2024). Exploring how politeness impacts the user experience of chatbots for mental health support. *International Journal of Human-Computer Studies*, 184, 103181.
- Bruner, J. (1990). *Acts of meaning: Four lectures on mind and culture* (Vol. 3). Harvard university press.
- Chalnick, A., & Billman, D. (1988). Unsupervised learning of correlational structure. In *Proceedings of the tenth annual conference of the cognitive science society* (pp. 510–516). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Du, X., An, P., Leung, J., Li, A., Chapman, L. E., & Zhao, J. (2024). Deepthink: Designing and probing human-ai co-creation in digital art therapy. *International Journal of Human-Computer Studies*, 181, 103139.
- Feigenbaum, E. A. (1963). The simulation of verbal learning behavior. In E. A. Feigenbaum & J. Feldman (Eds.), *Computers and thought*. New York: McGraw-Hill.
- Fitzpatrick, K. K., Darcy, A., & Vierhile, M. (2017). Delivering cognitive behavior therapy to young adults with symptoms of depression and anxiety using a fully automated conversational agent (woebot): a randomized controlled trial. *JMIR mental health*, 4(2), e7785.
- Frank, J. D., & Frank, J. B. (1993). *Persuasion and healing: A comparative study of psychotherapy*. JHU Press.
- Ghavibazou, E., et al. (2024). Dialogical origin of meaning making and co-construction of voices in psychotherapy. *Journal of Preventive Counselling*, 5(3).
- Heinrichs, J.-H. (2020). Artificial intelligence in extended minds: intrapersonal diffusion of responsibility and legal multiple personality. *Technology, anthropology, and dimensions of responsibility*, 159–176.
- Hill, J. A. C. (1983). A computational model of language acquisition in the two-year old. *Cognition and Brain Theory*, 6, 287–317.
- Hua, Y., Na, H., Li, Z., Liu, F., Fang, X., Clifton, D., & Torous, J. (2024). Applying and evaluating large language models in mental health care: A scoping review of human-assessed generative tasks. *arXiv preprint arXiv:2408.11288*.
- Hutchins, E. (1995). *Cognition in the wild*. MIT press.
- Inkster, B., Sarda, S., Subramanian, V., et al. (2018). An empathy-driven, conversational artificial intelligence agent (wysa) for digital mental well-being: real-world data evaluation mixed-methods study. *JMIR mHealth and uHealth*, 6(11), e12106.
- Jiang, Z., Xu, Y., Wang, Z., Han, Y., Pan, Z., Wang, Q., ... others (2024). Performances of popular large language models from china and abroad in chinese myopia care queries: A comparative analysis. *Investigative Ophthalmology & Visual Science*, 65(7), 1509–1509.
- Kivi, H. G., & Ghavibazou, E. (2023). Dialogical origin of meaning making and co-construction of voices in psychotherapy.
- Koulouri, T., Macredie, R. D., & Olakitan, D. (2022). Chatbots to support young adults' mental health: an exploratory study of acceptability. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 12(2), 1–39.
- Matlock, T. (2001). *How real is fictive motion?* Doctoral dissertation, Psychology Department, University of California, Santa Cruz.
- McAdams, D. P. (2008). *The person: An introduction to the science of personality psychology*. John Wiley & Sons.
- Merilehto, J. (2024). Inter-team cognitive diversity through distributed cognition. Available at SSRN 4712181.
- Miner, A. S., Laranjo, L., & Kocaballi, A. B. (2020). Chatbots in the fight against the covid-19 pandemic. *NPI digital medicine*, 3(1), 65.
- Molli, V. L. P. (2022). Effectiveness of ai-based chatbots in mental health support: A systematic review. *Journal of Healthcare AI and ML*, 9(9), 1–11.
- Nadarzynski, T., Miles, O., Cowie, A., & Ridge, D. (2019). Acceptability of artificial intelligence (ai)-led chatbot services in healthcare: A mixed-methods study. *Digital health*, 5, 2055207619871808.
- Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.
- Ohlsson, S., & Langley, P. (1985). *Identifying solution paths in cognitive diagnosis* (Tech. Rep. No. CMU-RI-TR-85-2). Pittsburgh, PA: Carnegie Mellon University, The Robotics Institute.
- Pickering, M. J., & Garrod, S. (2004). Toward a mechanistic psychology of dialogue. *Behavioral and brain sciences*, 27(2), 169–190.
- Sacks, H. (1978). *A simplest systematics for the organization of turn taking for conversation*. Academic Press.
- Shrager, J., & Langley, P. (Eds.). (1990). *Computational models of scientific discovery and theory formation*. San Mateo, CA: Morgan Kaufmann.
- Sools, A. M., & Mooren, J. H. (2012). Towards narrative futuring in psychology: Becoming resilient by imagining the future. *Graduate journal of social science*, 9(2), 203–226.

- Steger, M. F. (2022). Making meaning in life: A thematic review of successful experimental psychological and psychotherapeutic interventions. In *Meaning in life international conference 2022-cultivating, promoting, and enhancing meaning in life across cultures and life span (mil 2022)* (pp. 5–20).
- Sweeney, C., Potts, C., Ennis, E., Bond, R., Mulvenna, M. D., O’neill, S., . . . others (2021). Can chatbots help support a person’s mental health? perceptions and views from mental healthcare professionals and experts. *ACM Transactions on Computing for Healthcare*, 2(3), 1–15.
- Vaidyam, A. N., Wisniewski, H., Halamka, J. D., Kashavan, M. S., & Torous, J. B. (2019). Chatbots and conversational agents in mental health: a review of the psychiatric landscape. *The Canadian Journal of Psychiatry*, 64(7), 456–464.