MEDINFO 2025 — Healthcare Smart × Medicine Deep
M.S. Househ et al. (Eds.)

doi:10.3233/SHTI250919

Advancing Patient-AI Conversations with Recursive Learning Memory Using Relational Frame Theory

Rishu GOYAL, Yifei DONG, Kun YU[1] and Fang CHEN
University of Technology Sydney, Australia
ORCiD ID: Rishu Goyal https://orcid.org/0009-0002-7007-119X, Yifei Dong https://orcid.org/0000-0002-7995-9657, Kun Yu[1] https://orcid.org/0000-0001-5138-6749, Fang Chen https://orcid.org/0000-0003-4971-8729

Abstract. The future of Artificial General Intelligence (AGI) in bedside care relies on integrating human psychological principles to foster advanced cognitive abilities such as reasoning, problem-solving, and spatial awareness. Key features like recursive learning memory and human-like context awareness allow AI systems to continuously learn and recall patient information in a relational manner. To achieve this, we propose an AI framework based on Relational Frame Theory (RFT), which organizes patient data to reflect human relational patterns. This approach enables iterative and context-sensitive information retrieval through a dynamic knowledge graph. Our evaluations demonstrate that this method enhances patient interactions, offering deeper, more personalized engagement that surpasses traditional Retrieval Augmented Generation in its ability to emulate nuanced, human-like understanding.

Keywords. Bedside AI, Patient Interactions, Relational Frame Theory, Knowledge Graph

1. Introduction

In medical practice, maintaining coherent and meaningful conversations with patients over time is crucial for providing effective care. Leveraging advanced AI techniques to support this process can ensure consistent care while significantly reducing the workload for healthcare staff. Traditionally, a general challenge for AI models is to process and understand extensive context in text, but the advent of large language models (LLMs) like GPT-3 and GPT-4 has transformed human-computer interactions, moving us closer to achieving Artificial General Intelligence (AGI). For AGI to be truly impactful in medical settings, however, it must extend beyond language comprehension to include the ability to store and recall patient-specific information in a manner that mirrors human memory and understanding [1]. Therefore, this study aims to design an AI tool that accommodates these needs and provides customized conversations to the patients.

In recent advancements in AGI, Retrieval Augmented Generation (RAG) [2] has been employed to store and retrieve vast amounts of semantically relevant medical information. While RAG excels at searching extensive datasets, it struggles in the context of medical conversations, where much of the essential information is relational and not explicitly documented. For instance, if an AI system notes that a patient has high blood pressure and recognizes its connection to heart disease, it should be able to infer the associated risk. Developing this ability to identify and generate relational connections is

[1] Corresponding Author: Kun YU, email: kun.yu@uts.edu.au

vital for improving AI-driven interactions and providing more meaningful and insightful healthcare support.

Relational Frame Theory (RFT) offers a psychological framework for understanding how humans relate concepts through language and context [3]. Building on this framework, this paper introduces a method that utilizes graph databases to model complex relationships within medical data. By continuously learning and adapting through patient interactions, this approach enables the AI to generate more detailed and contextually relevant responses, facilitating deeper, more human-like engagement. This methodology has the potential to outperform traditional RAG, allowing AI agents to develop meaningful relationships with patients and significantly enhancing communication and care outcomes in medical environments.

2. Methods

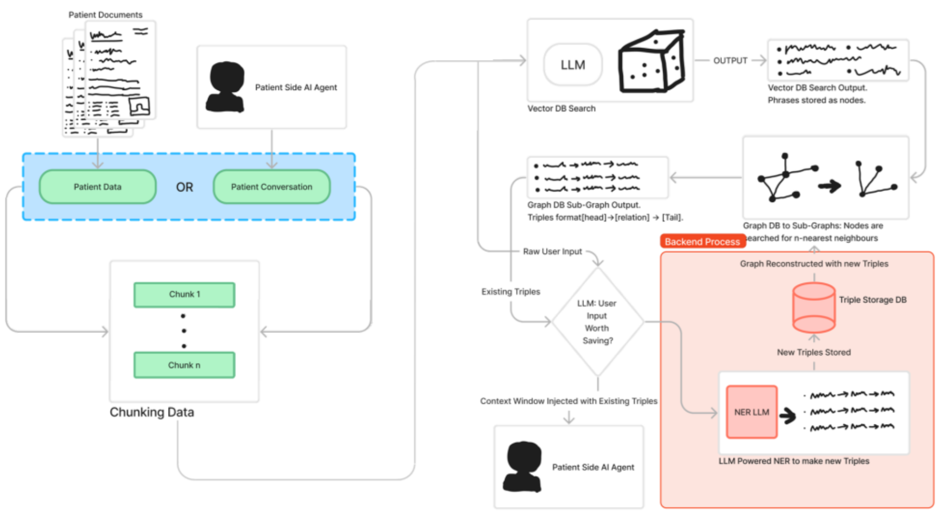

Figure 1. Pipeline for patient data processing involving data chunking/ontology creation, knowledge graph database generation and update, and patient conversational data integration.

2.1. Ontology Creation and Patient Data Injection

The ontology serves as the foundation of the patient's personal knowledge graph, organizing patient-specific nodes into categories to enable more effective inferencing using Graph Query Language (GQL) and standard graph algorithms, such as community detection and pathfinding. This ontology is designed to encompass the most common topics the patient bedside assistant is expected to address. It is synthetically constructed in a triple format (Head/Relation/Tail, as the example below) using GPT-4 to create a

Diabetes → DIAGNOSED_BY→ HbA1c Blood Test Lisinopril→ PRESCRIBED_FOR→ Hypertension Cardiac Catheterization→ PERFORMED_BY → Interventional Cardiologist.

highly interconnected graph. The prompt for generating this structure explicitly defines nodes as actionable nouns and relationships as adjectives or verbs, ensuring clarity and precision in the graph's design.

Patient data, including medical history, appointments, and personal demographics, is used to generate a foundational skeleton ontology in bulk, allowing personalization for each patient. A knowledge graph is then constructed using a base LLM, which processes text snippets into node/relation/node triples. These triples are concatenated and stored in a relational database, which is subsequently used to update the knowledge graph with relevant additions.

When building the knowledge graph, it is essential to ensure it grows in a branching structure with robust interconnected relationships. Given the use of an LLM, addressing issues like hallucinations and duplicate nodes is critical. To prevent redundancy, all existing nodes are stored in a vector space, enabling semantic retrieval of similar nodes. The LLM adds new nodes only if they can establish meaningful connections with the retrieved nodes, ensuring the graph remains interconnected, accurate, and unique.

2.2. Enriching Graph with Conversational Data

Once a comprehensive and detailed knowledge graph is constructed, the AI agent can engage in conversations with patients, using the chat history to continuously perform Named Entity Recognition (NER), which identifies and categorizes entities within the conversation. The knowledge graph evolves alongside the patient, incorporating their thoughts, feelings, emotions, desires, and experiences. After each conversation, the system checks for new information that needs to be saved or updated.

> Patient: Hey! I've started to take some **anti-inflammatory medication** for my **Arthritis**.
> AI: Hey Alice! Looks like that pain is getting worse huh. I'll keep them meds in mind!

2.3. Searching Knowledge Graph

For each patient response, relevant topic nodes are identified using Retrieval Augmented Generation (RAG). Once the relevant nodes are found, the *n*-nearest nodes within the knowledge graph are located, creating sub-graphs that capture the contextual relationships between the topics. For example, if the conversation shifts to the question, "What can I cook today?", the most semantically similar nodes might include "Food Allergies," "Low-Glycemic Diet," and "High-Fiber Diet." Based on this, the bedside assistant can recommend recipes tailored to the patient's prescribed dietary plan. It will then explore neighboring nodes connected by relationships, which may include conditions such as "Diabetes" or "Constipation." These contextual triples are

> User: Hey! What can I cook today?
> *Extracted Triples:*
> Diabetes → DIET_REQUIRES → Low-Glycaemic Food
> High-Fiber Diet →CONSIDERING → Constipation.
> Foods →DESPISED →Beans.
> AI: Hey Alice! Now we want to stick with a low-glycaemic, high-fibre diet to control your diabetes and help relieve that pesky constipation! And don't worry, we'll stay away from the beans. How about a heart-warming spiced Lentil Soup?

incorporated into the LLM's context window before generating a response. As a result, the assistant can provide highly personalized dietary recommendations that align with the patient's specific diagnosis, ensuring the advice is both relevant and customized to their health needs.

2.4. Methodology Assessment

The algorithm is accessed in a long conversational agent domain as a typical testcase similar to MemGPT [4]. Specifically, we adapted a human curated Objective Structured Clinical Exam (OSCE) case with 8 synthetically generated conversation-opener question and answer pairs. It is a real clinical case scenario document for medical students [5]. It encompasses everything in a patient's diagnosis from clinical inspections, test results and medical history, to final diagnosis, and management steps. This is to evaluate an agent's ability to craft engaging messages to the user that draws from knowledge accumulated in prior conversations. It will reveal how RFT performs in comparison with other AI agents in patient-AI conversations while also keeping the patient's personalization in mind. This experiment is similar to the Conversation Opener Task carried out by MemGPT [4].

3. Results

The comparison was conducted between GPT-4-o+RAG (File Search function in-built Open AI), Cohere's Command R+ with Grounding (Advanced RAG), and RFT+GPT-4o. System Prompt stayed consistent with all three. The ROUGE-L values were calculated, and by nature it's a matrix corresponding to the ratio of overlapping words between reference summary and candidate summary [6].

ROUGE-L (R)	***RFT AI***	***GPT4-o + file-search***	***Command R+***
Q1	**0.644**	0.534	0.5249
Q2	0.5901	**0.592**	0.486
Q3	0.5892	**0.6395**	0.5729
Q4	**0.621**	0.5623	0.4813
Q5	**0.5596**	0.4672	0.4870
Q6	**0.632**	0.5463	0.4755
Q7	**0.6158**	0.4886	0.551
Q8	**0.597**	0.5249	0.5897

Table 1. Recall rate of retrieved information from 8 *q-a* pairs.

4. Discussion

The results indicate that for most questions, Relational Frame Theory (RFT) AI outperformed Cohere's Command R+ and GPT-4-o with file search, demonstrating its ability to extract and synthesize knowledge from various sources within a clinical case document. Although the proposed method did not drastically outperform its competitors, with further parameter adjustments, even better results could be realized. Notably, GPT-4-o with file search placed nearly all of its text into the context window, with a token count exceeding 5000—equivalent to approximately 3000 words from a single input. In contrast, our method used only 1000 tokens yet still achieved good performance.

AI companions for patient care should be designed to accompany patients throughout their entire medical journey, from initial hospital consultations to in-home nursing and long-term post-discharge monitoring. This highly personalized assistant needs access to an ever-expanding knowledge base and the ability to intuitively understand and respond to the patient's thoughts, emotions, and experiences as they evolve. Two critical criteria must be met: (1) **Personalized engagement**—the AI's interactions must align with the patient's diagnosis, medical history, past conversations, and preferences, which are derived from previous interactions; and (2) **Indirect relationships**—the AI should be able to draw on information from past conversations to form responses that reflect deeper, implicit connections. The examination suggests that the proposed RFT AI solution meets both criteria, making it well-suited for bedside care.

Beyond this study, our future work will explore additional key considerations for AI, including ethics, fairness, long-term adaptability, patient privacy protection, and seamless integration with healthcare systems. Working with an AI bedside assistant could be new for healthcare providers, requiring tailored training to enhance adoption and trust in AI technology.

5. Conclusion

In this study, we introduced an AI agent system to introduce recursive learning memory, by building a Relational Framing for AI patient-bedside assistants. It is able to provide more human-like interactions. The RFT inspired approach has been evaluated alongside mainstream RAG applications, and solving the problem with very little prompting. It maintained personal engagement with the patient, post-diagnosis by knowing lots of facts about the patient's medical case.

Introduction of targeted retrieval from KGs and unlocking the implementation of GQL/ algorithms infrastructure to graph DBs in memory storage can have massive effect on making super-long-term memory for AI-human interactions through hyper-personalization. By bridging phycological concepts like RFT, it's likely to build AI agents that think and act like humans.

References

[1] Liu, N.F., Lin, K., Hewitt, J., Paranjape, A., Bevilacqua, M., Petroni, F. and Liang, P., 2024. Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*, *12*, pp.157-173.

[2] Lewis P, Perez E, Piktus A, Petroni F, Karpukhin V, Goyal N, Küttler H, Lewis M, Yih WT, Rocktäschel T, Riedel S. Retrieval-augmented generation for knowledge-intensive nlp tasks. Advances in Neural Information Processing Systems. 2020;33:9459-74.

[3] Blackledge, J.T., 2003. An introduction to relational frame theory: Basics and applications. *The Behavior Analyst Today*, *3*(4), p.421.

[4] Packer, C., Wooders, S., Lin, K., Fang, V., Patil, S.G., Stoica, I. and Gonzalez, J.E., 2023. Memgpt: Towards LLMs as operating systems. *arXiv preprint arXiv:2310.08560*.

[5] McGrath, A. (2023). Cough, Chest Pain and Fever – OSCE Case. https://geekymedics.com/cough-chest-pain-and-fever-osce-case/ (accessed Nov 30, 2024).

[6] Ganesan, K., 2017. An intro to ROUGE, and how to use it to evaluate summaries https://www.freecodecamp.org/news/what-is-rouge-and-how-it-works-for-evaluation-of-summaries-e059fb8ac840/ (accessed May 27, 2024).