

An Exploratory Study of Agentic AI for Long-Term Healthcare Navigation.

Gargi Nandy^{a*}, Chun-Hua Tsai^b

^aInformation Systems and Quantitative Analysis, College of Information Systems and Technology, University of Nebraska at Omaha, Omaha, Nebraska, U.S.A., gnandy@unomaha.edu, ORCID: 0009-0000-2665-1440.

^bInformation Systems and Quantitative Analysis, College of Information Systems and Technology, University of Nebraska at Omaha, Omaha, Nebraska, U.S.A., chunhuatsai@unomaha.edu, ORCID: 0000-0001-9188-0362.

Publishing history: Submitted*: 07 February 2026, Revised*: e.g 01 June 2021, Accepted*: 05 April 2026, Published*: e.g 01 August 2021

Abstract. Long-term healthcare poses significant challenges for families managing chronic and developmental conditions, requiring sustained coordination, information management, and preparation across repeated clinical encounters. For many families, care does not occur through a single clinical visit but unfolds over months and years through ongoing appointments, therapy sessions, follow-ups, and changing developmental needs. This makes caregiving a longitudinal information-navigation process in which families must remember past concerns, interpret new guidance, and prepare for future decisions. Family caregivers, particularly those supporting children with complex conditions such as hypoxic-ischemic encephalopathy (HIE), bear substantial cognitive and emotional burdens while navigating fragmented healthcare systems. Recent advances in generative and agentic artificial intelligence suggest new opportunities to support longitudinal caregiving work by helping caregivers connect past concerns with upcoming care needs. In this paper, we present an exploratory study of a memory augmented, multi agent conversational AI system designed to support HIE family caregivers in navigating multiple medical appointments. The system employs a three-layer architecture integrating a caregiver-facing user interface, an agentic coordination layer, and long-term memory to enable continuity, proactive preparation, and context-aware interaction over time. We conducted a qualitative, scenario-based simulation study using focus groups (N=9) to examine participants' perceptions of the system's usefulness and design implications. Findings indicate that participants valued the system's ability to organize prior discussions, maintain continuity across appointments, and reduce cognitive burden through structured summaries, appointment-centered guidance, and long-term memory retrieval. Participants also emphasized the importance of relational personalization, emotionally sensitive information presentation, proactive guidance, and human oversight to ensure trust and safety. This work contributes design implications for agentic AI in long-term healthcare and highlights the potential of memory-augmented systems to support caregivers' preparedness, sense-making, and engagement across complex lifelong healthcare journeys without replacing professional clinical judgment.

Keywords: Agentic AI, Long-term Healthcare, Care Coordination, Conversational AI, Hypo-Ischemic Encephalopathy

1. Introduction

Long-term healthcare needs are a growing concern for both aging adults and children with chronic conditions, posing critical challenges for patients, families, and the public spending. Global populations are aging rapidly – by 2030, one in six people will be over age 65, and this proportion will double by 2050 (Khalil et al., 2026). As people live longer, many will require ongoing assistance with daily activities and management of chronic

illnesses, creating an urgent need for sustainable long-term care solutions. Concurrently, the prevalence of chronic health conditions in young populations has risen: nearly one in three young people (ages 5–25) now lives with a chronic condition or functional limitation, a dramatic rise over the past two decades (Wisk and Sharma, 2025). These children often face lifelong health challenges, from asthma and diabetes to neurodevelopmental disabilities, and will need continuous engagement with healthcare services into adulthood. Yet current healthcare systems struggle to support such long-term needs – for example, many youth “are at risk of disengaging with care” during the transition from pediatric to adult services due to poor care coordination (Rivero, 2025). Without better support, gaps in care can lead to exacerbations of disease and diminished quality of life for these individuals and families.

Family caregivers bear much of the burden in long-term care, coordinating complex care regimens over years or decades. Parents of children with chronic illnesses, for instance, serve as “around-the-clock” care coordinators and advocates for their child. This role entails managing numerous medical appointments/decisions, therapies, and medications, often while juggling work and other responsibilities. Studies show that caring for a child with long-term health needs places significant strain on families, leading to financial hardship, social isolation, and mental health impacts and self-grieving (Piran et al., 2017; Raap et al., 2025). For example, hypoxic-ischemic encephalopathy (HIE) – a birth-related brain injury – is a leading cause of long-term disability in children and can result in conditions such as cerebral palsy, epilepsy, and cognitive impairments requiring lifelong care (Korf et al., 2023). HIE families bear a substantial and enduring financial burden. This includes the costs of continuous caregiving, specialized therapeutic services, and the procurement and maintenance of assistive devices required to sustain the child’s functional independence and quality of life.

Generative Artificial intelligence (Gen-AI) has emerged as a tool to help address the complexities of long-term care. In particular, conversational agents – AI systems that interact through natural language – offer a human-friendly interface to support patients and caregivers in managing health over extended periods such as ChatGPT (Fui-Hoon Nah et al., 2023). The advances in large language models (LLMs) enable agents to engage in nuanced, context-aware dialogues, making them suitable for guiding users through health decisions, providing education, and offering emotional support (Zheng et al., 2025). However, traditional chatbots have had limitations, such as short memory spans and scripted responses, which reduce their usefulness in longitudinal care. To truly assist with lifelong healthcare navigation, an AI agent must exhibit agentic qualities: it should operate with a degree of autonomy, pursue long-term goals, adapt to changing needs, and crucially, maintain persistent long-term memory of user information and interactions (Sumers et al., 2023). Agentic AI refers to AI systems capable of independent decision-making and action towards goals. Key characteristics include being autonomous (acting with minimal human input), goal-driven, adaptive to changing circumstances, and persistent beyond isolated tasks (Acharya et al., 2025).

In this paper, we present an exploratory study of a memory-augmented, multi-agent conversational system. The system is designed to support HIE family caregivers in preparing for multiple healthcare appointments in long-term care settings. The system employs a three layer architecture consisting of a user interface layer, an agentic coordination layer, and a long term memory layer that together enable longitudinal continuity, proactive preparation support, and context aware interactions. Specialized agents coordinate conversational response generation, memory retrieval, and structured summarization, allowing the system to reference prior discussions and upcoming appointments without exposing full conversation histories. Our design of a caregiver facing interface that supports recalling past topics, managing future appointments, and receiving structured follow up prompts to guide clinical preparation. To explore the system’s perceived usefulness and design implications, we conducted a qualitative, scenario based simulation study using focus groups (N=9), examining how participants envisioned the system supporting caregivers’ ongoing planning, sense making, and readiness across complex and evolving care journeys.

Based on the focus group discussion, the participants viewed the agentic, memory augmented system as a promising support tool for long term caregiving, particularly for organizing information, maintaining continuity across appointments, and reducing cognitive burden. They valued features such as topic based summaries, long term memory retrieval, appointment centered organization, and source transparency, which helped surface relevant prior discussions without requiring caregivers to revisit full conversation histories. Participants emphasized that effective personalization should function as a relational signal—acknowledging caregivers’ histories, identities, and evolving concerns—rather than merely retrieving relevant content, and it should be reinforced through both language and interface design. Participants also highlighted the importance of emotional sensitivity, noting that caregivers often seek support under stress or fatigue; this makes information

overload a critical concern and underscores the need for concise, skimmable, and prioritized outputs. Finally, participants responded positively to the system's proactive, appointment oriented guidance, particularly its ability to suggest next steps and questions for upcoming visits, while consistently stressing that such guidance should remain bounded by human oversight to ensure trust, safety, and alignment with professional expertise.

This work contributes to digital government, public health, and AI in healthcare by illustrating how agentic, memory-augmented systems can support longitudinal care coordination in complex, multi-stakeholder health ecosystems such as Health Information Exchange (HIE). By foregrounding caregivers' preparatory work across repeated medical and therapy appointments, the study highlights the role of AI as a public-facing infrastructure that can enhance continuity, transparency, and readiness without supplanting professional authority, aligning with digital government goals of accountable, human-centered service delivery. From a public health perspective, the findings support how agentic AI could reduce cognitive and emotional burden for caregivers navigating long-term care, potentially improving engagement, preparedness, and equity in access to health information. The study also offers design implications for future agentic AI development, emphasizing the need for longitudinal memory that balances continuity with data minimization, personalization that functions as a relational and empathetic signal rather than mere relevance, emotionally sensitive information presentation to mitigate overload, and proactive yet bounded guidance that incorporates human oversight. Together, these insights inform the responsible design of agentic AI systems that operate within healthcare and public-sector contexts, supporting complex care journeys while maintaining trust, safety, and ethical accountability.

2. Related Work

Caregiver appointment preparation is increasingly recognized as a distinct and ongoing form of health-related work that unfolds longitudinally across providers, appointments, and evolving care needs, rather than as a task associated with a single clinical encounter (Tennant et al., 2022). Caregivers—particularly those supporting children with complex or neurodevelopmental conditions—must continuously manage, interpret, and carry forward information across sequences of interrelated visits. This work includes recalling prior clinical guidance, tracking changes in symptoms or developmental progress, coordinating recommendations from multiple providers, and anticipating questions or decisions that may arise in future encounters (Schurgin et al., 2021). As a result, preparation for any given appointment is inherently shaped by what occurred in earlier visits and by expectations for future care, creating a cumulative information and coordination burden over time (Tennant et al., 2022).

These demands are especially pronounced in long-term pediatric and neurodevelopmental contexts such as hypoxic-ischemic encephalopathy (HIE). HIE is a neonatal brain injury associated with a complex and heterogeneous set of antenatal and intrapartum risk factors, many of which are non-hypoxic and non-ischemic in origin (Korf et al., 2023). Depending on severity, outcomes may include mortality, cerebral palsy, epilepsy, global developmental delay, intellectual disability, and later-emerging cognitive or behavioral difficulties that may not be reliably assessed in early childhood (Schreglmann et al., 2020; Lee and Glass, 2021). Care for children with HIE typically involves sustained engagement with multidisciplinary services, including pediatric specialists, physical therapists, occupational therapists, and speech or developmental therapists, with recommendations that evolve as the child grows (Pierrat et al., 2005; Huntingford et al., 2024; Songire et al., 2024). Family caregivers are often responsible for maintaining continuity across these encounters, ensuring that guidance from one provider informs subsequent visits, and monitoring developmental outcomes across multiple domains. As a result, each appointment is rarely an isolated event; instead, it is embedded within a broader caregiving trajectory that demands ongoing preparation, coordination, and sense-making over months or years (Debillon et al., 2018).

Prior research has documented that caregivers managing complex care frequently experience information overload and uncertainty about how to engage effectively during clinical encounters. Caregivers must remember a wide range of details—including therapy goals, developmental milestones, medication changes, and follow-up instructions—often under conditions of emotional stress, fatigue, or time pressure (Kiiski et al., 2025). These cognitive and emotional demands can limit caregivers' ability to recall information accurately, identify what is most relevant for a particular appointment, or formulate questions during visits. Studies of caregiving work emphasize that this coordination and memory labor is largely invisible within healthcare systems and is rarely supported by formal tools designed to assist caregivers in preparing for or reflecting on appointments (Schurgin et al., 2021).

Research on appointment preparation and pre-visit information seeking further highlights these challenges. Early work on the concept of the “prepared patient” demonstrates that patients and caregivers actively seek information in advance of appointments to feel more confident, informed, and engaged during consultations (Hu et al., 2012). This preparatory work may include searching for explanations of medical terms, learning about treatment options, or attempting to formulate questions ahead of time. While such efforts can empower caregivers, they are typically self-directed and fragmented, relying on disparate sources such as online resources, peer communities, personal notes, or informal digital tools rather than being systematically supported by healthcare technologies (Hu et al., 2012; Lu and Schulz, 2024). In longitudinal care contexts, these limitations are amplified, as caregivers must not only prepare for upcoming visits but also integrate insights from past encounters into future decision-making.

Recent developments in artificial intelligence (AI)—particularly in generative AI—have sparked considerable interest in their potential to enhance appointment preparation and care coordination. Empirical studies of AI driven appointment management platforms report several advantages, including greater scheduling efficiency, a reduced administrative workload, and improved patient perceptions of access and convenience (Knight et al., 2023; Wood et al., 2025). These advances give rise to increasingly autonomous AI agents. The term *agentic AI* denotes systems that can make independent decisions and take actions directed toward predefined objectives. Such systems are characterized by autonomy (minimal human intervention), goal orientation, adaptability to evolving circumstances, and persistence that extends beyond discrete tasks (Acharya et al., 2025). Nevertheless, the application of agentic AI to long term care remains an under explored area of research.

These literature highlights the complexity and longitudinal nature of caregiver appointment preparation, particularly in therapy-intensive and neurodevelopmental care contexts such as HIE. Caregivers’ preparation work extends beyond individual visits and involves ongoing coordination, memory work, and sense-making across multiple appointments and providers. While caregivers actively seek information to support this work, existing tools—including emerging AI-based systems—offer limited support for maintaining continuity, integrating past and future care contexts, or reducing cognitive and emotional burden over time. This gap points to the need for new approaches that explicitly account for the longitudinal, cumulative nature of caregiving and appointment preparation in long-term healthcare settings.

3. Methodology

3.1. User Scenario

To ground the study in a realistic long term healthcare context, we introduced a user scenario involving Maya, a 32 year old parent caring for her two year old son, Arin, who was diagnosed with hypoxic ischemic encephalopathy (HIE). HIE is a serious neonatal disorder that arises when an infant’s brain experiences both hypoxia and ischemia—insufficient oxygen delivery coupled with diminished cerebral blood flow—during the peripartum period. Following his diagnosis, Arin has continued to receive long term follow up care involving multiple health care providers and therapy services. The scenario emphasizes the longitudinal and cumulative nature of caregiving work, highlighting how preparation for one appointment is shaped by information and recommendations from earlier encounters. This scenario guided the proposed AI system design and served as a shared reference point during the focus group discussions, enabling participants to reflect on the system’s potential role in supporting caregiver preparation without requiring personal caregiving experience. Participants were asked to consider how a caregiver like Maya might use a conversational AI system to revisit prior discussions, organize information across appointments, and prepare for upcoming medical or therapy visits within this ongoing care journey.

3.2. System Design

In this exploratory study, we developed a multi agent conversational architecture to support longitudinal healthcare, especially the preparation for medical and therapy related appointments for patients and caregivers in long-term healthcare settings such as HIE. The system consists of three layers: the User layer, the Agentic layer, and the Memory layer. The User layer provides a front end interface through which users submit queries and receive AI generated responses. The Agentic layer orchestrates several specialized agents—a domain specific (HIE) agent, a memory agent, and a summary agent—to coordinate the workflow. Finally, the




Fig. 1 – System Workflow Diagram

Memory layer manages long term episodic memory of past interactions and integrates domain knowledge drawn from the literature and other authorized data sources. The overall system workflow is presented in Figure 1. The interaction flow is as follows: user queries are processed by the HIE agent, which generates context aware responses by integrating the current input with relevant domain knowledge and previously stored interaction history. A memory agent retrieves relevant past conversations from long term memory, while a summary agent distills salient information from effective interactions and persists concise summaries for future use.

The **Agent layer** was implemented using the Dify platform to define workflow logic and provide access to the open-source large language model (GPT-OSS 120B) (Mirzac et al., 2025). The agentic layer comprised of 3 agents- HIE agent, memory agent, and summary agent.

- **HIE agent:** this agent functions as the primary conversational agent and is responsible for generating caregiver-facing responses. The prompt instructs the model to respond as a supportive caregiving assistant, to reference relevant prior interactions when available, and to frame responses in relation to upcoming medical or therapy appointments. Rather than receiving full conversation transcripts, the HIE agent is provided with a compact, injected context block assembled at runtime from previously generated summaries and stored appointment information. This bounded context design enables the agent to maintain continuity across interactions while avoiding prompt overloading and unnecessary exposure to raw historical data.
- **Memory agent:** this agent is responsible for retrieving and persisting longitudinal information across sessions. It operates over the long-term conversational knowledge base and retrieves structured summaries, topic labels, and appointment metadata associated with the current user. These retrieved elements are selectively surfaced to the HIE agent as contextual input, allowing prior discussions to inform response generation without requiring caregivers to restate background information. The Memory agent also governs what information is written back to long-term storage, ensuring that only summarized and structured representations—rather than full conversational transcripts—are persisted. This design reflects an intentional balance between continuity and data minimization in a healthcare context.
- **Summary agent:** this agent is invoked after each interaction to generate a standardized representa-

tion of the exchange for long-term storage. Its prompt is schema-constrained and designed to guide the model toward producing concise, structured summaries that capture salient caregiving-related information while remaining consistent across interactions. Output is restricted to a strict JSON format to support reliable storage and retrieval by the Memory agent. Domain constraints embedded in the prompt help ensure that summaries remain appropriately scoped to HIE-related care, therapy contexts, and early-intervention settings, supporting longitudinal reuse without storing full conversational transcripts

The **Memory layer** incorporates two distinct knowledge bases, both implemented using a vector database to support efficient semantic retrieval. The first knowledge base stores long-term conversational memory, including summarized interaction histories that enable continuity across sessions. The second knowledge base consists of pre-existing domain knowledge derived from a prior HIE-focused study, which includes curated and validated content related to hypoxic-ischemic encephalopathy and caregiving practices from (Nandy et al., 2025). This separation allows the system to dynamically combine personalized caregiver context with stable, domain-specific reference information during response generation.

The **User layer** consisted of a Streamlit-based interface ¹ that supported interaction and visualization, enabling rapid prototyping and evaluation of the system's memory-augmented, agent-based design, presented in Figure 2. The interface presents a list of recent topics discussed in a persistent left-hand panel (A), enabling caregivers to view and select summarized prior conversations related to specific concerns (e.g., therapy scheduling or feeding challenges). Selecting a topic allows the system to pull and display the corresponding prior discussion, supporting recall and continuity without requiring caregivers to reread entire conversation histories. The interface also includes an appointment management component (B), where caregivers can save and view upcoming medical or therapy appointments. These appointments are displayed directly within the UI and surfaced proactively during conversations, allowing the system to reference relevant upcoming visits when supporting preparation. Caregivers can add, view, or manage appointments, and the system uses this information to contextualize responses and suggest visit-specific guidance. Upcoming appointments are automatically prioritized, enabling proactive reminders and personalized preparation guidance informed by memory-derived conversation summaries. By linking structured appointment data with historical caregiver context, the system maintains continuity between past interactions and future clinical events, thereby enhancing planning support and visit readiness.

Once a prior topic is selected, caregivers can continue the conversation in the main interaction panel (C), where the system integrates long-term conversational memory with appointment context to generate responses that reflect previously discussed issues. This design allows caregivers to ask new questions while maintaining continuity with earlier interactions. For example, caregivers may inquire about long-term outcomes related to an earlier concern and receive responses grounded in both prior discussions and upcoming appointments. The system further supports preparation by offering structured follow-up prompts and suggested questions for clinicians (D), helping caregivers identify what to ask during upcoming visits without requiring them to formulate questions entirely on their own.

3.3. Study Design

This study employed a qualitative, scenario-based simulation design using focus group discussions to examine how an AI-based conversational system could support family caregivers in preparing for medical and therapy-related appointments. To ground the discussion in a realistic caregiving context, participants were introduced to the Maya scenario. The scenario was designed to reflect the ongoing and cumulative nature of caregiving work, particularly the challenges of organizing information, recalling prior discussions, and preparing for multiple, evolving appointments over time. Scenario-based simulation enabled participants to reason about the system's potential role in appointment preparation without requiring direct personal caregiving experience (Rosson and Carroll, 2007).

A total of nine participants (P1–P9) took part in the study across two focus group sessions. One session included four participants (P1–P4), and the other included five participants (P5–P9). Participant demographic information is summarized in Table 1. Participants were recruited through direct email invitations and included a mix of AI experts and university graduate students, enabling diverse perspectives on both the technical design of the system and its potential applicability in real-world caregiving contexts. All participants

¹<https://streamlit.io/>




Fig. 2 – The interface includes (A) a persistent topic history panel for recalling prior conversations, (B) an appointment management module that surfaces upcoming visits and informs preparation, (C) a main interaction panel that integrates long-term memory with appointment context during ongoing dialogue, and (D) structured prompts that suggest follow-up questions and preparation guidance for upcoming clinical encounters.

were adults and fluent in English. Participation was voluntary, and each participant received a 20\$ gift card upon completing the session. Following consent, the investigator introduced the purpose and structure of the session and presented a scenario-based caregiving simulation of Maya that framed the remainder of the discussion. The simulation depicted a parent engaged in long-term caregiving for a child with ongoing developmental needs and preparing for multiple upcoming therapy-related appointments. The scenario was designed to reflect realistic caregiving challenges, such as organizing prior information, anticipating questions, and preparing for sequential medical encounters.

One author of this paper acted as moderator and performed a live demonstration of the AI driven chatbot using a scenario based simulation. The demonstration showed how the system can retrieve and organize information from prior conversations to support appointment preparation, surfacing relevant past interactions, summarizing key details, and linking historical context to upcoming appointments. After the demo, participants took part in a guided group discussion. Using a semi structured protocol, the moderator asked participants to reflect on the system's clarity, relevance, and perceived usefulness, as well as its potential role in helping caregivers prepare for medical or therapy appointments over time. Although each discussion followed the same overall structure, participants were encouraged to elaborate on issues they deemed important, allowing unanticipated themes to emerge organically.

Both focus group sessions were conducted remotely via Zoom, with each session lasting approximately 60 minutes, and were moderated by the moderator. A remote format was selected to reduce logistical barriers to participation and to accommodate participants' varied schedules and locations. Prior research has shown that videoconference-based qualitative methods, including remote focus groups, can support rich discussion and data quality comparable to in-person sessions, while offering increased flexibility and accessibility for participants (Archibald et al., 2019). Before the session began, participants were informed about the study procedures and the nature of data collection. Explicit consent was obtained from all participants prior to starting the presentation and system demonstration, including consent to audio-record and transcribe the session for research purposes.

Code	Gender	Role/Title	Education
P1	M	Postdoctoral Researcher	Artificial Intelligence
P2	M	Doctoral Student	Computer Science
P3	M	Doctoral Student	Public Administration
P4	M	Doctoral Student	Biomedical and Translational Sciences
P5	F	Doctoral Student	Biomedical Informatic
P6	F	Graduate Student	Data Science
P7	M	Doctoral Student	Computer Science
P8	M	Doctoral Student	Health and Kinesiology
P9	M	Doctoral Student	Biomechanics

Tab. 1 – Participants Demographics.

3.3.1. Data Collection and Analysis

With participants' consent, both Zoom focus group sessions were audio-recorded and transcribed verbatim. The resulting transcripts constituted the primary data for this study and were analyzed by the primary investigator, conducting a thorough thematic analysis. Prior research has shown that the qualitative method is the flexible way for identifying and interpreting patterns across data (Braun and Clarke, 2006). Thematic analysis was selected due to its suitability for exploratory, design-oriented research and its ability to capture commonalities and variations in participants' perspectives. Analysis focused on participants' reflections on the AI-based system's perceived usefulness, clarity, relevance, and potential role in supporting caregiver appointment preparation within the scenario-based simulation. Themes were iteratively identified through close reading of the transcripts and refinement across the dataset by two co-authors of this paper. This study was reviewed and approved by the University's Institutional Review Board (IRB). All study procedures were conducted in accordance with institutional guidelines and ethical standards for research involving human participants. Participants provided informed consent prior to participation, and appropriate measures were taken to protect participant privacy and confidentiality throughout the study.

4. Findings

Analysis the focus group discussions revealed several recurring themes related to how participants perceived the AI-based system's potential to support caregiver appointment preparation. Overall, participants viewed the system as promising for organizing information, supporting longitudinal care, and reducing caregiver cognitive burden, while also identifying important design considerations around transparency, personalization, and safety. The findings are organized into the following thematic areas.

4.1. Agentic Long-Term Memory and Personalization

The participants expressed appreciation for several existing feature level capabilities of the system, particularly those that supported organization, continuity, and transparency. Rather than viewing the prototype as a purely conceptual demonstration, participants recognized concrete design elements that they believed would already be useful for caregivers managing ongoing care. One commonly appreciated capability was the system's ability to organize prior conversations into topics and summaries, which participants described as reducing the effort required to revisit long interaction histories. Participants noted that having conversations already grouped and summarized helped surface relevant information without requiring caregivers to manually search through earlier interactions. As P3 explained, the value lay in how the system structured information for quick access:

"Here are a list of things that you could choose... based on whatever the challenge was... giving the caregiver options... at least two or three, then go from there." — P3

Similarly, P2 highlighted that organizing past conversations into identifiable sections would make it easier for caregivers to understand what information belonged together:

“When you’re recalling all that previous conversation, the user should need to know the different section... so they know what this is about.” — P2

Participants also valued the system’s long-term memory functionality, particularly its ability to retrieve past conversations and integrate them into current responses. This capability was described as enabling continuity across appointments and supporting a sense of progression in the caregiving journey, rather than treating each interaction as isolated. P7 articulated how this continuity could support preparation over time:

“The parent can look over the calendar and see the appointments... maybe a summary of the appointment... what can I expect in this appointment.” — P7

Participants recognized that retrieving past interactions could help caregivers avoid repeatedly explaining background information and instead focus on current concerns and next steps. Another positively received feature was the system’s emphasis on source transparency. Participants appreciated being able to see where information originated, especially in a healthcare context where credibility is critical. P9 explicitly noted that the existing system already provided meaningful transparency:

“It is providing the transparency of the sources... where the information and the knowledge base is coming from.” — P9

Participants consistently emphasized that personalization was critical to engagement, but they framed personalization as more than retrieving relevant content. Instead, they described a need for interactional personalization—where responses explicitly acknowledge the caregiver’s history, concerns, and identity. P2 suggested that directly addressing the caregiver persona and referencing prior challenges would make the system feel more supportive and engaging:

“You can say that, ‘Maya, your child... has been having difficulty in swallowing, and what we recommend is...’ This is more personalized and more engaging for me.” — P2

Participants viewed personalization as a relational signal rather than a technical function. P2 elaborated that personalization required the system to sound less descriptive and more directly engaged with the caregiver’s situation, noting that generic explanations felt emotionally distant even when they were accurate:

“It’s not personally empathetic... it’s a little bit descriptive type. It could be more empathetic to the user, like, ‘Oh, Maya, you have discussed about this thing...’” — P2

This reflection highlights that participants interpreted personalization as an indicator of attentiveness and care, not simply relevance. Addressing the caregiver by name and explicitly recalling prior discussions was seen as a way to acknowledge the caregiver’s ongoing journey and evolving concerns, reinforcing a sense of one-to-one support over time. Participants further emphasized that personalization should be supported by visible interface cues that help caregivers recognize continuity across interactions. P2 suggested using visual differentiation, such as color-coding, to make relationships between past and present conversations more legible:

“For each topic, you can select another color... so the user can easily recognize, okay, these two topics look the same color, so this conversation is from this topic.” — P2

These suggestions indicate that participants expected personalization to be reflected not only in language, but also in how information is stored, organized and presented, helping caregivers quickly reorient themselves when revisiting prior discussions. These reflections indicate that participants recognized several existing design elements—such as structured summaries, long-term memory retrieval, source transparency, and appointment-centered organization—as meaningful supports for long-term caregiving works.

4.2. Agentic Informational and Emotional Support

While participants generally valued the overall information provided by the system, they also raised consistent concerns about information overload, particularly in light of the emotional contexts in which caregivers are likely to seek support. Participant’s emphasized that caregiving-related information seeking does not occur

in a neutral cognitive state; rather, caregivers may engage with the system while experiencing stress, fatigue, anxiety, or uncertainty related to their child's condition or upcoming appointments. These emotional factors were viewed as shaping caregivers' capacity to process and act on information, especially when responses are dense or highly detailed. Several participants explicitly connected the effectiveness of information presentation to caregivers' psychological readiness. P1 highlighted that the same volume of information may be perceived very differently depending on emotional state, noting that caregivers in heightened or "unconventional" emotional contexts may struggle to absorb extensive textual responses:

"It depends a lot on the psychological mentality... it's a lot of information that is being provided. If there is a possibility of concising this information and highlighting the keywords, that can be useful... especially in an unconventional emotional state." — P1

This reflection suggests that the participants did not view information overload solely as a function of quantity, but as the interaction between content volume and emotional readiness. In this framing, even accurate and even relevant information may become burdensome if presented without consideration for caregivers' fluctuating emotional and cognitive capacity. Other participants echoed this concern, noting that as questions become more complex, responses may become difficult to navigate. P6 warned that responses drawing from multiple sources could quickly become overwhelming:

"If you're asking more complicated questions and it's drawing from more information, that could get a little unwieldy." — P6

P5 similarly highlighted that even summarized content could feel cognitively demanding when caregivers are trying to quickly orient themselves, suggesting that summaries themselves require additional visual prioritization:

"This summary is actually great... but sometimes it is a little wordy... some of these words could be highlighted, helping me read through these long summaries faster, so that I can prioritize on some of these keywords." — P5

Across these reflections, participants consistently framed effective information support as requiring sensitivity to caregiver's emotional and situational contexts. Rather than advocating for more comprehensive responses, participants emphasized the importance of the outputs that are skimmable, prioritized, and structured in ways that allow caregivers to quickly identify what matters most in that moment. This suggests that emotional readiness functions as an implicit design constraint—systems intended to support caregiving must account not only for informational needs, but also for caregivers' limited attention, emotional state, and varying capacity to engage deeply with the content. Participants appeared to recognize a tension between leveraging rich contextual memory and maintaining outputs that remain cognitively manageable, particularly for caregivers who may already be under significant emotional strain.

4.3. Agentic Recommendation and Proactive Guidance

Participants also responded favorably to the system's appointment-oriented framing, particularly the idea of linking retrieved information and summaries to upcoming medical or therapy visits. This alignment with real caregiving workflows helped participants envision how the system could fit into caregivers' routines. P7 described this framing as helping caregivers anticipate and prepare rather than react:

"You're not just asking questions randomly... you're preparing for something specific that's coming up." — P7

Several participants described proactive guidance as especially valuable given the long-term and cumulative nature of caregiving. P2 noted that caregivers often need help recalling what was previously discussed and identifying what should happen next:

"You may recall what was our last interaction about... what was lastly suggested, and what should I do next... so the user can easily go back and recall." — P2

This framing positions proactive guidance as a form of cognitive scaffolding, helping caregivers bridge past interactions and future appointments without requiring them to independently reconstruct context or action

steps. P7 similarly emphasized the value of anticipatory support, suggesting that guidance before and after appointments could help caregivers feel more prepared and less overwhelmed:

“I really like that you can ask, like, what can I expect in this appointment... when you’re taking care of a child that is mentally and physically draining you.” — P7

These reflections suggest that participants viewed proactive guidance as a way to reduce uncertainty and emotional burden, particularly in moments when caregivers may struggle to identify appropriate questions or preparation strategies on their own. At the same time, participants stressed the importance of human oversight, particularly in a healthcare context. P1 proposed incorporating clinician validation into the system’s workflow:

“You can make these assertions get validated by the healthcare expert... and then feed that info within the agent itself to adapt further...” — P1

Additional comments reinforce that participants did not view oversight as a limitation of the system, but as a necessary condition for trust and responsible use of AI. Participants expressed concerns about relying solely on automated outputs in sensitive healthcare situations and emphasized that expert involvement could help ensure accuracy and appropriateness while preserving caregiver confidence in the system’s guidance.

Together, these reflections suggest that participants envisioned the system as an attentive, guided support tool that personalizes interaction, proactively assists with preparation, and remains responsibly bounded through human oversight. Rather than expecting the system to replace professional expertise or caregiver judgment, participants described an ideal role for AI as augmenting caregivers’ preparatory work through structured guidance, attentiveness, and accountable support mechanisms. Participants especially valued the system’s ability to proactively suggest next steps, follow-up questions, or preparation strategies, reducing the burden on caregivers to formulate questions on their own. This form of structured prompting was seen as particularly beneficial under conditions of stress or uncertainty, reinforcing expectations that AI should support caregivers in navigating complex care journeys while remaining aligned with professional expertise and ethical responsibility.

5. Discussion

This paper adopts a scenario-based, qualitative approach to examine how an agentic, memory-augmented conversational AI system might support family caregivers in long-term healthcare contexts. Grounded in a realistic caregiving scenario involving a parent managing ongoing care for a child with HIE, the study combines a multi-layer system design with focus group discussions to explore participants’ interpretations of the system’s role, strengths, and limitations. By demonstrating a three-layer architecture that integrates a caregiver-facing interface, coordinated specialized agents, and long-term memory, the approach emphasizes longitudinal continuity, proactive appointment preparation, and context-aware interaction. The scenario-based simulation enabled participants to reason about caregiving needs and system behaviors without requiring personal caregiving experience, making it possible to surface design considerations relevant to complex, evolving care trajectories.

Our study work highlights several key findings regarding the perceived value and design implications of agentic AI in long-term healthcare. Participants viewed the system as a promising support tool for organizing information, maintaining continuity across appointments, and reducing caregivers’ cognitive burden, particularly through topic-based summaries, memory retrieval, and appointment-centered guidance. At the same time, they emphasized that effective personalization should function as a relational and empathetic signal, supported by both language and interface design, and that information presentation must be sensitive to caregivers’ emotional states to avoid overload. Participants further valued proactive guidance that helps caregivers anticipate next steps and formulate questions for upcoming visits, while consistently underscoring the importance of human oversight to ensure trust, safety, and alignment with professional expertise. Based on these findings, we argue the following three design implications for adopting agentic AI in long-term healthcare.

Design Implication 1: Support Long-Term Growth Through Persistent, Agentic Memory. Agentic AI systems designed for long-term healthcare should incorporate persistent, longitudinal memory that evolves alongside

patients and caregivers over time. Such memory should integrate caregivers' previously asked questions, appointment histories, medical records and contextual concerns with up-to-date medical literature and domain-specific knowledge. By combining lived caregiving context with current clinical guidance, the agent can help caregivers navigate complex healthcare systems more effectively, maintain continuity across providers and appointments, and reduce the need to repeatedly restate background information as care needs evolve.

Design Implication 2: Tailor Information to Caregivers' Changing Cognitive and Emotional Capacity. Agentic systems should provide information that is not only personalized in content but also calibrated to caregivers' varying informational needs and emotional states at different stages of the care journey. At times, caregivers may need high-level, digestible guidance to manage overwhelming situations such as coordinating multiple upcoming appointments, while at other times they may have the capacity to engage more deeply with detailed medical knowledge or emerging research. Agents should dynamically adjust the depth, format, and pacing of information and actively guide caregivers toward learning and updating themselves when appropriate, rather than presenting uniformly dense or generic responses.

Design Implication 3: Enable Proactive, Guided Support With Human Oversight. Beyond responding to caregiver-initiated questions, agentic AI systems should take on a proactive role by preparing the upcoming medical appointments, anticipating caregivers' evolving needs and offering timely recommendations and guidance. For example, as a child—especially those who require chronic long term care—transitions across developmental stages (from infancy to toddlerhood to school age), the agent could proactively surface relevant considerations, developmental expectations, and suggested questions for clinicians. Importantly, such proactive guidance should be developed in collaboration with healthcare professionals and incorporate mechanisms for expert validation, ensuring that recommendations remain accurate, trustworthy, and aligned with clinical standards while preserving caregivers' confidence in the system.

This paper contributes to digital government and public health by demonstrating how agentic, memory-augmented AI systems can function as accountable, human-centered infrastructure to support caregivers navigating long-term healthcare systems. By foregrounding longitudinal continuity, proactive preparation, and emotionally sensitive information delivery, the study highlights how AI can enhance caregivers' capacity to engage with complex life-long care ecosystems without displacing professional authority. The findings support the potential of such systems to reduce cognitive burden, improve preparedness for clinical encounters, and promote more equitable access to health information—key priorities for public health and digital government service design. Moreover, the emphasis on agentic long-term memory, tailored information delivery, and proactive guidance bounded by human oversight aligns with digital government principles of transparency, trust, and responsible innovation. Together, this work illustrates how agentic AI can extend beyond efficiency-focused applications to support sustained caregiving work, offering a model for integrating AI into public-facing health services in ways that strengthen continuity, accountability, and caregiver empowerment.

6. Limitation

One limitation of this study lies in the relatively homogeneous demographic profile of the participant group, which primarily consisted of individuals with similar educational and professional backgrounds (e.g., AI, computer science, and health-related fields). While this composition may introduce potential bias in perspectives—particularly toward greater familiarity with technology and AI systems—it was intentionally aligned with the exploratory nature of this study. Specifically, participants with higher technical literacy were better positioned to critically engage with the system demonstration, reflect on its design, and articulate nuanced feedback regarding its functionality and potential use in longitudinal caregiving contexts. Nevertheless, future work would expand to include a more diverse caregiver population, particularly individuals with varying levels of digital literacy, socioeconomic backgrounds, and lived caregiving experience, to enhance the generalizability and ecological validity of the findings.

7. Conclusion

This paper presents an exploratory investigation of an agentic, memory augmented conversational AI system designed to support family caregivers in long term healthcare contexts. Through a scenario based qualitative study, we examined how such a system might help caregivers organize information, maintain continuity across appointments, and prepare for evolving care needs over time. The findings suggest that participants valued

the system's ability to integrate long term memory, provide structured and proactive guidance, and support sense making under conditions of cognitive and emotional strain, while also emphasizing the importance of personalization, digestibility, and human oversight. In sum, this work highlights the potential of agentic AI to function as a supportive, longitudinal companion in complex care journeys and highlights the need for future research to further evaluate these systems in real world settings, explore mechanisms for clinical validation, and refine AI agent design that balance continuity, empathy, and responsible use in healthcare environments.

8. Contributor Statement

Gargi Nandy: Conceptualization, Methodology, Investigation, Data Curation, Formal Analysis, Writing – Original Draft, Writing – Review and Editing.

Chun-Hua Tsai: Supervision, Conceptualization, Methodology, Validation, Writing – Review and Editing.

9. Acknowledgment

We would like to thank all study participants for their time, thoughtful feedback, and valuable reflections. Their perspectives helped inform the design implications presented in this work.

10. Conflict of Interest

The authors declare that there are no conflicts of interest related to this work.

11. Use of AI

Generative AI tools were used to support language refinement, grammar checking, and overall clarity during the writing and revision process. All research design decisions, data analysis, interpretation of findings, and final conclusions were made by the authors.

References

- Acharya, D. B., Kuppan, K., and Divya, B. (2025). Agentic ai: Autonomous intelligence for complex goals—a comprehensive survey. *IEEE Access*.
- Archibald, M. M., Ambagtsheer, R. C., Casey, M. G., and Lawless, M. (2019). Using zoom videoconferencing for qualitative data collection: perceptions and experiences of researchers and participants. *International journal of qualitative methods*, 18:1609406919874596.
- Braun, V. and Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative research in psychology*, 3(2):77–101.
- Debillon, T., Bednarek, N., Ego, A., and Elie, L. W. G. A. P.-Y. B. O. C. M.-L. D. C. F. C. G. P. K. G. M. S. P. J. S. L. S. (2018). Lytopenal: long term outcome of neonatal hypoxic encephalopathy in the era of neuroprotective treatment with hypothermia: a french population-based cohort. *BMC pediatrics*, 18(1):255.
- Fui-Hoon Nah, F., Zheng, R., Cai, J., Siau, K., and Chen, L. (2023). Generative ai and chatgpt: Applications, challenges, and ai-human collaboration.
- Hu, X., Bell, R. A., Kravitz, R. L., and Orrange, S. (2012). The prepared patient: information seeking of online support group members before their medical appointments. *Journal of health communication*, 17(8):960–978.
- Huntingford, S. L., Boyd, S. M., McIntyre, S. J., Goldsmith, S. C., Hunt, R. W., and Badawi, N. (2024). Long-term outcomes following hypoxic ischemic encephalopathy. *Clinics in Perinatology*, 51(3):683–709.
- Khalil, R. A., Ahmad, K., and Ali, H. (2026). Redefining elderly care with agentic ai: challenges and opportunities. *IEEE Open Journal of the Computer Society*.
- Kiiski, A., Luoma, E., Airaksinen, M., Pohjanoksa-Mäntylä, M., Desselle, S., and Kivelä, S.-L. (2025). Caregivers' challenges in engaging with the health system to optimise medication management of older care recipients: a qualitative study including home visits. *BMJ open*, 15(6):e093122.
- Knight, D. R., Aakre, C. A., Anstine, C. V., Munipalli, B., Biazar, P., Mitri, G., Valery, J. R., Brigham, T., Niazi, S. K., Perlman, A. I., et al. (2023). Artificial intelligence for patient scheduling in the real-world health care setting: A metanarrative review. *Health Policy and Technology*, 12(4):100824.

-
- Korf, J. M., McCullough, L. D., and Caretti, V. (2023). A narrative review on treatment strategies for neonatal hypoxic ischemic encephalopathy. *Translational Pediatrics*, 12(8):1552.
- Lee, B. L. and Glass, H. C. (2021). Cognitive outcomes in late childhood and adolescence of neonatal hypoxic-ischemic encephalopathy. *Clinical and experimental pediatrics*, 64(12):608.
- Lu, Q. and Schulz, P. J. (2024). Physician perspectives on internet-informed patients: systematic review. *Journal of Medical Internet Research*, 26:e47620.
- Mirzac, V., Thallemer, A. M., Yu, M.-C., Chen, Y.-T., and Tsai, C.-H. (2025). Evaluating the use of ai agents in teaching doctor-patient communication skills. In *International Conference on Disruptive Technologies, Tech Ethics and Artificial Intelligence*, pages 60–72. Springer.
- Nandy, G., Gupta, S., Mohammad Afzali, F., Peeples, E., Pilon, B., and Tsai, C.-H. (2025). Balancing health information-seeking through retrieval-augmented generation-based llm chatbot. In *Adjunct Proceedings of the 33rd ACM Conference on User Modeling, Adaptation and Personalization*, pages 249–254.
- Pierrat, V., Haouari, N., Liska, A., Thomas, D., Subtil, D., Truffert, P., et al. (2005). Prevalence, causes, and outcome at 2 years of age of newborn encephalopathy: population based study. *Archives of Disease in Childhood-Fetal and Neonatal Edition*, 90(3):F257–FF261.
- Piran, P., Khademi, Z., Tayari, N., and Mansouri, N. (2017). Caregiving burden of children with chronic diseases. *Electronic physician*, 9(9):5380.
- Raap, E., Weille, K. L., and Dedding, C. (2025). “i am the mother of the kind of child you dread having”: Experiences of living with chronic sorrow among parents with a disabled child. *Journal of health psychology*, 30(9):2158–2171.
- Rivero, E. (2025). Pediatric chronic disease prevalence has risen to nearly 30
- Rosson, M. B. and Carroll, J. M. (2007). Scenario-based design. In *The human-computer interaction handbook*, pages 1067–1086. CRC Press.
- Schreglmman, M., Ground, A., Vollmer, B., and Johnson, M. J. (2020). Systematic review: long-term cognitive and behavioural outcomes of neonatal hypoxic-ischaemic encephalopathy in children without cerebral palsy. *Acta Paediatrica*, 109(1):20–30.
- Schurgin, M., Schlager, M., Vardoulakis, L., Pina, L. R., and Wilcox, L. (2021). Isolation in coordination: Challenges of caregivers in the usa. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–14.
- Songire, D. S., Sharath, H., and HV, S. (2024). Tailored rehabilitation protocol on hypoxic-ischemic encephalopathy: A case report. *Cureus*, 16(9).
- Sumers, T., Yao, S., Narasimhan, K. R., and Griffiths, T. L. (2023). Cognitive architectures for language agents. *Transactions on Machine Learning Research*.
- Tennant, R., Allana, S., Mercer, K., and Burns, C. M. (2022). Exploring the experiences of family caregivers of children with special health care needs to inform the design of digital health systems: formative qualitative study. *JMIR Formative Research*, 6(1):e28895.
- Wisk, L. E. and Sharma, N. (2025). Prevalence and trends in pediatric-onset chronic conditions in the united states, 1999–2018. *Academic Pediatrics*, 25(4):102810.
- Wood, K. V., Frings, D., Flood, C., and Thomas, N. (2025). Artificial intelligence machine learning-driven outpatient appointment management: A qualitative study on acceptability. *Digital health*, 11:20552076251321016.
- Zheng, X., Li, Z., Gui, X., and Luo, Y. (2025). Customizing emotional support: How do individuals construct and interact with llm-powered chatbots. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, pages 1–20.