From Answers to Relationships: Dual Digital Twins for Well-Being-Oriented Face-to-Face Service

Masayuki Numao[1] and Keiki Takadama[2]
[1]Faculty of Information Design, Tokyo Information Design Professional University
[2]Information Technology Center, The University of Tokyo

[1]numao@tid.ac.jp, [2]takadama@g.ecc.u-tokyo.ac.jp

Abstract

Advances in artificial intelligence, particularly large language models, have accelerated the development of answer-oriented systems optimized for correctness, efficiency, and task completion. However, many face-to-face human services—such as caregiving, rehabilitation, and education—operate within relationship-oriented contexts in which value emerges through engagement, emotional alignment, and motivational support rather than informational accuracy alone. This paper introduces the Dual Digital Twin (DDT) framework as a computational architecture for modeling and analyzing such relational interactions. The framework simulates both participants in a service encounter through a Resident Twin and Staff Twin, while an Observer Twin evaluates relational dynamics including engagement persistence, emotional resonance, and motivational impact. Implemented as a multi-LLM agent system on the Dify platform, the framework enables controlled scenario simulation and reflective training. We conduct a rehabilitation motivation study centered on daily walking goals monitored via wearable devices, comparing task-oriented, relationship-oriented, and hybrid dialogue strategies. Simulation results indicate that while task-oriented feedback ensures performance clarity, relational continuity is essential for sustaining engagement following negative outcomes. Hybrid strategies integrating performance transparency with motivational scaffolding yield the most balanced results. These findings position AI not as a replacement for relational labor, but as a reflective partner that illuminates the hidden mechanisms sustaining meaningful human interaction and well-being.

Introduction

AI and the Question of Human Creativity

Recent advances in artificial intelligence, particularly large language models (LLMs), have dramatically expanded the scope of machine capabilities in language understanding, knowledge synthesis, and problem solving. Contemporary AI systems now perform a wide range of answer-oriented tasks, from technical consultation to medical triage and educational tutoring. These developments have intensified long-standing debates concerning the future role of AI in human society — especially whether AI will augment human creativity or render aspects of human labor obsolete.

While much attention has focused on knowledge work and analytical professions, comparatively less discussion has addressed face-to-face human services such as caregiving, rehabilitation, and education. These domains involve forms of interaction in which value is not determined solely by correctness or efficiency, but by relational continuity, emotional alignment, and motivational impact.

The goal of this research is to investigate how relational interaction processes can be computationally modeled and evaluated through the Dual Digital Twin framework.

Limits of Answer-Oriented AI in Human Services

Most contemporary AI dialogue systems are designed within answer-oriented paradigms. Task-oriented dialogue (ToD) architectures optimize structured information exchange, while LLM-based conversational agents aim to produce contextually appropriate responses. In both cases, system performance is evaluated primarily through correctness, task completion, and response relevance.

However, interactions in caregiving and rehabilitation contexts frequently lack singular correct answers. For example, supporting an elderly individual in maintaining rehabilitation motivation involves emotional reassurance, conversational persistence, and adaptive encouragement — processes that cannot be reduced to informational accuracy.

Empirical observations in elderly conversational care further suggest that sustained dialogue itself contributes to well-being. Even when informational content is minimal, the presence of an attentive conversational partner enhances emotional stability and satisfaction. These findings point to the existence of hidden relational tasks embedded within face-to-face services.

From Explicit Tasks to Hidden Relational Goals

This distinction motivates a conceptual shift from explicit task execution to relational interaction processes. In rehabilitation dialogue, explicit goals may involve performance monitoring (e.g., daily walking targets), while implicit goals include sustaining motivation and reinforcing self-efficacy.

Importantly, these two dimensions are not mutually exclusive. Rather, relational continuity often functions as the condition that enables task adherence. Corrective feedback delivered without relational scaffolding may reduce engagement, whereas motivational framing can sustain participation even when performance falls short.

Understanding this dual structure requires new computational frameworks capable of modeling both informational and relational dynamics.

Scenario-Based Evaluation in Rehabilitation Motivation

To examine the framework's analytical potential, we conducted simulation studies in a rehabilitation motivation scenario centered on daily walking goals monitored via wearable devices.

Three dialogue strategies were compared:

- Task-Oriented feedback.
- Relationship-Oriented encouragement.
- Hybrid interaction integrating both.

Observer-based evaluation revealed that while task-oriented dialogue ensured performance clarity, relational continuity was essential for sustaining engagement following negative outcomes. Hybrid strategies produced the most balanced motivational effects.

These findings operationalize relational continuity as a measurable interaction outcome.

Contributions of This Paper

This paper makes the following contributions:

1. Conceptual distinction between answer-oriented AI and relationship-oriented service interaction.
2. Introduction of the Dual Digital Twin framework for relational simulation.
3. Implementation of a multi-agent architecture on a large language model platform.
4. Scenario-based evaluation in rehabilitation motivation dialogue.
5. Empirical demonstration of hybrid task–relational interaction benefits.

Conceptual Background

Answer-Oriented AI: Task Completion as Success

Task-oriented dialogue (ToD) systems represent a canonical paradigm of answer-oriented artificial intelligence. Their primary objective is to accomplish predefined tasks by acquiring necessary information and producing appropriate responses within structured conversational flows. Classical implementations rely on slot-filling architectures in which dialogue progresses through the acquisition of required parameters, enabling task completion such as reservations, scheduling, or health monitoring.

In elderly care contexts, such task structures often include cognitive screening and daily monitoring activities. For example, structured diagnostic dialogues such as HDS-R–based questioning can be implemented as scripted scenarios in which the system sequentially gathers responses related to memory, orientation, and reasoning ability.

Recent research has sought to enhance ToD naturalness by incorporating neural conversational capabilities or chit-chat modules. However, these additions typically function to support task execution rather than redefine its evaluative criteria. Thus, ToD systems remain fundamentally grounded in answer-oriented optimization.

Relationship-Oriented Practices: Hidden Tasks and Relational Continuity

In contrast, non-task-oriented dialogue — often referred to as chit-chat — has traditionally been defined as conversation conducted for its own sake, without explicit task completion goals. However, when applied to elderly care, this classification becomes insufficient.

Kawamura et al. propose that seemingly non-task-oriented dialogue in care contexts actually contains *hidden tasks* embedded within relational interaction. These hidden tasks are not framed as explicit problem-solving objectives but as relational goals such as activating the user, reducing boredom, fostering comfort, or encouraging long-term engagement.

In this view, casual conversation serves functional purposes that remain implicit within the interaction design. For example, robot-initiated dialogue may introduce light topics or simple questions not to obtain factual information but to sustain conversational flow and stimulate cognitive engagement. Hidden tasks such as "getting close," "relaxing the user," or "activating participation" guide dialogue progression even in the absence of formal task structures.

We therefore conceptualize these practices as *relationship-oriented dialogue systems*. Their evaluative criteria differ fundamentally from ToD systems:

- Success is measured by engagement duration.
- Emotional resonance outweighs factual accuracy.
- Conversational continuity supersedes task completion.

In this sense, hidden tasks operationalize relational objectives within conversational AI. They provide a structural bridge between explicit task systems and human-centered interaction practices, enabling dialogue systems to support long-term care engagement and minimally invasive cognitive monitoring simultaneously.

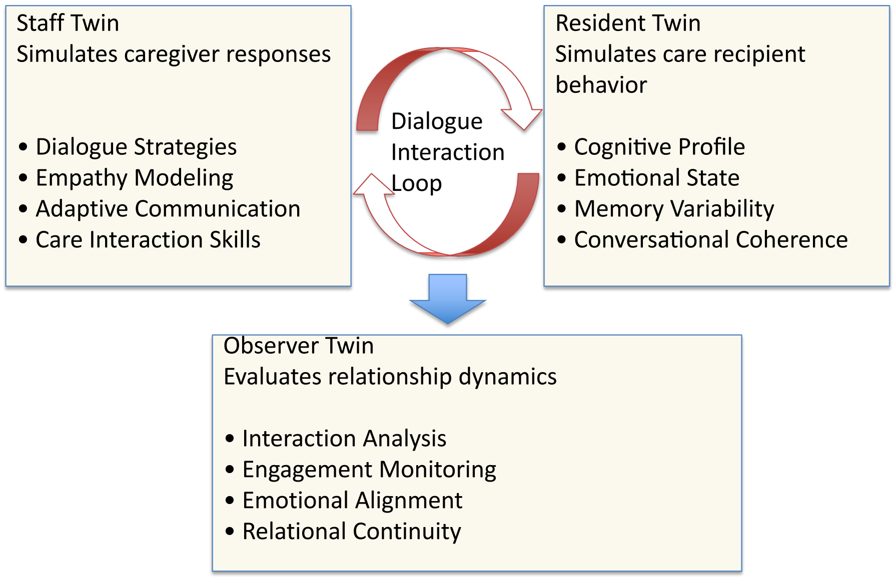

Figure 1. Dual Digital Twin Framework

Bridging Tasks and Relationships via Dual Digital Twins

The preceding sections distinguished between explicit task-oriented dialogue systems and relationship-oriented conversational practices characterized by hidden relational tasks. While these paradigms are often treated as separate design categories, real-world care and educational settings rarely conform strictly to one or the other. Instead, they exhibit hybrid characteristics in which explicit informational objectives coexist with implicit relational goals.

Conventional AI architectures, however, are typically optimized for one dimension. Task-oriented dialogue systems prioritize correctness and completion, while open-domain conversational systems emphasize fluency and naturalness without systematic relational evaluation. Neither paradigm alone adequately captures the dual nature of human-centered service interactions, where informational exchange and relational continuity operate simultaneously. The Dual Digital Twin (DDT) framework is proposed as a structural bridge between these dimensions. Rather than classifying dialogue as either task-oriented or non-task-oriented, the framework models interaction as a dynamic system involving both explicit and implicit objectives. The Resident Twin and Staff Twin simulate bidirectional interaction in which utterances carry both informational content and relational signals. The Observer Twin, in turn, evaluates the interaction not solely according to correctness, but across relational metrics such as engagement, emotional alignment, and conversational persistence.

Importantly, the DDT framework does not reject task-oriented AI. Instead, it embeds explicit tasks within a broader relational evaluative structure. In this sense, it provides a layered model:

- Layer-1: Informational Task Performance
- Layer-2: Relational Continuity and Well-Being Impact

The Observer Twin functions as the integrative mechanism that connects these layers. By generating feedback on relational metrics alongside task outcomes, the system supports reflective learning for caregivers and educators, allowing them to understand not only whether an interaction was informationally correct, but whether it strengthened the human connection.

Related Work

Research relevant to the proposed Dual Digital Twin framework spans four major domains: (1) answer-oriented dialogue systems and large language models, (2) socially assistive and care-oriented conversational systems, (3) human-centered and well-being–oriented AI, and (4) digital twin modeling. This section situates the present work within these intersecting research trajectories.

Answer-Oriented Dialogue Systems and Large Language Models

Task-oriented dialogue (ToD) systems have historically focused on structured information exchange, emphasizing slot-filling, belief tracking, and policy optimization. Early work applied reinforcement learning to optimize dialogue strategies in constrained domains such as in-car assistance and information services (Lemon et al., 2006). Subsequent research advanced statistical dialogue state tracking and partially observable Markov decision process (POMDP) frameworks to manage uncertainty in user intent (Young et al., 2013; Wang and Lemon, 2013).

Recent neural approaches have integrated generative encoder–decoder architectures and large language models (LLMs) to enhance flexibility and contextual understanding (Zhao et al., 2017; Brown et al., 2020). Hybrid dialogue systems incorporating chit-chat capabilities alongside task execution have also emerged to mitigate interaction rigidity (Sun et al., 2020; Yu et al., 2022]. Despite these advances, system evaluation remains predominantly answer-oriented, prioritizing correctness, task completion, and response relevance rather than relational continuity.

Socially Assistive and Care-Oriented Conversational Systems

Research in socially assistive robotics (SAR) has explored relational interaction in therapeutic and caregiving contexts. Early definitions of SAR emphasized assistance through social rather than physical intervention (Feil-Seifer and Matarić, 2005). Empirical studies involving companion robots such as PARO demonstrated measurable psychological and physiological benefits among elderly users, including reduced stress and increased social engagement (Wada and Shibata, 2007).

More recent conversational systems have extended these insights through multimodal dialogue support and dementia mitigation frameworks. State-machine-based multimodal dialogue systems and LLM-integrated monitoring robots have shown potential for sustaining daily conversational engagement among older adults (Suzaki and Numao, 2023; Numao and Kawamura, 2025). However, evaluation in these systems often remains qualitative, lacking computational frameworks for modeling relational dynamics.

Human-Centered and Well-Being–Oriented AI

Parallel to technological advances, human-centered AI research has emphasized alignment with human values, autonomy, and well-being. Russell's human-compatible AI framework highlights the importance of designing systems that support rather than override human agency [Russell, 2019]. Similarly, Shneiderman advocates human-centered design principles that prioritize augmentation over automation (Shneiderman, 2022).

Calvo et al. further propose ethical inquiry frameworks for AI systems that enhance psychological well-being and self-determination (Calvo et al., 2020). While these perspectives articulate normative design goals, they provide limited computational mechanisms for modeling relational interaction processes in applied service contexts.

Digital Twin Modeling and Human Simulation

Digital twin technologies originated in industrial informatics, where virtual replicas of physical systems enable monitoring, prediction, and optimization (Tao et al., 2019). More recently, the concept has expanded toward human-centric applications, including healthcare monitoring and behavioral simulation.

However, most digital twin implementations model individuals as passive data entities rather than relational agents embedded in social interaction. The Dual Digital Twin framework extends this paradigm by simulating both participants in a face-to-face encounter and incorporating evaluative observation of relational dynamics.

Dual Digital Twin Framework

Figure 1 illustrates the architecture of the Dual Digital Twin framework. The Dual Digital Twin (DDT) framework is proposed as a computational architecture designed to model, simulate, and reflect relationship-oriented face-to-face interactions. Unlike conventional dialogue systems that optimize task completion or response correctness, the DDT framework focuses on relational dynamics emerging through sustained human interaction. It achieves this by constructing digital representations of both participants in a service encounter and introducing an evaluative observer layer that analyzes interaction quality beyond informational exchange.

Design Principles

The framework is grounded in three design principles derived from the conceptual distinction between answer-oriented and relationship-oriented AI.

(1) Bidirectional Human Modeling

Face-to-face services involve mutual adaptation between participants. Therefore, modeling only one agent — as in conventional user simulation or dialogue policy optimization — is insufficient. The DDT framework represents both participants as behavioral twins, enabling simulation of reciprocal interaction dynamics.

(2) Integration of Explicit and Hidden Tasks

As discussed in Section 2, care and educational dialogue often involves both explicit informational tasks and implicit relational goals. The framework accommodates both by representing dialogue content and relational signals within the same interaction space.

(3) Observer-Based Relational Evaluation

Correctness-based evaluation cannot capture relational outcomes such as comfort or engagement. An independent observer layer is therefore required to analyze interaction continuity, emotional alignment, and well-being impact. Together, these principles establish the DDT framework as a reflective AI system rather than an automation engine.

Resident Twin Modeling

The Resident Twin represents the cognitive, emotional, and conversational characteristics of a care recipient or learner. It functions as a behavioral simulation model that generates responses conditioned on both dialogue input and internal state variables.

Key modeling dimensions include:

- **Cognitive Profile:** memory retention, orientation stability, recall variability
- **Emotional State:** affective tone, anxiety, comfort level
- **Conversational Coherence:** topic continuity, response latency, semantic drift
- **Engagement Disposition:** willingness to participate, fatigue, attentional fluctuation

These parameters allow the Resident Twin to simulate diverse interaction scenarios, including dementia progression stages or mood variability.

Importantly, the twin does not aim to produce factually correct responses but to emulate realistic conversational behaviors observed in care settings. This supports scenario-based training and interaction analysis.

Staff Twin Modeling

The Staff Twin models the communicative and relational strategies employed by caregivers or educators. It generates dialogue responses based on both informational and affective cues derived from the Resident Twin.

Key functional components include:

- **Dialogue Strategy Selection:** structured questioning, open-ended prompts
- **Empathy Modeling:** affective mirroring, reassurance, validation
- **Adaptive Communication:** pacing adjustment, topic shifting, repetition handling
- **Hidden Task Activation:** sustaining dialogue, encouraging participation

Through these mechanisms, the Staff Twin operationalizes both explicit tasks (e.g., cognitive assessment questions) and relational objectives (e.g., reducing anxiety or prolonging engagement).

The twin thus serves not merely as a conversational agent but as a simulation of professional practice.

Observer Twin and Relational Evaluation

The Observer Twin constitutes the evaluative core of the DDT framework. It monitors dialogue exchanges between the Resident and Staff Twins and assesses interaction quality through relational metrics rather than correctness measures.

Primary evaluation dimensions include:

- **Engagement Level:** dialogue duration, turn-taking balance
- **Emotional Alignment:** affective resonance, comfort signaling
- **Relational Continuity:** topic persistence, conversational flow
- **Well-Being Impact:** satisfaction proxies, emotional stabilization

By externalizing these metrics, the Observer Twin enables reflective analysis of interaction processes that are otherwise tacit in face-to-face services.

The observer also functions as a feedback generator. Evaluative outputs can be used to refine Staff Twin strategies, enabling iterative training loops and professional development simulations.

System Implementation on Dify

Overview

We implement the Dual Digital Twin framework as a multi-agent conversation simulator on Dify (advanced-chat mode), where two LLM agents—Resident Twin and Staff Twin—interact in a controlled loop, while maintaining separate conversation memories for each twin. The workflow is packaged as a Dify application named DualTwinSim configured in advanced-chat mode.

Core Components (Two Twins + Loop Controller)

(1) Loop Controller (simulation driver).
(2) Resident Twin (LLM node).
(3) Staff Twin (LLM node).

Memory Design (Twin-Specific State Separation)

To preserve role integrity and enable asymmetric knowledge/state, the workflow defines two conversation-scoped variables:

- conversation.staffMemory
- conversation.residentMemory

Turn Execution and Data Flow inside the loop

Within each iteration, the workflow executes the following sequence (as defined by the graph edges and node functions):

1. Loop-start → Resident LLM
2. Resident output → RESIDENT Answer
3. Resident answer → RESIDENT Memory
4. Update STAFF→ STAFF LLM
5. Staff output → STAFF Answer
6. Staff answer → STAFF Memory
7. Staff Memory → Update RESIDENT

Reset and Reproducibility

At the end of a run, a dedicated assigner clears both conversation.staffMemory and conversation.residentMemory, enabling reproducible repeated simulations.

Scenario Simulation and Training Protocol

The Dual Digital Twin (DDT) framework serves not only as an interaction simulator but also as a controlled experimental environment for analyzing how dialogue strategies influence relational and motivational outcomes in face-to-face services. This section describes the scenario configuration, interaction modes, and evaluation protocol used in the rehabilitation motivation study.

Scenario Design

Each simulation scenario begins with parameter initialization defining both cognitive and relational conditions of the Resident Twin. Core parameters include:

- Cognitive condition (e.g., normal aging, mild cognitive impairment),
- Emotional baseline (calm, anxious, discouraged),
- Motivational variability, and
- Interaction goal type (task-oriented, relationship-oriented, or hybrid).

The rehabilitation scenario centers on daily walking goals monitored through smartphone or smartwatch step-counting systems. The daily target is set at 6,000 steps, and each simulation instantiates one of two behavioral outcomes:

1. Goal Achieved
2. Goal Not Achieved

These conditions allow examination of dialogue effects under both positive and negative performance feedback.

Turn-Based Interaction Loop

Dialogue unfolds through a fixed-length interaction loop (maximum 12 turns) implemented in the Dify workflow. At each turn:

1. The Resident Twin generates an utterance conditioned on prior staff input, internal cognitive-emotional parameters, and resident-side memory.
2. The Staff Twin responds according to its configured interaction policy and accumulated staff-side memory.
3. Dialogue pairs are serialized and stored in twin-specific memory arrays.
4. The loop continues until termination criteria are met.

This structure ensures experimental reproducibility while allowing emergent conversational dynamics.

Interaction Modes

To model realistic care communication, three interaction modes are defined.

(1) Mode A: Task-Oriented Interaction

The Staff Twin prioritizes performance monitoring and corrective feedback. Dialogue characteristics include:

- Numerical reporting of step counts,
- Comparison against targets,
- Emphasis on improvement directives,
- Minimal topic deviation.

Evaluation in this mode focuses on task clarity and performance reinforcement. However, emotional responsiveness is limited.

(2) Mode B: Relationship-Oriented Interaction

The Staff Twin emphasizes encouragement and engagement continuity. Dialogue strategies include:

- Acknowledgment of effort,
- Praise-based reinforcement,
- Open-ended prompts,
- Acceptance of topic drift.

Success is measured by engagement duration, emotional stabilization, and willingness to continue interaction rather than performance accuracy alone.

(3) Mode C: Hybrid Interaction

In real-world caregiving, explicit and relational goals coexist. Hybrid mode integrates both paradigms. The Staff Twin provides performance transparency while embedding motivational scaffolding. Relational strategies are activated when signs of disengagement or discouragement appear.

This mode aims to balance:

- Informational clarity,
- Emotional alignment,
- Sustained behavioral motivation.

Observer-Based Evaluation Across Modes

Across all modes, the Observer Twin evaluates interaction logs using four relational metrics:

- Engagement Level (turn persistence, participation balance),
- Emotional Alignment (affective resonance),
- Relational Continuity (flow stability after negative feedback),
- Motivational Impact (expressed intention to continue rehabilitation).

When applicable, task performance accuracy is assessed as a secondary dimension.

By applying a unified evaluative framework across modes, the protocol enables systematic comparison of how dialogue strategies influence both behavioral adherence and relational sustainability.

Experimental Evaluation

To evaluate the effectiveness of the Dual Digital Twin (DDT) framework in modeling relationship-oriented service interaction, we conducted scenario-based simulations within a rehabilitation motivation context. The evaluation examines how different dialogue strategies influence engagement, emotional alignment, and motivational sustainability rather than task performance alone.

Simulation Setup

The experimental scenario involved daily walking rehabilitation monitored via smartphone or smartwatch step-counting systems. A target of 6,000 steps per day was established. Two outcome conditions were simulated:

- Goal Achieved
- Goal Not Achieved

All simulations were conducted under identical baseline parameters: mild cognitive decline, moderate motivational variability, and stable emotional disposition. Each interaction was executed through a 12-turn dialogue loop. Five simulation runs were conducted for each interaction mode with identical parameter settings (temperature = 0.7).
Three dialogue strategies were evaluated:

- Mode A: Task-Oriented Interaction
- Mode B: Relationship-Oriented Interaction
- Mode C: Hybrid Interaction

Quantitative Relational Evaluation

Dialogue outcomes were assessed by the Observer Twin using four relational metrics: Engagement Level, Emotional Alignment, Relational Continuity, and Motivational Impact.

Table 1. Comparative Relational Evaluation Scores

Interaction Mode	Engagement	Emotional Alignment	Relational Continuity	Motivational Impact
Mode A	Low	Low	Low	Medium
Mode B	High	High	High	High
Mode C	High	High	High	Very High

Task-oriented dialogue provided performance clarity but often reduced engagement after negative outcomes. Relationship-oriented dialogue sustained emotional rapport but sometimes lacked explicit behavioral reinforcement. Hybrid dialogue achieved balanced performance across relational and motivational dimensions.

Dialogue Pattern Comparison

Representative excerpts illustrate interaction differences under the "Goal Not Achieved" condition.

Mode A — Task-Oriented

Staff:
"You walked 4,200 steps today. The target is 6,000. Please increase activity tomorrow."
Resident:
"I see… I tried, but my legs were tired."
Observation:
Responses shortened quickly, and interaction typically terminated before reaching maximum turn length.

Mode B — Relationship-Oriented

Staff:
"You still walked quite a lot today. That effort really matters."
Resident:
"Do you think so? I was worried it wasn't enough."
Observation:
Residents elaborated emotionally, and dialogue persistence remained high, though performance targets were less emphasized.

Mode C — Hybrid

Staff:
"You reached 4,200 steps today — slightly below target. But you've been improving this week. That's excellent progress."
Resident:
"Really? I didn't realize I was improving."
Observation:
Residents accepted corrective feedback without discouragement, sustaining both engagement and goal orientation.

Engagement Persistence

Turn-length analysis further differentiated interaction modes.

Table 2. Turn-Length Analysis

Mode	Average Sustained Turns
Mode A	6.4
Mode B	11.2
Mode C	10.8

Task-oriented dialogue frequently terminated early following corrective feedback, whereas relational and hybrid interactions sustained conversational continuity.

Motivational Response Indicators

Observer analysis also examined forward-looking behavioral intent expressed by the Resident Twin.
Indicative responses included:

- Mode A: "I'll try… but it's difficult."
- Mode B: "I want to keep walking more."
- Mode C: "I think I can aim a little higher tomorrow."

Hybrid dialogue most consistently generated proactive rehabilitation intent.

Summary of Findings

Three key insights emerged:

1. Task clarity alone does not sustain engagement. Corrective feedback may reduce participation persistence.

2. Relational continuity stabilizes emotional response. Encouragement mitigates discouragement following negative outcomes.
3. Hybrid interaction maximizes motivational sustainability.

Integrating performance transparency with relational scaffolding supports both behavioral adherence and emotional well-being.

These results demonstrate the analytical value of modeling hidden relational tasks within simulated care interaction.

Conclusion

This paper examined the distinction between answer-oriented AI and relationship-oriented human services, arguing that contemporary artificial intelligence systems are primarily optimized for correctness, efficiency, and task completion, whereas face-to-face professions such as caregiving and rehabilitation depend fundamentally on relational continuity and emotional alignment.

Through the Dual Digital Twin (DDT) framework, we introduced a computational architecture capable of modeling both explicit task execution and hidden relational goals within a unified simulation environment. By representing both participants in an interaction as behavioral twins and incorporating an Observer Twin for relational evaluation, the framework enables systematic analysis of engagement, emotional resonance, and motivational impact beyond traditional correctness metrics.

Scenario-based simulations in rehabilitation dialogue demonstrated that purely task-oriented strategies often fail to sustain engagement following negative outcomes, while purely relational strategies may lack structured reinforcement. The Hybrid interaction mode—integrating performance transparency with motivational scaffolding—produced the most balanced outcomes, sustaining both behavioral adherence and emotional stability. These findings suggest that relational continuity functions not as an alternative to task execution but as a stabilizing mechanism that enables it.

More broadly, this work reframes the role of AI in human-centered services. Rather than automating relational labor, AI can illuminate and support the tacit mechanisms that sustain meaningful interaction. By externalizing hidden tasks and relational dynamics, the Dual Digital Twin framework positions AI as a reflective partner in professional development and well-being support.

As automation increasingly transforms answer-oriented domains, the enduring value of face-to-face services may lie in their capacity to foster human connection, motivation, and dignity. The future of AI in these contexts, therefore, should not be defined by replacement, but by its ability to enhance relational creativity and sustain meaningful work.

Acknowledgements

This work was supported by JSPS KAKENHI Grant Number JP22H00547 " Support system for early detection of symptoms and prevention of progression of dementia based on continuous biological and behavioural data."

References

Brown, T. B., et al. 2020: Language models are few-shot learners. Advances in Neural Information Processing Systems (NeurIPS), 33.

Russell, S. 2019: Human Compatible: Artificial Intelligence and the Problem of Control, Viking.

Calvo, R. A., et al. 2020: Supporting human autonomy in AI systems: A framework for ethical enquiry. AI & Society, 35, 523–535.

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

Suzaki, T., and Numao, M. 2023: State Machine-Based Multimodal Dialogue System for the Elderly Care Service, AAAI Spring Symposia 2023.

Numao, M., and Kawamura, M. 2025: Interactive Monitoring Robot for Mitigating Dementia Through Daily Conversations Using Multiple LLMs AAAI Spring Symposia.

Wada, K., and Shibata, T. 2007: Living with seal robots—Its sociopsychological and physiological influences on the elderly at a care house. IEEE Transactions on Robotics, 23(5), 972–980.

Lemon O., et al. 2006: An ISU dialogue system exhibiting reinforcement learning of dialogue policies: Generic slot-filling in the TALK in-car system, in: Demonstrations, pp.119–122. URL: https://aclanthology.org/E06-2009.

Wang Z., and Lemon O. 2013: A simple and generic belief tracking mechanism for the dialog state tracking challenge: On the believability of observed information, in: Proceedings of the SIGDIAL Conference, Association for Computational Linguistics, Metz, France, 2013, pp. 423–432. URL: https://aclanthology.org/W13-4067.

Young S., et al. 2013:Pomdp-based statistical spoken dialog systems: A review, Proceedings of the IEEE 101 1160–1179. doi:10.1109/JPROC.2012.2225812.

Zhao T., et al. 2017: Generative encoder-decoder models for task-oriented spoken dialog systems with chatting capability, URL: http://arxiv.org/abs/1706.08476. arXiv:1706.08476.

Yu C., et al. 2022: A chit-chats enhanced task-oriented dialogue corpora for fuse-motive conversation systems, URL: https://arxiv.org/abs/2205.
doi:10.48550/arXiv.2205.05886.