Inside Family Caregivers' Conversations with a Well-being Chatbot in the Wild: Intent, Affect, and Supportive Moves from an 8-Week Deployment

Rongqi Bei
University of Michigan
Ann Arbor, USA
rbei@umich.edu

Predrag Klasnja
University of Michigan
Ann Arbor, USA
klasnja@umich.edu

Mark W. Newman
University of Michigan
Ann Arbor, USA
mwnewman@umich.edu

Noelle E. Carlozzi
University of Michigan
Ann Arbor, USA
carlozzi@med.umich.edu

Abstract

Family caregivers of people living with dementia face accumulating stressors, yet traditional mental health support is hard to reach amid unpredictable, demanding caregiving. Large language model (LLM)-based well-being chatbots have emerged as an on-demand stopgap: in our 8-week deployment, caregivers' self-reported anxiety, depression, anger, perceived stress, and sleep improved (medium-to-large effects). Yet how caregivers use such tools, and how the tools respond, remains underexplored. Through an LLM-assisted workflow, we analyzed the deployment's 13,637 chat-log messages between 22 caregivers and the chatbot for caregivers' communicative intent (speech act theory) and affect, and the chatbot's supportive moves (emotional-support and helping-skills frameworks). Caregivers used the chatbot mainly to express feelings and report situations, with individual variation; their conversations moved through three paths resembling coping processes. The chatbot enacted a consistent reflective style, mirroring emotion and asking questions, and left roughly a third of explicit guidance requests without advice in its immediate response. We contribute an empirical account of caregiver-chatbot interaction and raise three questions: how to make the chatbot's approach context-aware so it meets caregivers where they are; what to code, and at what unit, to tell whether supportive moves land; and how these interactions connect to what caregivers do beyond the conversation.

CCS Concepts

• **Human-centered computing → Human computer interaction (HCI)**.

Keywords

Family Caregivers; Conversational Agents; Large Language Models; Well-being Support; Chat Log Analysis

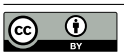

This work is licensed under a Creative Commons Attribution 4.0 International License.
UbiComp Companion '26, Shanghai, China
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2533-3/2026/10
https://doi.org/10.1145/3798063.3837289

ACM Reference Format:
Rongqi Bei, Predrag Klasnja, Mark W. Newman, and Noelle E. Carlozzi. 2026. Inside Family Caregivers' Conversations with a Well-being Chatbot in the Wild: Intent, Affect, and Supportive Moves from an 8-Week Deployment. In *Companion of the 2026 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp Companion '26), October 11–15, 2026, Shanghai, China.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3798063.3837289

1 Introduction

Family caregivers who support their loved ones living with dementia experience elevated levels of stress, depression, anxiety, and sleep disruption that compromise their mental well-being [1]. Yet, the same demands that undermine caregivers' well-being (e.g., unpredictable schedules and competing responsibilities) also limit their access to traditional support like in-person therapy and psychoeducation, which require significant time and advance planning [4, 13, 17]. Large language model (LLM)-based chatbots have emerged as a candidate stopgap, with on-demand availability that fits caregivers' schedules and dynamic responses to the diverse challenges they face [2, 16]. Our recently completed 8-week field deployment with 23 dementia caregivers offers preliminary evidence of this promise: findings showed medium-to-large improvements pre- to post-deployment on self-report measures across anxiety, depression, anger, perceived stress, and sleep, and companion interviews revealed positive but mixed experiences with how well the chatbot met caregivers' well-being needs. However, much less is known about the nature of these interactions over weeks of everyday use, including what caregivers seek to accomplish, what emotions they express, how the chatbot responds, and how conversations unfold. To date, caregiver-facing chatbots have mostly been studied through interviews and surveys that capture caregivers' self-reported needs and preferences after short, study-based sessions with a prototype [5, 16], rather than analyzing the conversations themselves. Naturalistic, longitudinal logs from this vulnerable population are accordingly rare, and the methods to characterize them are still developing.

We analyze the conversation logs from this deployment; one of the 23 caregivers did not share their logs, so our analysis includes the remaining 22 participants. To characterize the caregiver side of these interactions, we coded messages for communicative intent,

drawing on speech act theory [7, 15] and paired this with sentiment analysis. To characterize the chatbot side, we apply a move-coding scheme grounded in emotional-support dialogue and helping-skills frameworks [6, 9]. Both analyses were scaled with LLM assistance through U-M GPT, the University of Michigan's secure institutional LLM service.

This work makes three contributions. First, we provide an empirical account of caregiver-chatbot interactions across 8 weeks of real-world use, characterizing what caregivers do, how the chatbot responds, and how both sides of these naturalistic, longitudinal exchanges unfold in a population where such interactions are rarely examined. Second, we demonstrate an LLM-assisted approach for coding both sides of these support conversations on a secure institutional service. Third, we use this account to identify three open questions for community discussion about the responsible design and evaluation of conversational AI for caregivers' mental health how a chatbot can infer a caregiver's momentary capacity and adjust what it asks of them; what unit of exchange is needed to assess whether a supportive move was taken up; and how conversational data can be linked to behavior outside the app.

2 Methods

2.1 Study Dataset

The chat log data come from our 8-week deployment of an off-the-shelf LLM-based mental health chatbot with family caregivers of people living with dementia who reported mild-to-moderate caregiver burden, as assessed by the Zarit Burden Interview [3]. The chatbot, Earkick,[1] is designed around reflective listening and non-prescriptive emotional probing. It was introduced as a tool to support participants' own well-being, with no prescribed frequency of use or topic. Of the 23 participants who completed the 8-week study, 22 shared their chat logs, which form our corpus: 463 conversation sessions and 13,637 messages, including 3,147 caregiver messages, with the remainder generated by the chatbot. The chatbot accounts for more messages because it often splits a single turn into several shorter messages. The study was approved by the University of Michigan Institutional Review Board (HUM00271412), and all LLM-assisted coding described below was run on U-M GPT, the University of Michigan's secure institutional LLM service.

2.2 Coding Caregiver Communicative Intent and Chatbot Moves

We coded each caregiver message for its communicative *intent*. Drawing on speech act theory [15], which classifies utterances by their illocutionary force, and Jakobson's phatic function [7] for contact-maintaining acts, we organized intent into five directions: Assertive, Expressive, Directive, Commissive, and Phatic. We omitted declarations because they require institutional authority absent from this setting. Coding proceeded in three stages, combining LLM assistance with researcher judgment. First, we drew a 239-message sample using participant-stratified proportional allocation, with a floor of three messages per participant. GPT-5.2 open-coded each message with a free-text intent label grounded in the five intent directions and informed by preceding context (the prior caregiver message and intervening chatbot turns). The first author then consolidated these labels into a 13-category codebook (Table 1) within the five directions, assigning each message a single primary intent. Categories were refined using two heuristics: each had to reflect a distinct communicative purpose, and candidates with fewer than five exemplars were held for review and merged where appropriate. We then applied the codebook to the full corpus with GPT-4o; this reached 74.9% agreement with the author's labels on the 239-message sample.

To characterize the chatbot's turns, we adapted the emotional-support strategy scheme of Liu et al. [9], which builds on Hill's helping-skills taxonomy [6]. We tested Liu's eight strategies on a proportionally stratified 127-message sample, and found that they adequately captured the chatbot moves observed in our data. Because several recurrent moves (opening/check-in, closing, acknowledgment/rapport, safety/redirect) would otherwise collapse into Liu's residual category, we broke them out as dedicated codes. This yielded 12 categories that retain all eight of Liu's original strategies (Table 2). We coded one primary move per message, defined as the move carrying the message's main supportive work. We then applied the codebook to the full corpus with LLM assistance. On a held-out set of 100 messages, model labels agreed with the author's labels at 85.0% (Wilson 95% CI [76.7, 90.7]; Cohen's $\kappa = 0.816$).

A limitation is that the codebooks, although developed within published frameworks, reflect a single researcher's interpretation. Model and author labels diverged when a message carried more than one act, and the single primary code forced a choice, most often between venting and reflection on the caregiver side and between restatement and reflection of feelings on the chatbot side. These disagreements exposed implicit codebook boundaries, which we converted into explicit decision rules before coding at scale.

2.3 Sentiment Analysis

We ran two off-the-shelf RoBERTa classifiers on every caregiver message. First we used a sentiment classifier[2] to compute continuous valence as $P(\text{positive}) - P(\text{negative})$, a score in $[-1, +1]$. Second, we used a GoEmotions classifier[3] trained on the 27 emotion categories to obtain discrete emotion labels, which we use descriptively to characterize the emotions underlying the observed valence shifts. Both classifiers were used as published.

3 Findings

3.1 What Caregivers' Messages Do

Across the 8-week deployment, caregivers' messages were dominated by Expressive and Assertive acts, each accounting for roughly 30-40% of messages. Phatic acts made up about 15-22%, Directive acts were relatively rare at about 2-10%, and Commissive acts accounted for only a small remainder (Figure 1a). Caregivers were predominantly narrating situations, expressing feelings, and reflecting, rather than requesting guidance. This pattern likely reflects two co-occurring forces: how caregivers preferred to use the space, and what the chatbot's style encouraged (Earkick's reflective stance tends to invite continued narration and disclosure rather

[1] https://earkick.com/

[2] https://huggingface.co/cardiffnlp/twitter-roberta-base-sentiment-latest

[3] https://huggingface.co/SamLowe/roberta-base-go_emotions

Table 1: Caregiver intent codebook: 13 categories across five communicative directions, grounded in speech act theory [15].

Direction	Category	Definition
Assertive	Daily status/situational reporting	Reports what is happening, what happened, or what is planned (e.g., events, activities, routines, states).
	Explaining/reasoning	Provides a cause, hypothesis, or mechanism to explain the external situation or another person's behavior.
	Stating/evaluating coping strategies	Identifies or evaluates a strategy or approach for managing the caregiving situation, emotions, or well-being.
Expressive	Emotional state disclosure	Briefly names or labels a current emotional state, without extended narrative.
	Venting/narrating distress	Recounts stressful events with the primary purpose of offloading emotional weight.
	Reflecting/meaning-making	Interprets the caregiver's own relationship to an experience (e.g., meaning, identity, values, or growth).
	Ambivalence/moral conflict	Articulates conflicted feelings or moral tension (e.g., guilt, push-pull about decisions, competing values).
	Positive affect/gratitude	Shares positive emotion (e.g., joy, relief, hope, gratitude, or spiritual comfort) rather than a state.
Directive	Requesting practical guidance	Asks for concrete advice, information, instructions, options, or strategies.
	Managing chatbot interaction	Instructs the chatbot on communication style, sets behavioral boundaries, grants/withholds consent, or signals (dis)engagement.
Phatic	Acknowledging/affirming	Signals agreement, validation, appreciation, or receipt without adding substantive content.
	Opening/closing the channel	Greets, introduces, signs off, or signals unavailability to continue.
Commissive	Committing to future action	Pledges a specific future action (e.g., trying a suggestion, implementing a plan, returning later).

Table 2: Chatbot move codebook: 12 categories, grounded in emotional-support dialogue and helping-skills frameworks [6, 9].

Move	Description
Question	Invites the caregiver to share, elaborate on, or explore their situation, thoughts, or feelings.
Reflection of feelings	Names or mirrors back the caregiver's emotion, including one implied but not stated.
Restatement/paraphrasing	Rephrases or mirrors the content of what the caregiver said.
Affirmation/reassurance	Affirms strengths or efforts, validates or normalizes a reaction, or reassures.
Providing suggestions	Offers a concrete coping strategy, action, or next step.
Information	Provides facts, resources, psychoeducation, or an explanation.
Self-disclosure	Expresses, in the first person, a simulated reaction or shared experience.
Acknowledgment/rapport	Minimal receipt or relational move; brief acknowledgment or banter.
Opening/check-in	A session-initiating greeting that invites the caregiver to begin.
Closing	Signs off, wishes the caregiver well, or wraps up the session.
Safety / redirect	Flags risk or crisis and urges professional or emergency help.
Miscellaneous	Primary function fits none of the above (e.g., app logistics).

than question-and-answer exchange). This aggregate mix was relatively consistent across the 8 weeks, with only a modest shift observed. However, individuals differed in how they engaged with the chatbot. Caregivers' intent profiles followed several recognizable patterns: some used the chatbot primarily for Expressive acts, others leaned toward Directive or Assertive, and a few were more Phatic or balanced (Figure 1b).

3.2 How the Chatbot Responds

The chatbot's support moves were consistent with a reflective, or active, listening stance (Figure 2). Its most common responses were asking a question that invited the caregiver to say more (`Question`) and naming the emotion the caregiver had expressed (`Reflection of feelings`), followed by validating or reassuring the caregiver (`Affirmation`) and paraphrasing the content of what they had said (`Restatement`). Together these made up roughly three-quarters of all chatbot messages. In contrast, concrete problem-solving was uncommon: offering advice (`Providing suggestions`) and sharing factual information (`Information`) together account for less than 8% of chatbot messages.

Two patterns also emerged in how the chatbot sequenced moves within a turn: the consecutive messages it sent before the caregiver replied. First, `Reflection of feelings`, `Affirmation`, and `Restatement` formed a tightly connected cluster, transitioning among one another more often than to any other move. Second, a binomial test on move position showed that this cluster appeared above chance in the opening message of a turn, and `Question` and `Providing suggestions` above chance in its closing messages. The characteristic turn thus opens by reflecting the caregiver's utterance back and closes by handing the floor back with a question.

Within each turn, the chatbot's opening support move was contingent on the caregiver's communicative intent ($\chi^2(132) = 1711.4$, $p < .001$, Cramér's $V = .23$), a pattern concentrated in a few strongly coupled intent–move pairs.[4] Across four of the caregivers' Expressive intents (i.e., venting, emotional state disclosure, ambivalence and moral conflict, and reflecting and meaning-making), the chatbot's first move was predominantly `Reflection of feelings` (56-71%; adjusted residuals +4.1 to +9.8), returning the expressed emotion to the caregiver. When caregivers instead committed to a future action, the dominant opening move shifted to `Affirmation` (71%; adjusted residual +12.3), validating the stated commitment. Directive requests for practical guidance, by contrast, elicited no single dominant opening move. Widening the lens to the full chatbot turn, 69% of multi-message turns following a guidance request contained at least one advice-bearing move (i.e., `Providing`

[4] p-values were Benjamini-Hochberg corrected for multiple comparisons; all reported residuals are significant at $q < .001$.

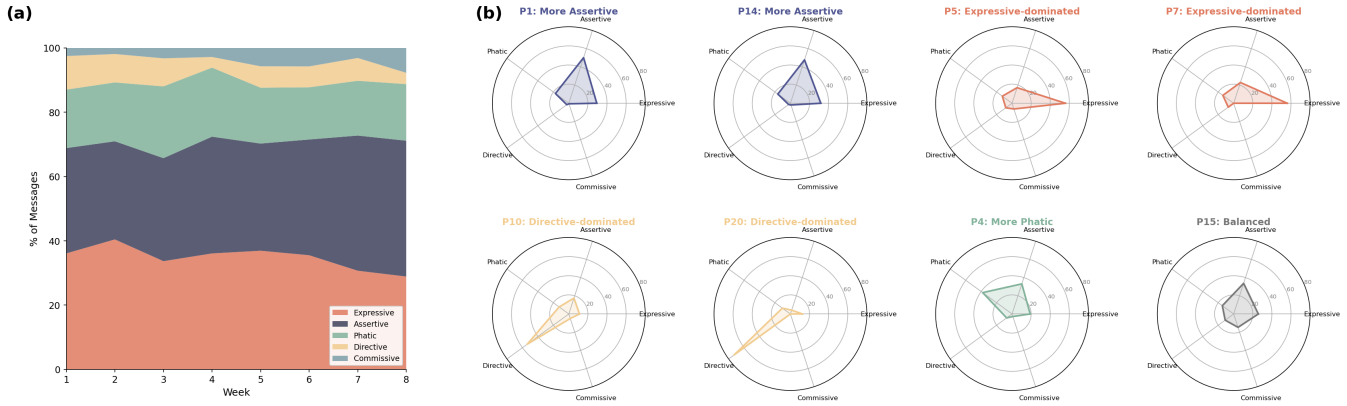

Figure 1: Caregiver communicative intent across the deployment. (a) Weekly share of messages by intent act, aggregated across caregivers. (b) Per-caregiver intent profiles for eight representative participants.

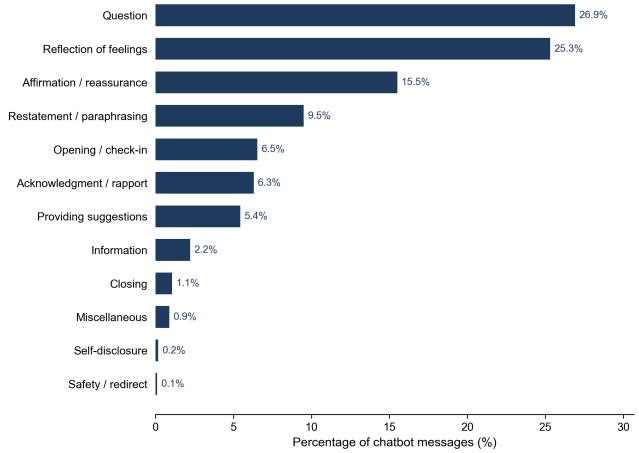

Figure 2: Distribution of chatbot moves (N = 10,490 messages).

`suggestions` and `Information`). Yet, this also means that nearly one-third of these turns contained no advice-bearing move at all.

3.3 How Conversations Unfold

We first asked where each intent tended to appear within a session. For sessions with at least three caregiver messages, we assigned each message a within-session position (i.e., opening, middle, or closing) based on its relative ordinal location. We then compared each intent category's share at each position with its corpus-wide share. Openings over-represented daily status reporting (+17.6%) and emotional state disclosure (+9.5%), suggesting that caregivers often began by setting the scene, reporting what was happening and how they felt. Requesting practical guidance was also over-represented in the openings (+8.7%), indicating that when caregivers came with a question, they tend to raise it up front. Middles carried more deeper processing intent, over-representing reflecting or meaning-making (+2.4%), which was notably rare in openings (-9.3%). Closings shifted toward conversational maintenance, over-representing acknowledging (+8.6%), opening/closing acts (+7.4%), and committing to future action (+3.6%). Deeper-processing intents (i.e., reflecting/meaning-making, explaining/reasoning, and stating/evaluating coping) also increased from 7.7% of messages in single-message sessions to 33.6% in sessions of eight or more messages.

We next examined whether caregivers' communicative intents unfolded in recurring sequences. We collapsed runs of consecutive same-intent messages so that transitions captured movement between distinct communicative intents. For each ordered intent pair, we compared how often one intent followed another with the following intent's overall frequency in the collapsed sequence. The fifteen most over-represented transitions were organized into three recurring paths (Figure 3).[5] **(a) A reflective processing loop** moved from daily status reporting through reasoning about the situation and evaluating coping strategies into reflective meaning-making, with reflecting sometimes looping back to explaining. This suggests caregivers oscillated between interpreting personal meaning and re-examining the external situation that produced it. **(b) Expression into contextualization**: after emotional disclosure, caregivers often grounded the feeling in a concrete event or moved into deeper interpretation; ambivalence similarly led to reflecting or explaining. **(c) A short task-oriented exchange**: requests for practical guidance were often followed by acknowledgment, likely reflecting caregivers posing a question, receiving an answer, and signaling receipt.

3.4 What Happens to Affect

We examined within-session affective movement. For sessions of at least five caregiver messages, we partitioned messages into opening, middle, and closing thirds and stratified sessions by the mean valence of their opening third (Figure 4a). Sessions that began with negative affect moved toward neutral over time; sessions that began with positive affect drifted mildly downward toward neutral; and sessions that began neutral remained neutral. For sessions that began with negative affect, the discrete emotions showed that this

[5]These paths are the caregiver-side projection of a dyadic exchange, in which each caregiver message follows a chatbot reply; future work should examine how chatbot replies shape these paths.

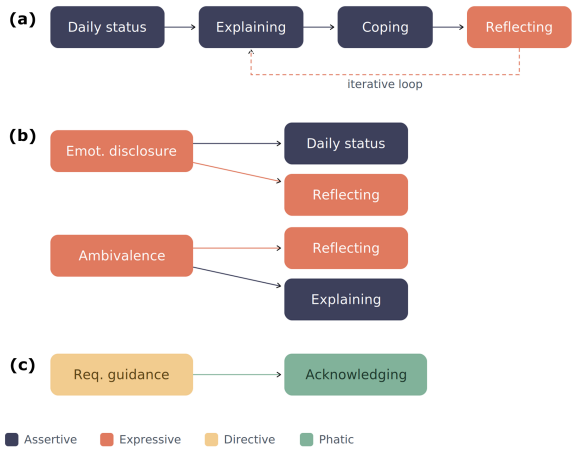

Figure 3: Three recurring within-session transition paths among caregiver communicative intents: (a) a reflective processing loop, (b) emotional expression into contextualization, and (c) a short task-oriented exchange.

movement was a recomposition rather than a simple fading of affect (Figure 4b). Openings carried acute negative emotions, especially sadness, alongside disappointment and annoyance. By the middle of sessions, these emotions had softened: sadness fell sharply, while approval became the most prominent emotion. Closing segments centered more clearly on approval, admiration, and gratitude. We are cautious about interpreting this shift as evidence of therapeutic mechanism: stratifying on negative openings invites regression toward the mean, and closing turns often carry a positive register of acknowledgment and farewell. We therefore treat the pattern as descriptive regularity and return in the discussion to what would be required to measure in-session change.

4 Discussion

In this section, we examine the implications of these findings for the design and evaluation of well-being chatbots for family caregivers.

The three within-session paths can be interpreted in light of the literature on stress and coping, which distinguishes problem-focused, emotion-focused, and meaning-focused coping [8, 11]. In the stress-process model, these coping approaches mediate between caregiving stressors and caregivers' well-being [12]. The reflective processing loop resembles meaning-focused coping; the path from emotional disclosure into contextualization sits between emotion regulation and meaning-making; and the request-to-acknowledgment exchange resembles problem-focused coping. When conversations followed these paths, caregivers' messages reflected forms of communicative work that the coping literature associates with adjustment. This offers a conversation-level view of what the medium-to-large improvements in self-reported outcomes looked like during the 8 weeks caregivers spent with the chatbot, while stopping short of identifying the mechanisms that produced those outcomes.

The chatbot side of these conversations had a consistent character. It relied on questions, reflection, and affirmation and rarely offered advice. This pattern aligns with active-listening practice [14] and with the core skills of motivational interviewing [10], in which withholding immediate advice is a deliberate supportive move rather than a failure to help. This approach fits expressive and reflective intents well, but is less aligned with explicit requests for guidance: as §3.2 shows, when caregivers requested guidance, nearly a third of the chatbot's subsequent turns contained no advice-bearing move. This can be viewed as a gap in responsiveness, or as faithful reflective practice, and how we interpret this depends on what a given caregiver needed in that moment. Two considerations follow. The first concerns the caregivers' momentary availability. Our findings show that reflecting/meaning-making intent appeared primarily in longer exchanges (§3.3), yet caregiving is often fragmented, interrupted, and cognitively taxing [13]. The moments when a caregiver most needs to work through a difficult feeling may also be the moments when they are least able to sustain a multi-turn conversation. The most therapeutically promising path the chatbot offers may therefore be the one caregivers are least able to follow when they need it. Inferring caregivers' momentary availability from signals such as time of day, location, input modality, and message length is a context-inference problem that ubiquitous computing is well positioned to address. A system attuned to these cues can compress reflective work when a caregiver's bandwidth is low, or ensure a concrete answer arrives when someone is plainly rushed. The second concerns differences between caregivers rather than within a session: caregivers differed in the intents they brought to the chatbot, with some noticeably more directive than others (§3.1). If such differences prove stable enough to act on, this raises the possibility that a structured, advice-forward style serves some caregivers while a more reflective style may serve others. Whether and how the reflect-advise balance should adapt remains open.

Moving toward answers to questions would require evaluation at a finer grain, examining what happens from one message to the next (e.g., whether a caregiver has the bandwidth to take up a reflective prompt, whether a `Question` actually draws them into the exchange it opens). The move coding we demonstrate supports a fidelity check: it shows whether the chatbot behaved consistently with its intended reflective design, as it did in our deployment. What it does not tell us is whether caregivers took up those moves. Understanding which moves caregivers take up, and whether those moves support therapeutically desirable forms of engagement, would allow us to shape the chatbot's behavior more deliberately. The caregiver side poses a parallel question about the appropriate coding target. Our current communicative-intent scheme captures what a message does, not whether the caregiver is engaging in a therapeutically desirable process (e.g., reframing a distressing thought toward a more balanced appraisal or reporting that a previously discussed strategy worked in practice). Capturing such processes would require a different coding target. Pursuing this finer grain is not straightforward. It raises additional questions for the field, including whether LLM-assisted coding remains reliable at that level of granularity, and how to interpret message-to-message correspondence in real deployments, where a chatbot turn may be split across several messages, and a caregiver may take up a prompt several turns later, or not at all.

Evaluation can also extend beyond the conversation, linking chatbot interactions to changes in caregivers' real-world behavior, both beneficial and harmful. This is another place where ubiquitous computing could contribute. Together, these questions indicate

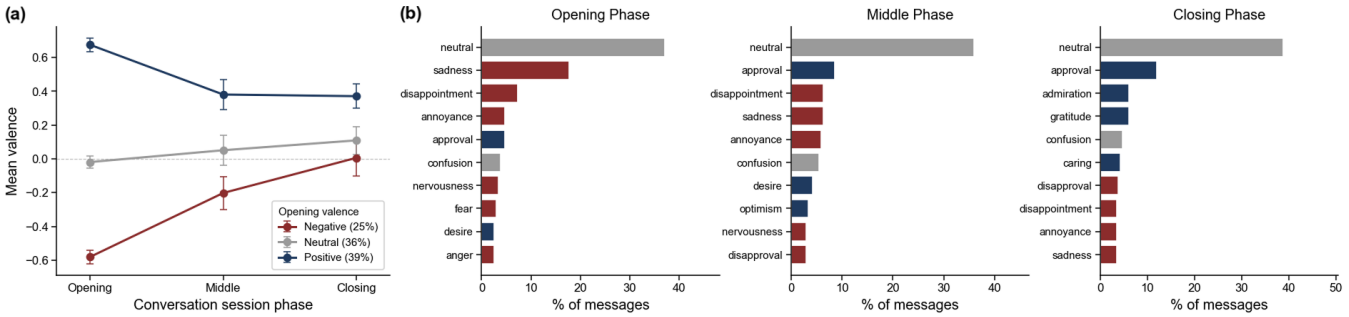

Figure 4: (a) Within-session valence by opening valence; (b) top emotions at the opening, middle, and closing phase of negatively-starting sessions.

where evaluation must develop: in grain, by examining message-to-message dynamics, and in reach, by connecting conversational patterns to behavior beyond the app. We offer them to this community as questions worth working through collectively.

5 Conclusion

We analyzed 13,637 messages from an 8-week deployment of an LLM-based mental health chatbot with 22 dementia caregivers, characterizing caregivers' communicative intent and affect alongside the chatbot's support moves. Across caregivers, messages more often expressed feelings and reported situations than requested guidance, though individuals varied in how they used the chatbot. The chatbot enacted a consistent reflective style that aligned with caregivers' predominant expressive intents, yet nearly a third of explicit guidance requests drew no advice in its immediate response. Within sessions, caregivers' messages moved through three recurring intent paths that we read against emotion-, meaning-, and problem-focused coping; and in sessions that began negative, valence moved toward neutral as the mix of expressed emotions shifted. Based on these findings, we raise three questions. The first is a design question: because reflection asks for bandwidth that caregivers cannot always spare, how should a chatbot balance reflective support with more substantive input, adapting to what a caregiver can take up in the moment? The other two questions concern evaluation. One asks what to code, and at what unit, to determine whether the chatbot's moves are taken up, and whether therapeutically desirable or undesirable patterns follow in the caregiver's response. The other extends beyond the conversation to real-world behavior, asking how to connect these interactions to what caregivers do outside the conversation. This work contributes empirical findings from field deployment and raises open questions for designing and evaluating chatbots for caregivers' mental health.

Acknowledgments

We thank our participants for giving their time amid caregiving responsibilities.

References

[1] Wilson Abreu, Debbie Tolson, Graham A Jackson, and Nilza Costa. 2020. A cross-sectional study of family caregiver burden and psychological distress linked to frailty and functional dependency of a relative with advanced dementia. *Dementia* 19, 2 (Feb. 2020), 301–318. doi:10.1177/1471301218773842

[2] Sahar Borna, Michael J. Maniaci, Clifton R. Haider, Cesar A. Gomez-Cabello, Sophia M. Pressman, Syed Ali Haider, Bart M. Demaerschalk, Jennifer B. Cowart, and Antonio Jorge Forte. 2024. Artificial Intelligence Support for Informal Patient Caregivers: A Systematic Review. *Bioengineering* 11, 5 (May 2024), 483. doi:10.3390/bioengineering11050483

[3] Michel Bédard, D. William Molloy, Larry Squire, Sacha Dubois, Judith A. Lever, and Martin O'Donnell. 2001. The Zarit Burden Interview: A New Short Version and Screening Version. *The Gerontologist* 41, 5 (Oct. 2001), 652–657. doi:10.1093/geront/41.5.652

[4] Noelle E. Carlozzi, Sung Won Choi, Zhenke Wu, Jonathan P. Troost, Angela K. Lyden, Jennifer A. Miner, Christopher M. Graves, Jitao Wang, Xinghui Yan, and Srijan Sen. 2022. An app-based just-in-time-adaptive self-management intervention for care partners: The CareQOL feasibility pilot study. *Rehabilitation Psychology* 67, 4 (2022), 497–512. doi:10.1037/rep0000472

[5] Fernanda Espinoza, Darren Cook, Chris R. Butler, and Rafael A. Calvo. 2024. Supporting dementia caregivers in Peru through chatbots: generative AI vs structured conversations. In *Proceedings of the 36th International BCS Human-Computer Interaction Conference* (University of York, UK) *(BCS HCI '23)*. BCS Learning & Development Ltd, Swindon, GBR, 89–98. doi:10.14236/ewic/BCSHCI2023.11

[6] Clara Hill. 2020. *Helping skills: Facilitating exploration, insight, and action (5th ed.)*. doi:10.1037/0000147-000

[7] Roman Jakobson. 1960. Closing Statement: Linguistics and Poetics. In *Style in Language*, Thomas A. Sebeok (Ed.). MIT Press, Cambridge, MA, 350–377.

[8] Richard S. Lazarus and Susan Folkman. 1984. *Stress, Appraisal, and Coping*. Springer Publishing Company.

[9] Siyang Liu, Chujie Zheng, Orianna Demasi, Sahand Sabour, Yu Li, Zhou Yu, Yong Jiang, and Minlie Huang. 2021. Towards Emotional Support Dialog Systems. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (Eds.). Association for Computational Linguistics, Online, 3469–3483. doi:10.18653/v1/2021.acl-long.269

[10] William R Miller and Stephen Rollnick. 2012. *Motivational interviewing: Helping people change*. Guilford press.

[11] Crystal L. Park. 2010. Making sense of the meaning literature: An integrative review of meaning making and its effects on adjustment to stressful life events. *Psychological Bulletin* 136, 2 (2010), 257–301. doi:10.1037/a0018301

[12] Leonard I. Pearlin, Joseph T. Mullan, Shirley J. Semple, and Marilyn M. Skaff. 1990. Caregiving and the Stress Process: An Overview of Concepts and Their Measures1. *The Gerontologist* 30, 5 (Oct. 1990), 583–594. doi:10.1093/geront/30.5.583

[13] Francine N. F. R. Queluz, Emily Kervin, Lori Wozney, Pamela Fancey, Patrick J. McGrath, and Janice Keefe. 2020. Understanding the needs of caregivers of persons with dementia: a scoping review. *International Psychogeriatrics* 32, 1 (Jan. 2020), 35–52. doi:10.1017/S1041610219000243

[14] Carl R. Rogers and Richard E. Farson. 1987. Active Listening. In *Communicating in Business Today*, R. G. Newman, M. A. Danziger, and M. Cohen (Eds.). D. C. Heath. Original work published 1957.

[15] John R. Searle. 1976. A Classification of Illocutionary Acts. *Language in Society* 5, 1 (1976), 1–23. doi:10.1017/S0047404500006837

[16] Jiayue Melissa Shi, Dong Whi Yoo, Keran Wang, Violeta J. Rodriguez, Ravi Karkar, and Koustuv Saha. 2026. Mapping Caregiver Needs to AI Chatbot Design: Strengths and Gaps in Mental Health Support for Alzheimer's and Dementia Caregivers. *ACM Trans. Comput. Healthcare* (March 2026). Just Accepted. doi:10.1145/3803549

[17] Kyra J Waligora, Mona N Bahouth, and Hae-Ra Han. 2019. The Self-Care Needs and Behaviors of Dementia Informal Caregivers: A Systematic Review. *The Gerontologist* 59, 5 (Sept. 2019), e565–e583. doi:10.1093/geront/gny076