

ChatWise: A Strategy-Guided Chatbot for Enhancing Cognitive Support in Older Adults

Zhengbang Yang
zyang30@gmu.edu
George Mason University
Fairfax, Virginia, USA

Junyuan Hong
University of Texas at Austin
Austin, Texas, USA

Yijiang Pang
Michigan State University
East Lansing, Michigan, USA

Jiayu Zhou
University of Michigan
Ann Arbor, Michigan, USA

Zhuangdi Zhu
zzhu24@gmu.edu
George Mason University
Fairfax, Virginia, USA

Abstract

Cognitive health in older adults presents a growing challenge. Although conversational interventions show feasibility in improving cognitive wellness, human caregiver resources remain overloaded. AI-based chatbots have shown promise, yet existing work is often limited to implicit strategies or heavily depends on training and label resources. In response, we propose a strategy-guided AI chatbot named CHATWISE that follows a dual-level conversation reasoning framework. It integrates macro-level strategy planning and micro-level utterance generation to enable engaging, multi-turn dialogue tailored to older adults.

CCS Concepts

• **Human-centered computing** → **Empirical studies in interaction design**.

Keywords

Conversational modeling, Cognitive health, LLM-driven chatbot, Conversational interventions

ACM Reference Format:

Zhengbang Yang, Junyuan Hong, Yijiang Pang, Jiayu Zhou, and Zhuangdi Zhu. 2025. ChatWise: A Strategy-Guided Chatbot for Enhancing Cognitive Support in Older Adults. In *Proceedings of Large Language Models for Scientific and Societal Advances (KDD'25 SciSoc LLM Workshop)*. ACM, New York, NY, USA, 12 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 Introduction

The cognitive well-being of the elderly population is a pressing social concern, as evidenced by the prevalence of cognitive disorders within this population, often exacerbated by loneliness and isolation [18, 31]. According to WHO, approximately 14% of adults aged 60 and over experience mental health disorders, projected to affect 2.1 billion individuals by 2050 [19]. Compared with other age

groups, older adults are more vulnerable due to age-related changes in cognitive reserves [24] and reduced social connections [18, 31]. Such impact extends beyond older individuals to families and society, resulting in a reduced life quality and increased medical burden. In the meantime, modest delays in cognitive decline can significantly reduce dementia prevalence, and addressing social isolation could prevent 4% of dementia cases [16]. Interventions through guided conversations have shown efficacy in reducing loneliness and mitigating cognitive decline [9, 38–40]. However, effective intervention requires sustained interaction and monitoring, which is limited by the availability of human companions, leading to inconsistent access or effectiveness.

Advances in artificial intelligence (AI) particularly Large Language Models (LLMs) have shown promise in augmenting human expertise with conversational support. Existing efforts range from sentiment-based chatbots [1, 15, 17, 23, 32, 37] to audio assistants [12, 36], ascribed to mature text-audio transformation techniques [22, 33]. However, these systems primarily default to interactions with implicit goals, which may fail to drive engaging conversations that necessitate strategic multi-turn interactions tailored to older adults. Some work has explored fine-tuning LLMs on domain-specific datasets to adapt to elderly care [29], which requires extensive labeled data and resources, while they may struggle to generalize across conversational contexts. In contrast, we take an orthogonal approach by focusing on the inherent reasoning capabilities of LLMs through in-context learning, which offers more resource efficiency and adaptability.

Our goal is to provide AI-powered, *engaging conversational support* for older adults that serves as an accessible complement to human companions, with the aspiration to improve their cognitive function and reduce social loneliness. Previous clinical studies aimed at socially isolated older adults showed the existence of causal relationships between interviewer strategy and interviewee response [5], which revealed that conversation behavior can have a *measurable* influence on interlocutors, and thus inspired us to develop principled yet efficient methods that transform clinical insights into dialogue design by leveraging the advanced reasoning ability of LLMs.

In response, we propose an LLM-driven chatbot named CHATWISE. It employs *dual-layer* conversation generation that first derives categorized macro-level information to suggest meta-conversational strategies, which then guides the micro-level utterance generation

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD'25 SciSoc LLM Workshop, Toronto, ON, Canada

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-XXXX-X/2018/06
<https://doi.org/XXXXXXX.XXXXXXX>




Figure 1: Alignment between conversational strategies of AI chatbots and human professionals across diverse participant data. Our design shows closer alignment with caregiver behavior in real conversation contexts compared to baseline chatbots. See Sec 4 for full results.

to improve both user engagement and cognitive outcomes over multi-turn dialogue interactions. CHATWISE is evaluated using both real, de-identified real-world dialogues between older adults and professional human caregivers of a clinical trial [13], and synthetic interactive conversations using simulated users (*i.e.* digital twins) [12] modeled on such data.

Our work provides multifold contributions: (i) We introduce a clinically grounded chatbot that highly aligns with the behavior of professional human caregivers to older adults (Sec 4.1). Figure 1 overviews its alignment performance. It also empirically enhances simulated users’ engagement and cognitive status, which significantly outperforms baseline chatbots (Sec 4.2). (ii) CHATWISE follows a tuning-free, in-context learning framework for daily conversational support. Comparative studies demonstrated that providing macro-level strategies to guide conversation generation is the key contributor to enhancing user engagement. (iii) Its dual-level policy design can be readily applied to various LLMs, with consistent user cognitive gains across different backbone LLMs.

2 Related Work

AI-powered Chatbots for Older Adults. have shown feasibility in preventing or detecting the cognitive issues to assist human professionals. Recent efforts span commercial products [1, 2] and research prototypes focused on emotional support and audio assistance. [29] performed supervised fine-tuning on LLMs to enhance their performance on specialized nursing and elderly care tasks. Yang et al. [36] introduced an LLM-based voice assistant designed to bridge communication between older adults and their healthcare providers. Liu et al. [15] built an annotated dataset to tune LLMs for emotional support tasks. Ryu et al. [23] developed a chatbot for the mental health of the elderly. Hong et al. [12] constructed digital twins of the elderly. Unlike prior work, our approach prioritizes user conversational *engagement* through in-context reasoning, which has been empirically shown to enhance user cognitive status.

Dialogue Systems for Mental Health or Cognitive Stimulation: Recent studies have leveraged LLMs for the augmentation of emotional support conversations through *generated* dialogue data [14, 44]. Xyngkou et al. [35] assessed the acceptability of GPT-4-based conversational agents among people with dementia (PwD), and

highlighted the importance of design considerations for this sensitive population. Similarly, Favela et al. [8] implemented a conversational robot embedded with ChatGPTs for reminiscence therapy, with authors as simulated users, which showcased the potential of generative AI to support cognitive stimulation through interactive storytelling. Distinct from these works, our research introduces a dual-level dialogue approach combined with comprehensive evaluation methodologies to provide a finer-grained understanding of conversational engagement.

Conversational Strategies: Liu et al. [15] curated a dataset with annotated strategies, demonstrating the effectiveness of *Helping Skills Theory* [11] in providing emotional support. Yuan et al. [42] examined the causal relationships between dialogue acts (DAs) and participants’ emotional states in a clinical trial [13], emphasizing the impact of strategic interventions in tele-mediated dialogues. Seo et al. [27] identified key strategies for improving child patient-provider communication through semi-structured interviews. However, few works have systematically integrated these strategies into the reasoning flow of AI chatbots. Our approach fills this gap by contextualizing structured conversational strategies for automatically enhancing user engagement.

Multi-Turn Chatbot Exploration: Recent advances in AI dialogue systems predominantly focused on short-turn interactions [6, 10, 20], with relatively limited attention to the challenges of multi-turn exchanges. While some pioneering research has explored multi-turn optimization through Reinforcement Learning (RL) approaches for LLMs [3, 10, 34, 45], these methods were not specifically designed for supporting senior dialogue engagement. Our inference-based method introduces a strategy-compatible framework that aligns with the principles of RL-driven optimization, which can enable future extension to further enhance multi-turn conversation engagement.

3 CHATWISE Design

CHATWISE encompasses a hierarchical framework that integrates the insights of traditional clinical studies to drive LLM reasoning. It features a dual-level design with a *strategy provider* π_s that generates macro-level conversation strategies, which guides a *utterance generator* π_u to produce micro-level responses. In this paper, we interchangeably refer to CHATWISE as the *interviewer* or the *moderator*, a role commonly defined in communication studies that guides and facilitates the conversation [30]. For clarity, we denote a multi-turn conversation as $\tau = \{\mu_0, u_0, x_1, u_1, x_2, u_2, \dots\}$ which comprises interviewee (user) utterances x_t and interviewer (CHATWISE) utterances u_t at time step t , with a conversation starting point μ_0 . $s_t = \{\mu_0, u_0, \dots, x_t\}$ denotes the historical conversation up to timestep t . We then formulate CHATWISE as a dual policy $\pi = \{\pi_s, \pi_u\}$, which maps the conversational *state* s_t to the next utterance u_t provided to the user, conditioned on meta information z given to the strategy provider π_s :

$$u_t \sim \pi(\cdot | s_t, z) \equiv \pi_u(\cdot | \pi_s(s_t | z)).$$

Below, we illustrate the process of this hierarchical conversation generation.

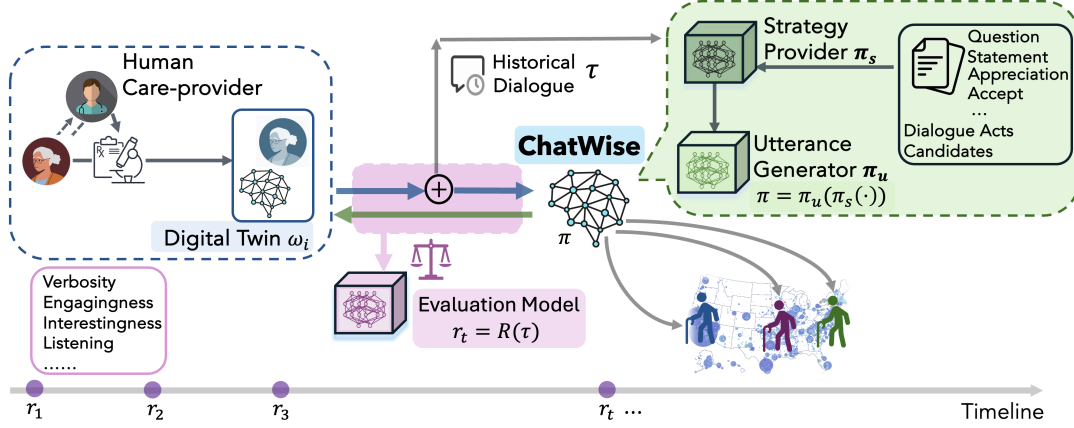


Figure 2: CHATWISE employs a dual-level policy design that contextualizes a clinically-derived *strategy pool* to generate *macro-level* actions, which then guide *micro-level* utterance generation. It is iteratively developed using digital twins as simulated users. Prior to a future real user study, real and simulated data analysis show that CHATWISE closely aligns with professional human caregivers’ behavior, leading to cognitive gains of simulated users accumulated over multi-turn interactions.

3.1 Strategy Pool

To first develop clinically grounded strategies, we establish a strategy pool containing dialogue acts (DA) curated from clinical intervention studies involving older adults. These DAs, such as *asking an open question* (e.g. “What is your favorite movie”) or *showing empathy* (e.g. “That sounds really meaningful”), serve as macro-level actions and atomic communicative units that convey distinct conversational intentions [25], which have also shown causal relationships with participant emotions in conversations [5]. Specifically, we extract DAs that exist in conversations from a real tele-health clinical trial [42], which represents a subset of the complete DAMSL taxonomy [4, 25]. We further augment this DA with strategies from prior emotional support dialogue dataset [15] to form a comprehensive strategy pool \mathcal{A} . We then treat each strategy serves as a macro-level action candidate $a \in \mathcal{A}$ and refine it with a *definition* and an *example* to serve as in-context information z for enhancing the reasoning of strategy π_s (See Appendix E for a complete DA set).

3.2 Emotion Annotation

To enhance context-awareness during conversations and support post-conversation analysis, the strategy provider π_s is additionally prompted to infer the user’s current emotion based on historical conversations s_t . At timestep t of the conversation, π_s will categorize user emotion into one of the five classes: *joy*, *neutral*, *sadness*, *anger*, and *surprise*, given a calibrated system prompt. Our comparative study shows that CHATWISE can lead to detectable positive emotion changes over multi-turn conversations, which aligns with conversations provided by human professionals (Sec 4.5).

3.3 From Strategy to Utterance Generation

The strategy provider π_s processes the dialogue history s_t , contextualized by the strategy pool information z to decide the most appropriate strategies for continuing a conversation. We limit each interviewer’s utterance to contain one or two strategies, i.e. $\mathbf{a}_t =$

$\{a_t^i\}_{i \leq 2} \sim \pi_s(\cdot | s_t, z)$. Our rationale follows counseling studies [43] that dialogue intentions can be *forward*, e.g. initiating new topics via an open question, or *backward*, e.g. responding with acknowledgment, or both to transit between intentions or topics. Thus, π_s is prompted to comply with one of the three conditions: (i) a forward strategy, (ii) a backward strategy, or (iii) a backward strategy followed by a forward strategy.

The selected strategies \mathbf{a}_t , and *optionally*, the user’s current emotion label e_t , and the conversation history s_t are used as inputs to the utterance generator π_u . Our ablation study showed that providing meta-level strategies has a dominant effect over providing emotion information alone for conversational engagement (Sec 4.3). To ensure natural flow, π_u will first improvise a few rounds as warm-ups before adopting the suggested strategies. We also draw on clinical guidance to let π_u encourage users to choose topics rather than imposing them. Figure 2 overviews this design.

4 Experiments

We conducted various experiments focusing on answering the following questions:

Q1. Does CHATWISE’s behavior align with human professionals when responding to real conversation context sampled from clinic trials?

Q2. How does CHATWISE compare to baseline AI alternatives in terms of multifaceted conversational engagement metrics?

Q3. How do factors such as the conversation turns and different user characteristics influence CHATWISE’s performance?

To address these questions, we evaluated CHATWISE on two complementary scenarios: we first conducted *offline* evaluations using data from a real clinical trial (Sec 4.1), then we performed *interactive* experiments, in which conversations are stochastically generated between CHATWISE and digital twins as simulated users (Sec 4.2). Below, we summarize the configurations, metrics, and main results of each experimental setting and defer more details to the Appendix.

4.1 Offline Evaluation on Real Clinical Data

4.1.1 Data Preprocessing. We extracted dialogues from the I-CONNECT clinic trials [13], with each in text format converted from video conversations. Short dialogues below 40 *turns* were excluded, where a *turn* represents an uninterrupted, continuous utterance by a single speaker. Two dialogue sets were subsampled and de-identified: (1) 150 randomly sampled dialogues to assess the overall strategy alignment, and (2) one dialogue per week from 7 randomly selected participants to assess CHATWISE’s alignment robustness to different participants and time periods over the clinical trial.

4.1.2 Offline Evaluation Metric. We define Strategy Match Percentage (SMP) to quantify the alignment between the conversational strategies generated by a chatbot and those generated by a human professional caregiver from the I-CONNECT clinic data. Let \mathbf{a}_t denote the set of caregiver-provided strategies annotated from a real conversation τ at turn (timestep) t , and \mathbf{a}'_t the strategies from a simulated moderator (CHATWISE or utterance generator), conditioned on the same real context s_t . Let $\mathbb{I}(a'_t \in \mathbf{a}_t)$ be an indicator function that returns 1 if a chatbot-generated strategy a'_t matches one of the human strategies \mathbf{a}_t , and 0 otherwise. The SMP is computed as following:

$$\text{SMP}_t(\mathbf{a}_t, \mathbf{a}'_t) = \frac{1}{|\mathbf{a}'_t|} \sum_{a'_t \in \mathbf{a}'_t} \mathbb{I}(a'_t \in \mathbf{a}_t).$$

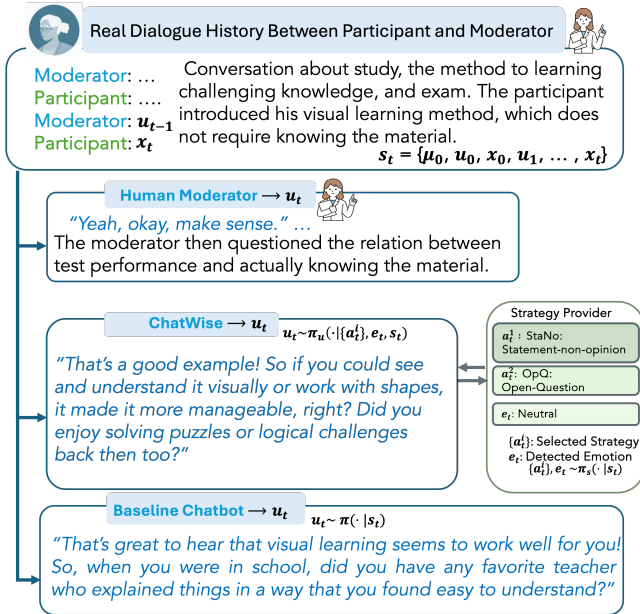


Figure 3: Comparison of moderator strategies: Given real dialogue contexts, responses from a baseline (generated), CHATWISE (generated), and the human moderator (real) are compared w.r.t. underlying strategies. A real sample is shown here while omitting participant data.

4.1.3 Strategy Label Annotation. Given a real dialogue denoting as $\tau = \{\mu_0, u_0, x_1, u_1, x_2, u_2, \dots\}$ with caregiver (moderator) utterance u_t and participant utterance x_t at turn t , we leveraged the macro-level strategy definition in Sec 3.1 to prompt an LLM as annotator (GPT-4o) and detect strategies \mathbf{a}_t behind u_t as the golden label. Denoting $s_t = \{\mu_0, u_0, x_1, \dots, x_t\}$ as the real conversation context, we feed s_t to a **baseline** chatbot to get utterance action \tilde{u}_t . The baseline is identical to the utterance generator π_u of CHATWISE yet does not receive strategy guidance from π_s . Similarly, we get strategy labels $\tilde{\mathbf{a}}_t$ behind \tilde{u}_t using the above annotator. We then input s_t to CHATWISE’s strategy provider π_s to get proposed strategy $\mathbf{a}'_t \sim \pi_s(s_t)$. We report the strategy alignment of $\text{SMP}_t(\mathbf{a}_t, \mathbf{a}'_t)$ (w/ CHATWISE) and $\text{SMP}_t(\mathbf{a}_t, \tilde{\mathbf{a}}_t)$ (w/o CHATWISE) over different turns t and average the result across dialogues. Figure 3 illustrates the process of collecting each moderator’s response based on real conversation history.

4.1.4 Overall Strategy Alignment. As shown in Figure 5, which averages results over 150 dialogues, CHATWISE consistently achieves an SMP close to 1.0, indicating a strong alignment with the strategies employed by real caregivers throughout the conversation. In contrast, the baseline chatbot deviates from human behavior when lacking action guidance. Both chatbots used GPT-4o as the utterance generator π_u , and o3-mini was used for π_s in CHATWISE. We do not employ micro-level utterance similarity, as the I-CONNECT interventions involve spontaneous daily conversations, rather than fixed question answering, unlike domains in programming or math. Whereas, on a macro level, CHATWISE mostly mirrors the strategic choices of human professionals who are well trained to follow clinical protocols.

To illustrate how our method behaves differently from the baseline, we collect the real dialogue context and different moderator responses where $\mathbf{a}_t = \mathbf{a}'_t \neq \tilde{\mathbf{a}}_t$, and show the discrepancy heatmap. Raw dialogue content is omitted to comply with the data usage agreement. Results in Figure 4 demonstrate that while “open question” is a commonly dominant strategy (forward DA) for both human and chatbot moderators, the primary strategy difference lies in how they lead into the question (e.g. “acknowledge”, “restatement”), which may in turn influence the specific question ultimately asked.

Robustness to Participant Heterogeneity: Figure 6 presents the SMP of CHATWISE across dialogue turns among participants. CHATWISE consistently outperforms the baseline and maintains notably higher SMP. These results demonstrate the robust alignment of CHATWISE with human caregiver strategies across diverse participants and their interaction contexts, which implies its potential for personalized dialogue systems.

Alignment Consistency Over Timeline: Figure 7 presents the SMP over time, where dialogues were sampled from different weeks over the clinical intervention. CHATWISE maintains consistently high SMP (near 1.0) across all timesteps, where the baseline shows lower and unstable performance. This indicates CHATWISE’s temporal robustness and its potential to facilitate the development of long-term cognitively supportive dialogues.




Figure 4: Strategy distribution for dialogue samples where $a_t = \tilde{a}_t \neq \hat{a}_t$. The vertical axis indicates ground-truth strategies a_t extracted from human professional utterances, and the horizontal axis denotes strategies \tilde{a}_t detected from the baseline chatbot.




Figure 5: SMP between human caregiver and chatbots on real conversation data, where w/o CHATWISE denotes the baseline without strategy guidance. Our method strongly aligns with human caregiver behavior.




Figure 6: SMP broken down by participants over 40 conversation turns, where each color indicates a specific person. Solid lines represent CHATWISE, and dashed lines denote the baseline. CHATWISE consistently aligns to human caregivers given different participant contexts.

4.2 Evaluation with Digital Twins

4.2.1 Conversation Generation. While evaluation on real clinic data shows strong alignment between CHATWISE and professional caregivers, the conversations are fixed, thus lacking indications




Figure 7: Strategy alignment of CHATWISE is consistent as real conversations progress over weeks.

whether and how CHATWISE may influence user cognitive or emotional status during conversations. To enable controlled and responsive development of CHATWISE prior to future human user studies, we used 9 digital twins provided by Hong et al. [12] as simulated users with different personas, which are fine-tuned LLMs calibrated to mimic the linguistic behaviors of older adults, some of whom show MCI symptoms. All digital twins are trained using *real, de-identified* dialogue data from the I-CONNECT clinical trial. We collected 20 trajectories of conversations between CHATWISE (moderator) and each digital twin as a user, with each conversation containing 20 turns. The same settings were applied to the baseline chatbot for fair comparisons.

4.2.2 Metrics for Interactive Engagement. User Verbosity: Given a conversation sequence τ that consists of user utterance x and moderator utterance u , we primarily measure user engagement through the user’s talkativeness compared to the moderator, *i.e.* the user’s verbosity (v), defined as $v = \frac{\sum_{x \in \tau} |x|}{\sum_{u \in \tau} |u|}$, where $|x|$ ($|u|$) denotes the number of tokens in an utterance x (u). This metric follows a clinical study [39] that suggests reducing the moderator talkativeness while encouraging participant (user) expression.

Cognitive Win Rate: See et al. [26] defined different aspects for evaluating conversation quality, from which we select 3 that are focused on assessing user cognitive status: *Listening*, *Fluency*, and *Making Sense*. All three metrics evaluate the **user** behavior rather than the overall dialogue quality considering both user and moderator behaviors.

Instead of providing a numerical score for each cognitive metric that might be unstable, we adopt the **Win Rate (WR)** definition from prior work [21] to compare pairs of dialogues generated with and without CHATWISE. Dialogue pairs are randomly matched so long as they are collected from the same digital twin (simulated user) for evaluation. WR is defined as the proportion of pairs in which the dialogue generated with CHATWISE is preferred over the baseline using an LLM-as-judge, based on a given cognitive metric.

To mitigate the LLM’s *position bias* issues observed in prior work [28, 41], we excluded sample pairs with inconsistent preference labels when the order of the two samples was reversed and calculated WRs using only the remaining consistent pairs.

4.2.3 Models and Baseline. We evaluated CHATWISE with 3 different LLM backbones as strategy providers: GPT-4o, o3-mini, and Llama3.1-405B. For each setting, a baseline adopts the same utterance generator as CHATWISE while without receiving strategy

| Strategy Provider | Verbosity \uparrow | Listening \uparrow | Fluency \uparrow | Making Sense \uparrow |
|-------------------|----------------------|----------------------|--------------------|-------------------------|
| baseline | 0.7398 | 0.5249 | 0.4986 | 0.5024 |
| GPT-4o | 0.8635 | 0.4962 | 0.4748 | 0.4786 |
| o3-mini | 0.8643 | 0.4884 | 0.5368 | 0.5180 |
| Llama 3.1-405B | 0.8083 | 0.4407 | 0.4926 | 0.4963 |

Table 1: Multifaceted evaluations on cognitive engagement with data sampled by interacting with digital twins over 20 turns, using different LLMs serving as strategy providers, where *Listening*, *Fluency*, and *Making Sense* metrics report win rates, and baseline denotes a conversation generator without receiving strategy guidance. All evaluations focus on user reactions.

guidance. The WR of *baseline* is reported as the complement of CHATWISE’s average WR: $WR_{\text{baseline}} = 1 - \text{average}(WR_{GPT-4o} + WR_{o3-mini} + WR_{Llama3.1-405B})$, where WR_X indicates the WR of CHATWISE, given model X as the strategy provider.

4.2.4 Performance Overview.

Effectiveness of Strategy-Guided Generation: As shown in Table 1, our design mostly demonstrated better user engagement compared to baselines across tested LLMs as the strategy provider. Particularly, the best-performing model in our setting, o3-mini, was optimized for STEM reasoning tasks, which indicates that strong reasoning ability can also benefit inferring dialogue strategies.

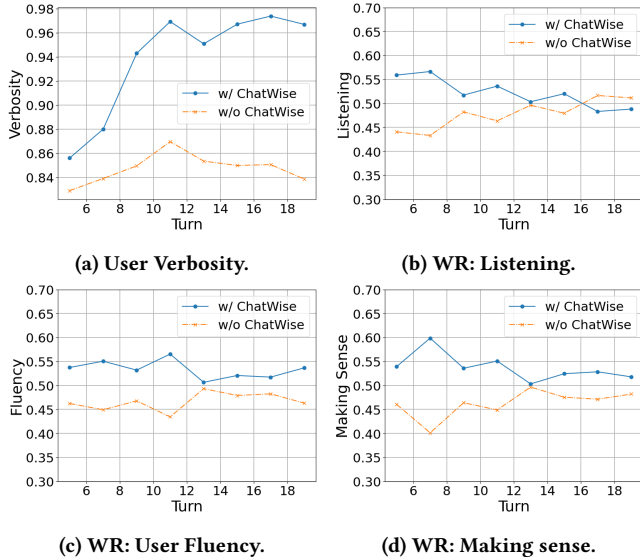


Figure 8: Multifaceted Engagement Evaluation across conversation turns, excluding the first two warm-up turns. CHATWISE’s performance gain over the baseline exists over multi-turn dialogues.

Performance with Dialogue Progression: We analyzed CHATWISE’s performance over increasing dialogue length by truncating conversations at turn t and computing metrics accordingly, as shown in Figure 8. Results revealed a clear upward trend in *verbosity*, which indicates increased talkativeness of the user rather than the CHATWISE, as opposed to the baseline. Meanwhile, the WR gains of

CHATWISE either remain stable or show a slight decline as conversations progress, which we infer are subject to the length bias of LLMs [7] as the overall number of tokens of CHATWISE-generated dialogues tends to decrease in later turns.

4.3 Ablation Study

We evaluated CHATWISE with and without user *emotion* in the strategy provider’s output. As shown in Table 2, where both strategy providers used GPT-4o as the backbone, removing emotion information led to only a slight drop of performance, while removing CHATWISE entirely resulted in a much lower verbosity score, which highlights that strategy guidance itself is the key driver of user engagement.

| Method | CHATWISE | CHATWISE w/o emotion | w/o CHATWISE |
|-----------|----------|----------------------|--------------|
| Verbosity | 0.8635 | 0.8463 | 0.7398 |

Table 2: Ablative study on the performance of user verbosity when removing different contextualized information. w/o emotion denotes the strategy provider without user emotion information as input.

4.4 CHATWISE’s Robustness to User Persona

As shown in Figure 9, the gain of CHATWISE interacting with 9 simulated users is consistently significant, in which the metrics were log-normalized. This indicates the robustness of CHATWISE given varying senior characteristics.




Figure 9: User verbosity log-normalized across 9 digital twins. Given varying strategy provider backbones, our design consistently outperforms the baseline (w/o CHATWISE) with non-strategic conversations.

4.5 Analysis of User Emotion Transitions

To analyze if CHATWISE can support positive emotional shifts during conversations, we selected four digital twins and collected 40 additional dialogues from each under both CHATWISE and the baseline settings for further analysis, with the main findings summarized below.

Transient user emotions: We define an emotion transition triplet as (e_t, u_t, e_{t+1}) , where e_t (e_{t+1}) represents the user’s emotion at turn t ($t+1$), and u_t is the moderator utterance in between. We computed the average occurrence of emotion triplets and showed the top 15 most frequent in Figure 10. Most triplets reflect unchanged user emotions, which indicates difficulty in either influencing or detecting emotional shifts within a short turn.






Figure 10: Average occurrence of each emotion transition triplet across the samples of each digital twin.

User Emotions over Multi-Turn Conversations: We then analyzed user emotion changes from the beginning to the end of dialogues and calculated averaged occurrence across digital twins (Figure 11b). Over 48% users experienced emotional shifts post-dialogue, with significantly more positive changes after engaging with CHATWISE. This implies that while transient emotions are difficult to track, strategic conversational support may improve user emotions over multi-turn conversations.



(a) *Real user emotion transitions summarized from 150 dialogues in the I-CONNECT study.*



(b) Averaged *digital twin* emotion transitions among 40 dialogues, summarized over 320 dialogues.

Figure 11: User emotion shifts in conversations between digital twin and CHATWISE (Figure 11b) align well with real-user emotion shifts in the clinical study (Figure 11a). More *neutral-to-joy* transitions were observed given CHATWISE compared to the baseline.

Predominant Strategies: We identified the top 10 most frequent strategies across digital twins (Figure 12) and found that predominant interaction strategies remained consistent across user types (See Appendix G). Particularly, *Open Question*, *Statement-non-opinion*, and *Acknowledgment* strategies dominate CHATWISE driven conversations.

5 Conclusion and Future Work

We present a dual-level framework for AI chatbots that supports multi-turn conversations for older adults by integrating clinical insights into LLM reasoning. Extensive evaluations based on both real clinical data and generated conversations showed that our method aligns well with the behavior of professional human caregivers while robustly enhancing the cognitive status of older adults through simulation studies. Our future work includes training-based methods for further optimizing multi-turn interaction experience and investigating the acceptability and feasibility of our design through real user studies.

6 Ethical Considerations

This study involved simulated dialogues using digital twins trained on *de-identified* conversation data from a clinical trial involving older adults. Offline analysis of this dataset was conducted to evaluate the similarity between chatbot and caregiver strategies. This project did not use any personally identifiable information, and all data processing followed institutional privacy and research ethics guidelines. Potential risks are controllable, which include privacy concerns of emotion detection and misinterpretation of conversation due to AI-based evaluation. To mitigate this, the emotional detection module in the system is made optional and can be omitted as configured. Win rate calculations are based on a sanity check where we omitted pair samples that receive inconsistent preference labels by LLM-as-the-judge after swapping their positions. We have manually cross-validated a sufficient subset of AI-generated analyses to reduce bias. A real-user study is planned for future validation and will undergo full Institutional Review Board (IRB) approval.

References

- [1] 2024. Senior Talk: AI Chatbot for Mental Health Support for Seniors. *Senior Talk*. <https://www.senior-talk.com/senior-mental-health>
- [2] 2024. Your AI sidekick for happier, healthier aging. *Elliq*. <https://elliq.com/>
- [3] Marwa Abdulhai, Isadora White, Charlie Snell, Charles Sun, Joey Hong, Yuexiang Zhai, Kelvin Xu, and Sergey Levine. 2023. Lmrl gym: Benchmarks for multi-turn reinforcement learning with language models. *arXiv preprint arXiv:2311.18232* (2023).
- [4] James Allen and Mark Core. 1997. Draft of DAMSL: Dialog act markup in several layers.
- [5] Shuyi Cao, Lizhen Qu, and Leimin Tian. 2021. Causal relationships between emotions and dialog acts. In *2021 9th International Conference on Affective Computing and Intelligent Interaction (ACII)*. IEEE, 1–8.
- [6] Sumit Kumar Dam, Choong Seon Hong, Yu Qiao, and Chaoning Zhang. 2024. A complete survey on llm-based ai chatbots. *arXiv preprint arXiv:2406.16937* (2024).
- [7] Yann Dubois, Balázs Galambosi, Percy Liang, and Tatsunori B Hashimoto. 2024. Length-controlled alpacaEval: A simple way to debias automatic evaluators. *arXiv preprint arXiv:2404.04475* (2024).
- [8] Jesus Favela, Dagoberto Cruz-Sandoval, and Mario O. Parra. 2023. Conversational Agents for Dementia using Large Language Models. In *2023 Mexican International Conference on Computer Science (ENC)*. 1–7. doi:10.1109/ENC60556.2023.10508610
- [9] Katherine L Fiori and Justin Jager. 2012. The impact of social support networks on mental and physical health in the transition to older adulthood: A longitudinal, pattern-centered approach. *International Journal of Behavioral Development* 36, 2 (2012), 117–129.

- [10] Zhaolin Gao, Wenhao Zhan, Jonathan Daniel Chang, Gokul Swamy, Kianté Brantley, Jason D. Lee, and Wen Sun. 2025. Regressing the Relative Future: Efficient Policy Optimization for Multi-turn RLHF. In *The Thirteenth International Conference on Learning Representations*. <https://openreview.net/forum?id=cVyELMpMRS>
- [11] Clara E Hill. 2020. *Helping skills: Facilitating exploration, insight, and action*. American Psychological Association.
- [12] Junyuan Hong, Wenqing Zheng, Han Meng, Siqi Liang, Anqing Chen, Hiroko H Dodge, Jiayu Zhou, and Zhangyang Wang. 2024. A-CONNECT: Designing AI-based Conversational Chatbot for Early Dementia Intervention. In *ICLR 2024 Workshop on Large Language Model (LLM) Agents*.
- [13] I-CONNECT. 2024. Study Details. <https://www.i-connect.org/study-details>
- [14] Jiyue Jiang, Sheng Wang, Qintong Li, Lingpeng Kong, and Chuan Wu. 2023. A Cognitive Stimulation Dialogue System with Multi-source Knowledge Fusion for Elders with Cognitive Impairment. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (Eds.). Association for Computational Linguistics, Toronto, Canada, 10628–10640. doi:10.18653/v1/2023.acl-long.593
- [15] 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
- [16] Gill Livingston, Jonathan Huntley, Andrew Sommerlad, David Ames, Clive Ballard, Sube Banerjee, Carol Brayne, Alistair Burns, Jiska Cohen-Mansfield, Claudia Cooper, et al. 2020. Dementia prevention, intervention, and care: 2020 report of the Lancet Commission. *The lancet* 396, 10248 (2020), 413–446.
- [17] Ashley ND Meyer, Traber D Giardina, Christiane Spitzmueller, Umber Shahid, Taylor MT Scott, and Hardeep Singh. 2020. Patient perspectives on the usefulness of an artificial intelligence-assisted symptom checker: cross-sectional survey study. *Journal of medical Internet research* 22, 1 (2020), e14679.
- [18] Nicholas R Nicholson. 2012. A review of social isolation: an important but underassessed condition in older adults. *The journal of primary prevention* 33 (2012), 137–152.
- [19] World Health Organization. 2024. Ageing and health. <https://www.who.int/news-room/fact-sheets/detail/ageing-and-health>
- [20] Risako Owan, Libby Ferland, Naome Etori, Wilma Koutstaal, and Maria Gini. 2023. Conversational agents for elderly users. In *Workshop Assistive Robotics for Citizens at IROS*.
- [21] Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2023. Direct Preference Optimization: Your Language Model is Secretly a Reward Model. In *Advances in Neural Information Processing Systems*, A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (Eds.), Vol. 36. Curran Associates, Inc., 53728–53741. https://proceedings.neurips.cc/paper_files/paper/2023/file/a85b405ed65c6477a4fe8302b5e06ce7-Paper-Conference.pdf
- [22] Yi Ren, Yangjun Ruan, Xu Tan, Tao Qin, Sheng Zhao, Zhou Zhao, and Tie-Yan Liu. 2019. FastSpeech: Fast, robust and controllable text to speech. *Advances in neural information processing systems* 32 (2019).
- [23] Hyeyoung Ryu, Soyeon Kim, Dain Kim, Soan Han, Keeheon Lee, and Younah Kang. 2020. Simple and steady interactions win the healthy mentality: designing a chatbot service for the elderly. *Proceedings of the ACM on human-computer interaction* 4, CSCW2 (2020), 1–25.
- [24] Timothy A Salthouse. 2009. When does age-related cognitive decline begin? *Neurobiology of aging* 30, 4 (2009), 507–514.
- [25] John R Searle. 1976. A classification of illocutionary acts1. *Language in society* 5, 1 (1976), 1–23.
- [26] Abigail See, Stephen Roller, Douwe Kiela, and Jason Weston. 2019. What makes a good conversation? How controllable attributes affect human judgments. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, Jill Burstein, Christy Doran, and Thamar Solorio (Eds.). Association for Computational Linguistics, Minneapolis, Minnesota, 1702–1723. doi:10.18653/v1/N19-1170
- [27] Woosuk Seo, Ayse G Buyuktur, Sanya Verma, Hyeryoung Kim, Sung Won Choi, Laura Sedig, and Sun Young Park. 2021. Learning from healthcare providers' strategies: Designing technology to support effective child patient-provider communication. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [28] Lin Shi, Chiyu Ma, Wenhua Liang, Weicheng Ma, and Soroush Vosoughi. 2024. Judging the judges: A systematic investigation of position bias in pairwise comparative assessments by llms. *arXiv preprint arXiv:2406.07791* (2024).
- [29] Qiao Sun, Jiexin Xie, Nanyang Ye, Qinying Gu, and Shijie Guo. 2025. Enhancing Nursing and Elderly Care with Large Language Models: An AI-Driven Framework. *Proceedings of the 31st International Conference on Computational Linguistics* (2025).
- [30] Maite Taboada. 2006. Spontaneous and non-spontaneous turn-taking. *Pragmatics* (July 2006), 329–360. doi:10.1075/prag.16.2-3.04tab
- [31] Rui Hong Teo, Wen Hui Cheng, Ling Jie Cheng, Ying Lau, and Siew Tiang Lau. 2023. Global prevalence of social isolation among community-dwelling older adults: a systematic review and meta-analysis. *Archives of gerontology and geriatrics* 107 (2023), 104904.
- [32] Chun-Hua Tsai, Yue You, Xinning Gui, Yubo Kou, and John M Carroll. 2021. Exploring and promoting diagnostic transparency and explainability in online symptom checkers. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–17.
- [33] Aaron Van Den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, Koray Kavukcuoglu, et al. 2016. Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499* 12 (2016).
- [34] Siddharth Verma, Justin Fu, Mengjiao Yang, and Sergey Levine. 2022. Chai: A chatbot ai for task-oriented dialogue with offline reinforcement learning. *arXiv preprint arXiv:2204.08426* (2022).
- [35] Anna Xyngkou, Chee Siang Ang, Panote Siriara, Jonasz Piotr Kopecki, Alexandra Covaci, Eiman Kanjo, and Wan-Jou She. 2024. MindTalker: Navigating the Complexities of AI-Enhanced Social Engagement for People with Early-Stage Dementia. 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 96, 15 pages. doi:10.1145/3613904.3642538
- [36] Ziqi Yang, Xuhai Xu, Bingsheng Yao, Ethan Rogers, Shao Zhang, Stephen Intille, Nawar Shara, Guodong Gordon Gao, and Dakuo Wang. 2024. Talk2Care: An LLM-based Voice Assistant for Communication between Healthcare Providers and Older Adults. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 8, 2 (2024), 1–35.
- [37] Yue You, Chun-Hua Tsai, Yao Li, Fenglong Ma, Christopher Heron, and Xinning Gui. 2023. Beyond self-diagnosis: how a chatbot-based symptom checker should respond. *ACM Transactions on Computer-Human Interaction* 30, 4 (2023), 1–44.
- [38] Kexin Yu, Benjamin M Hampstead, Katherine Wild, Lisa C Silbert, and Hiroko H Dodge. 2022. Examining I-CONNECT Intervention Effect on Psychosocial Wellbeing. *Alzheimer's & Dementia* 18 (2022), e059654.
- [39] Kexin Yu, Katherine Wild, Kathleen Potempa, Benjamin M Hampstead, Peter A Lichtenberg, Laura M Struble, Patrick Pruitt, Elena L Alfaro, Jacob Lindsley, Mattie MacDonald, et al. 2021. The Internet-based conversational engagement clinical trial (I-CONNECT) in socially isolated adults 75+ years old: randomized controlled trial protocol and COVID-19 related study modifications. *Frontiers in digital health* 3 (2021), 714813.
- [40] Kexin Yu, Chao-Yi Wu, Lisa C Silbert, Jeffrey A Kaye, and Hiroko H Dodge. 2023. I-CONNECT Intervention Effects on Weekly Time Spent Outside of Home and Social Contacts among Socially Isolated Older Adults with Normal Cognition and Mild Cognitive Impairment. *Alzheimer's & Dementia* 19 (2023), e077984.
- [41] Yijiong Yu, Huiqiang Jiang, Xufang Luo, Qianhui Wu, Chin-Yew Lin, Dongsheng Li, Yuqing Yang, Yongfeng Huang, and Lili Qiu. 2024. Mitigate position bias in large language models via scaling a single dimension. *arXiv preprint arXiv:2406.02536* (2024).
- [42] Fengpei Yuan, Wenjun Zhou, Hiroko H. Dodge, and Xiaopeng Zhao. 2023. Short: Causal structural learning of conversational engagement for socially isolated older adults. *Smart Health* 28 (2023), 100384. doi:10.1016/j.smhl.2023.100384
- [43] Justine Zhang and Cristian Danescu-Niculescu-Mizil. 2020. Balancing objectives in counseling conversations: Advancing forwards or looking backwards. *arXiv preprint arXiv:2005.04245* (2020).
- [44] Chujie Zheng, Sahand Sabour, Jiaxin Wen, Zheng Zhang, and Minlie Huang. 2023. AugESC: Dialogue Augmentation with Large Language Models for Emotional Support Conversation. In *Findings of the Association for Computational Linguistics: ACL 2023*, Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (Eds.). Association for Computational Linguistics, Toronto, Canada, 1552–1568. doi:10.18653/v1/2023.findings-acl.99
- [45] Yifei Zhou, Andrea Zanette, Jiayi Pan, Sergey Levine, and Aviral Kumar. 2024. Archer: Training language model agents via hierarchical multi-turn rl. *arXiv preprint arXiv:2402.19446* (2024).

A Prompt design

The prompt used in the experiments and structured output design are available at <https://anonymous.4open.science/r/ChatWise-8F53>, including:

- Strategy provider system prompt.
- Moderator initial system prompt.
- Moderator system prompt with strategies.
- Strategy provider system prompt for ablation study.
- Moderator system prompt with strategies for ablation study.
- Structured output class for OpenAI models as strategy provider.
- Structured output class for OpenAI models as strategy provider in ablation study.
- System prompt for GPT-4o to extract the strategy given by Llama3.1.

B Offline Data Preparation

We extracted dialogue content from the I-CONNECT video conversations and processed it into OpenAI-compatible dialogue history format. Dialogues with fewer than 40 turns were excluded. We then subsampled two subsets for evaluation purposes. To assess the overall strategy alignment of ChatWise, we randomly sampled 150 dialogues. To evaluate the robustness of ChatWise across different participants and time periods, we randomly selected 7 participants who completed the full study and sampled one dialogue from each of their sessions during weeks 1, 9, 17, 25, 33, and 41. All user identifiers and personally identifiable information have been anonymized. The OpenAI-compatible dialogue history format example: [{"role": "system", "content": "system_prompt"}, {"role": "user", "content": "user_content1"}, {"role": "assistant", "content": "assistant_content1"}, {"role": "user", "content": "user_content2"}, {"role": "assistant", "content": "assistant_content2"}, ...].

C Data Generation Configuration

By default, we used o3-mini as strategy provider. GPT-4o serves as an utterance generator in all settings. W/o CHATWISE denotes the baseline, which is using GPT-4o as the utterance generator only, without a strategy provider. We tested different LLM backbones as strategy providers in CHATWISE. Considering the Llama3.1-405B does not support structured output, we applied GPT-4o as the strategy extractor to structure the strategy output of Llama3.1-405B. We employed GPT-4o as the judge to select the preferred one when computing the Win Rate. The prompt for selecting the preferred response is listed in Figure 13. The following are the configurations of each model:

Utterance generator (GPT-4o):

n=1
max_tokens=1024
top_p=1
temperature=1

GPT-4o, o3-mini as strategy provider:

n=1
max_tokens=1024
top_p=1
temperature=1
response_format=Strategy

Llama3.1-405B as strategy provider:

top_p: 0.9
max_tokens: 1024
temperature: 0.6
presence_penalty: 0
frequency_penalty: 0
GPT-4o as strategy extractor:
n=1
max_tokens=1024
top_p=1
temperature=1
GPT-4o as judge:
n=1
max_tokens=1024
top_p=1
temperature=1

D Win rate for w/o CHATWISE

The Win rate for w/o CHATWISE is defined as:

$$1 - \text{average}(\text{WR}_{\text{GPT-4o}} + \text{WR}_{\text{o3-mini}} + \text{WR}_{\text{Llama3.1-405B}})$$

,where WR_X is CHATWISE's Win rate against the baseline with X as strategy provider.

E Dialogue Acts

The map of strategy to its corresponding abbreviated tag is listed in Table3.

There are two kinds of strategies: backward-looking and forward-looking. Backward-looking strategies reflect how the current utterance relates to the previous discourse. Forward-looking strategies reflect the current utterance constrains the future beliefs and actions of the participants and affects the discourse. Table 4 and Table 5 provide definitions and examples of each.

F Log-normalization

The following is the log-normalization function, where y is the normalized result, x is the input variable.

$$y = \ln(4x + 1)$$

G Primary Strategies

We calculated the average occurrence of each strategy across each digital twin and listed their top 10 most frequently occurring strategies, as shown in Figure 12.




Figure 12: Strategy occurrence across digital twins.

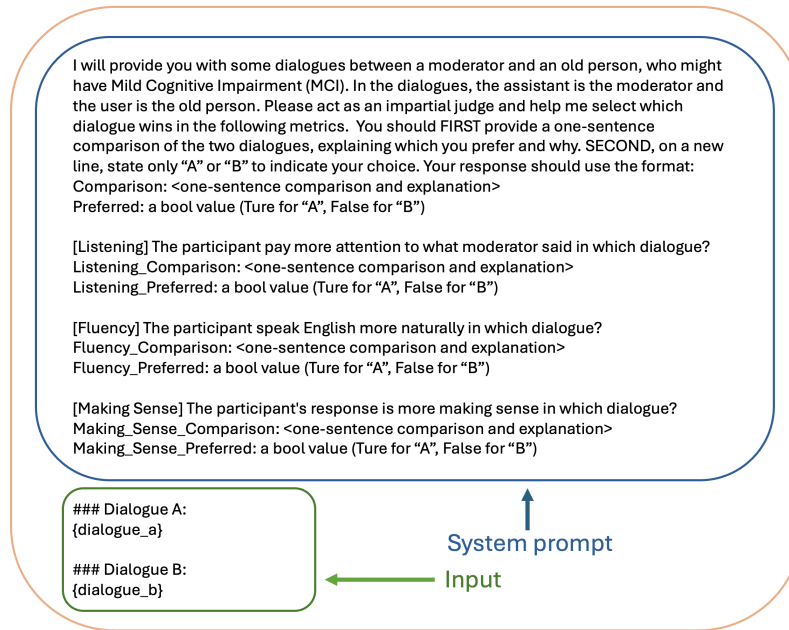


Figure 13: Prompt for GPT-4o as a Judge.

| Strategy | Tag |
|-----------------------------|-------|
| Acknowledge (Backchannel) | Ack |
| Statement-non-opinion | StaNo |
| Statement-opinion | Sta |
| Affirmation and Reassurance | Agr |
| Appreciation | App |
| Conventional-closing | ConC |
| Hedge | H |
| Other | Oth |
| Quotation | Quo |
| Action-directive | AcD |
| Collaborative Completion | CoC |
| Restatement or Paraphrasing | Rep |
| Offers Options Commits | Off |
| Self-talk | Sel |
| Apology | Apo |
| Reflection of Feelings | RoF |
| Yes-No-Question | YNQ |
| Wh-Question | WhQ |
| Declarative Yes-No-Question | DYNQ |
| Open-Question | OpQ |
| Or-Clause | OrC |
| Conventional-opening | CoO |
| Self-disclosure | Sd |
| Providing Suggestions | PS |
| Information | I |

Table 3: Strategy to its corresponding tag. The strategies are drawn from the DAs in DAMSL [4] that are used by telehealth clinical trials [42], integrated strategies from prior emotional support dataset [15] .

| Strategy | Definiton | Example |
|----------|--|--|
| StaNo | A factual statement or descriptive utterance that does not include an opinion. | Me, I'm in the legal department. |
| Ack | A brief utterance that signals understanding, agreement, or active listening. | Uh-huh. |
| Sta | A statement that conveys a personal belief, judgment, or opinion. | I think it's great |
| Agr | Affirm the help seeker's strengths, motivation, and capabilities and provide reassurance and encouragement. | That's exactly it. |
| App | An expression of gratitude, admiration, or acknowledgment of another's effort or input. | I can imagine. |
| ConC | A formal or socially standard utterance signaling the end of a conversation. | Well, it's been nice talking to you. |
| H | An expression that introduces uncertainty or qualification to a statement, often to soften its impact. | I don't know if I'm making any sense or not. |
| Oth | Exchange pleasantries and use other support strategies that do not fall into the above categories. | Well give me a break, you know. |
| Quo | A direct or indirect repetition of someone else's words. | Albert Einstein once said, "Imagination is more important than knowledge." |
| AcD | A command, request, or suggestion directing someone to take action. | Why don't you go first |
| CoC | A continuation or completion of someone else's utterance in a collaborative manner. | If we want to make it to the top of the mountain before sunset, we should... |
| Rep | A simple, more concise rephrasing of the help-seeker's statements that could help them see their situation more clearly. | It sounds like you're saying that you're struggling to stay on top of your work, and it's leaving you feeling overwhelmed. |
| Off | A statement proposing choices, making a commitment, or offering to do something. | I'll have to check that out |
| Sel | An utterance directed at oneself, often reflecting internal thought processes or problem-solving. | What's the word I'm looking for |
| Apo | An expression of regret or asking for forgiveness. | I'm sorry. |
| RoF | Articulate and describe the help-seeker's feelings. | It sounds like you're feeling really frustrated and drained because your efforts don't seem to be paying off. |


Table 4: Backward-looking strategies, definition, and example.

| Strategy | Definiton | Example |
|----------|--|---|
| YNQ | A question expecting a binary (yes/no) response. | Do you have to have any special training? |
| WhQ | A question beginning with a wh-word (e.g., what, who, where), seeking specific information. | Well, how old are you? |
| DYNQ | A statement posed as a question, expecting a yes/no answer. | So you can afford to get a house? |
| OpQ | A broad question inviting a wide range of responses, often conversational. | How about you? |
| OrC | A question offering explicit alternatives, often in the form of "or." | or is it more of a company? |
| CoO | A socially standard utterance used to initiate a conversation. | How are you? |
| Sd | Divulge similar experiences that you have had or emotions that you share with the help-seeker to express your empathy. | I completely understand how you feel. I remember feeling the same way before my first big presentation at work. I was so anxious, but I found that practicing a few extra times really helped calm my nerves. |
| PS | Provide suggestions about how to change, but be careful to not overstep and tell them what to do. | You can keep a note to stop your idea from going. |
| I | Provide useful information to the help-seeker, for example with data, facts, opinions, resources, or by answering questions. | Taking silver line from Washington D.C. to Dulles Intel Airport costs about 1 hour. |


Table 5: Forward-looking strategies, definition, and example.

H Personalized Dialogue Analysis


This section shows the personalized dialogue analysis. The Strategy occurrence across digital twins is shown in Figure 14. The Occurrence of each emotion transition triplet across digital twins is shown in Figure 15. The *Open Question*, *Statement-non-opinion*,




(a) Digital Twin 3.



(b) Digital Twin 5.




(c) Digital Twin 6.




(d) Digital Twin 9.

Figure 14: Strategy occurrence across digital twins.


and *Acknowledgment* strategies still dominate CHATWISE driven conversations, suggesting their potential effectiveness in fostering engagement in conversations. The user's emotion is detected as unchanged in most triplets, indicating the difficulty of altering or measuring the user's emotional movement within a short turn.




(a) Digital Twin 3.



(b) Digital Twin 5.



(c) Digital Twin 6.



(d) Digital Twin 9.

Figure 15: Occurrence of each emotion transition triplet across digital twins.