Integrating Conversational Entities and Dialogue Histories with Knowledge Graphs and Generative AI

Graham Wilcock
CDM Interact and
University of Helsinki
Helsinki, Finland
graham.wilcock@cdminteract.com

Kristiina Jokinen
AI Research Center, National Institute of
Advanced Industrial Science and Technology
Tokyo, Japan
kristiina.jokinen@aist.go.jp

Abstract

Existing methods for storing dialogue history and for tracking mentioned entities in spoken dialogues usually handle these tasks separately. Recent advances in knowledge graphs and generative AI make it possible to integrate them in a framework with a uniform representation for dialogue management. This may help to build more natural and grounded dialogue models that can reduce misunderstanding and lead to more reliable dialogue-based interactions with AI agents. The paper describes ongoing work on this approach.

1 Introduction

Recognising conversational entities and maintaining dialogue history are two crucial tasks in natural language understanding for dialogue management. They help to maintain dialogue context and support coherent continuation of the dialogue. The two tasks have traditionally been clearly separated.

Recent advances in knowledge graphs and generative AI suggest possibilities for integrating them into a model that not only facilitates the dialogue processing with knowledge that focuses on a particular domain and contributes to the relevance and reliability of the information, but simultaneously serves as a uniform representation for storing the dialogue history.

This paper presents ongoing work on a prototype design of dialogue history graphs. While domain knowledge graphs store relatively static, long-term information about well-established domain entities and relationships between them, we use rapidly-changing, short-term graphs to represent dialogue history, including all mentioned conversational entities even if they do not necessarily correspond to known entities in the domain graph.

We use LLMs to split dialogue turns into smaller semantic units and to extract conversational entities. We present examples from recent results, with the aim that this approach will improve conversational grounding, reduce misunderstandings, and result in more reliable spoken interaction with AI agents.

The paper is structured as follows. Section 2 surveys related work on conversational entities and dialogue histories. Section 3 describes our work on representing the structure, content and history of dialogues in knowledge graphs. We use an example dialogue, shown in Table 1, to illustrate incremental generation of a dialogue history graph. Evaluation is briefly discussed in Section 4, with reference to results in the Appendices. Section 5 concludes.

2 Related Work

This section surveys how conversational entities and dialogue histories are handled by existing conversational AI frameworks. Entity recognition helps to build semantic representations of user intents. Dialogue history records which entities have been talked about, to maintain coherence of the dialogue as it progresses towards the goal.

2.1 Entity tracking and dialogue states

Traditionally, conversational entities have been tracked in statistical dialogue systems with the help of dialogue frames or states, which record the important concepts and entities needed to carry out task-based dialogues successfully. For instance, RavenClaw (Bohus and Rudnicky, 2009), which was used in the Let's Go system (Raux et al., 2005), effectively separates domain-specific task logic from the domain-independent dialogue control, and the series of the Dialogue State Tracking Challenge (Williams et al., 2013; Yoshino et al., 2024) aimed at efficient and correct evaluation of the dialogue progress via the dialogue entities being correctly filled in during the dialogue.

PyDial (Ultes et al., 2017) represents dialogue state with a set of domain-specific slots that can be switched to different slots for different domains.

Proceedings of the 15th International Workshop on Spoken Dialogue Systems Technology, pages 290–298
May 27–30, 2025.

Human:	What is hypertension?
Robot:	*Hypertension (HTN) is a medical condition characterized by consistently elevated blood pressure, defined as a systolic blood pressure (SBP) of 140 mm Hg or higher and/or a diastolic blood pressure (DBP) of 90 mm Hg or higher. It remains a significant public health issue, affecting millions of adults in the United States.*
Human:	What is it caused by?
Robot:	*Hypertension can be caused by a variety of factors, including genetic predisposition, lifestyle choices such as poor diet, lack of physical activity, obesity, and excessive alcohol consumption. Other contributing factors may include chronic stress, certain medical conditions, and the use of specific medications.*
Human:	Is exercise helpful?
Robot:	*Yes, exercise is helpful in managing hypertension, as regular moderate-intensity endurance exercise can lower blood pressure both acutely and chronically. It is recommended for individuals with hypertension to engage in regular physical activity to maximize health benefits.*
Human:	What is moderate-intensity endurance exercise?
Robot:	*Moderate-intensity endurance exercise is physical activity that raises your heart rate and breathing but still allows you to talk comfortably, typically rated as 5 to 6 on a scale of 0 to 10 for exertion. Examples include brisk walking, cycling at a leisurely pace, or swimming at a moderate speed.*
Human:	How can I reduce my salt intake?
Robot:	*To reduce your salt intake, you can start by cooking at home using fresh ingredients, avoiding processed foods, and reading food labels to choose low-sodium options. Additionally, using herbs and spices for flavor instead of salt can help you cut back on sodium.*
Human:	What are the symptoms of diabetes?
Robot:	*Common symptoms of diabetes include increased thirst, frequent urination, extreme fatigue, blurred vision, and slow-healing sores or frequent infections. Other symptoms may include unexplained weight loss and tingling or numbness in the hands or feet.*
Human:	Is exercise helpful?
Robot:	*Yes, exercise is helpful for managing diabetes, as it improves blood glucose control, enhances insulin sensitivity, and can lower the risk of complications. Regular physical activity is essential for optimal health in individuals with diabetes.*

Table 1: A dialogue example in medical domain.

More flexibility was introduced for multi-domain dialogues (Budzianowski et al., 2018) by allowing frames to be switched on recognition of domain-specific entities, enabling flexible changes from one topic to another. However, these multi-domain dialogue models cannot handle structures such as relations between entities in different domains. To address this limitation, PyDial introduced a new Conversational Entity Dialogue Model that models both objects and relations in a conversational world. For further details, see (Ultes et al., 2018).

In the Empathic project (Vázquez et al., 2023) a modular dialogue management framework was developed using templates, dialogue acts, NLU entities and LLMs. Dialogue flow could be represented as a graph, but graphs were not used as a uniform format in the system.

AMR (Abstract Meaning Representation) is a labled graph -based representation for the sentence meaning, used in semantic annotation banks (Banarescu et al., 2013). Dialogue-AMR enriches AMR with dialogue acts, and with tense and aspect features for dialogues (Bonial et al., 2020). Standard AMR is also explored by (Bai et al., 2021) as a semantic representation for dialogue histories in order to better represent conversations.

2.2 Dialogue history in Rasa and LangChain

Rasa conversational AI (Bocklisch et al., 2017) offers a range of options for storing conversation history in *tracker stores*. The dialogue utterances are stored as simple text strings, with no annotation of the conversational entities that they mention.

The default `InMemoryTrackerStore` loses the history when the Rasa server is stopped or restarted. `SQLTrackerStore` provides persistent storage of histories in SQL databases. `RedisTrackerStore` is a fast in-memory store that can optionally persist the history. There are also NoSQL database stores.

For entity recognition, Rasa has transformer-based components that can be rapidly trained from small corpora of domain-specific examples. DIET Dual Intent and Entity Transformer (Bunk et al., 2020) combines recognition of user intents and mentioned entities. Recognised entities are inserted into domain-specific slots, but the entities in the slots are not recorded in the tracker store history. The slots are repeatedly overwritten by new entities as the dialogue progresses.

LangChain has a similar set of storage options for *chat message histories* (LangChain, 2023). In-memory history is lost when the session ends. SQL databases provide persistent storage. Redis is an in-memory store that can persist the data, and there are also NoSQL database stores.

Chat message histories store utterances as strings with no annotation of the conversational entities mentioned. However, LangChain has an alternative Conversation Entity Memory, that uses an LLM to extract entities and information about them.

LangChain can also summarize dialogues using an LLM, in order to reduce the amount of text to be stored. Dialogue history can now also be saved in LangGraph Memory (LangChain, 2025).

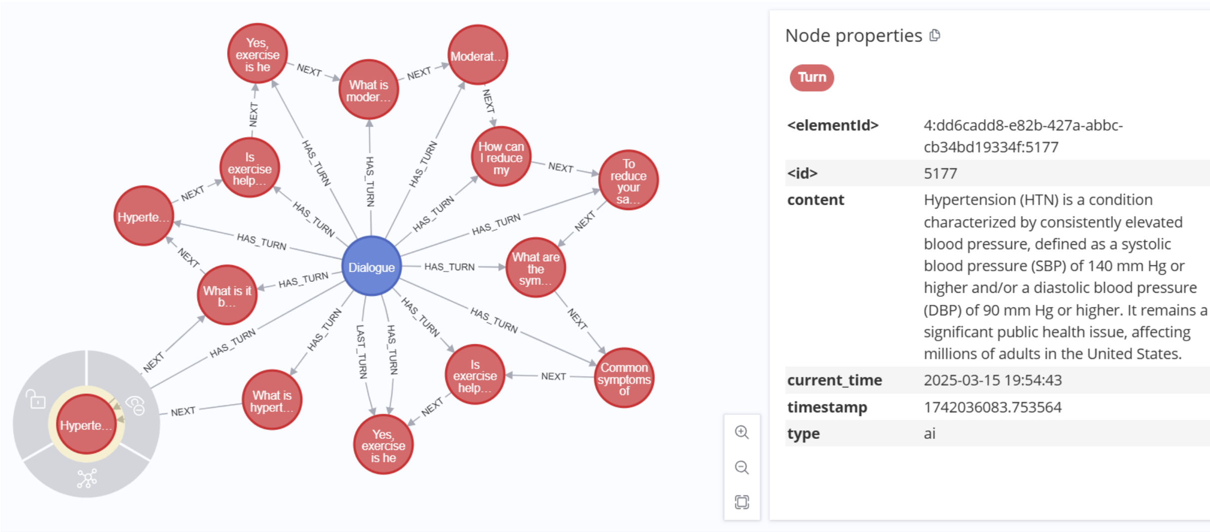

Figure 1: All the dialogue turns from the dialogue in Table 1 in a dialogue history graph. One turn by the AI agent is highlighted, showing its content (about hypertension), timestamp and type (Human or AI).

With Neo4jChatMessageHistory, LangChain can also store chat message history in Neo4j graph databases. Messages are nodes in the graph, linked in sequence by NEXT relationships. A Session node maintains an updated LAST_MESSAGE link to show where to insert the next message.

2.3 Dialogue graphs in GraphDial

The GraphDial project saw the benefits of encoding dialogue state as a flexible knowledge graph that can be updated by transformations expressed in a graph query language (Walker et al., 2022a). The project uses the same labeled property graph format as Neo4j, and the same Cypher query language, but with an in-memory graph database.

GraphDial was tested successfully in a case study, in which a Pepper robot is a receptionist with domain-specific tasks (Walker et al., 2022b, 2023). Rasa NLU was used for both domain-specific entity and domain-specific intent recognition.

Our work differs from GraphDial in two aspects because we combine the benefits of using graphs with the new possibilities of generative AI:

- We use LLMs to do entity recognition across open domains, and
- We use LLMs to generate responses that take account of user intents without needing to do intent classification as a separate task.

3 Dialogue History Graphs

In earlier work we used knowledge graphs to store domain information for spoken dialogue systems, and showed that more cooperative responses can be generated by adding semantic metadata such as taxonomies extracted from WikiData to the domain graphs (Wilcock and Jokinen, 2022a,b).

More recently we proposed a new role for knowledge graphs, aiming to reduce dialogue errors by better support for conversational grounding. This approach (Jokinen and Wilcock, 2025) uses both domain knowledge graphs and dialogue history graphs. The aim is to explore how to improve conversational grounding by constructing shared knowledge via entity linking.

Table 1 is a transcript of a human-robot dialogue about hypertension, exercise and diabetes that uses a domain graph generated from PDF documents by LLMGraphTransformer, as described by Jokinen and Wilcock (2025). The robot responses were generated by an LLM using vector-based RAG retrieval from the domain knowledge graph.[1]

In this paper, the dialogue serves as an example to illustrate incremental generation of a dialogue history graph. All the dialogue turns are visualised in Figure 1. We have added timestamps to each turn, which are linked in sequence by NEXT links. The first AI response is highlighted and LAST_TURN shows where to add the next turn.

Dialogue history has not usually been stored in

[1] A similar dialogue with a Furhat robot is on YouTube (https://www.youtube.com/watch?v=vs4Y5jjoIqM).

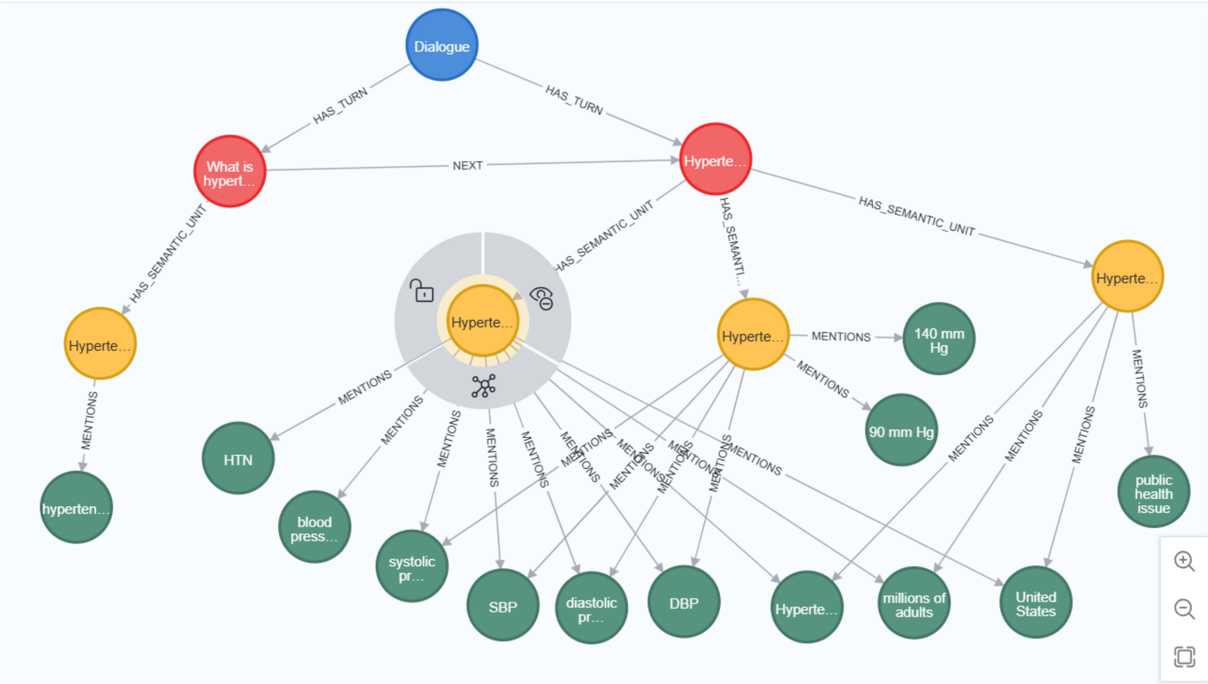

Figure 2: The four layers (Dialogue, Turns, Semantic Units, Conversational Entities) in a dialogue history graph.

knowledge graphs. Nevertheless, there are potential benefits from using graphs to store dialogue history, as already mentioned in Section 2. These include flexibility of entity tracking, domain switching, and the use of graph formalism for search and update of knowledge. Another benefit is to provide a flexible dialogue memory for cooperative dialogues and enable conversational grounding to help building of a shared context within which to enhance mutual understanding, clarify vague utterances, and resolve misunderstandings.

Recently there have been significant advances in graph-based retrieval from documents that go beyond basic RAG retrieval. One example is Graph RAG for local and global summarization over large document collections (Edge et al., 2024) and another advance is GraphReader (Li et al., 2024) for retrieval from long documents such as books. GraphReader represents documents using graphs with a generic schema. Document nodes are linked to text chunk nodes. Text chunk nodes are linked to "atomic fact" nodes. Atomic fact nodes are linked to "key element" nodes. Bratanic (2024) suggests that using this generic schema enables the approach to be applied to documents in any domain.

We propose a similar approach for representing the structure, content and history of dialogues, with the intention that a generic dialogue schema will be applicable to dialogues in any domain. This aligns with early dialogue research like RavenClaw (Bohus and Rudnicky, 2009) which emphasises the separation of task content and dialogue skills.

Figure 2 shows the four layers in the generic schema for a dialogue history graph. Dialogue nodes (blue) in the top layer are linked to turn nodes (red) in the second layer. Turn nodes are linked to semantic unit nodes (yellow) in the third layer. Semantic unit nodes are linked to conversational entity nodes (green) in the fourth layer.

Dialogue history graphs grow incrementally with each new dialogue turn. Figure 2 shows the state of the graph after the first human turn (*What is Hypertension?*) and the AI response. At this point in the dialogue there are only two turn nodes. The complete dialogue history graph for the example dialogue in Table 1 is shown in Figure 3.

When each turn is added to the dialogue history graph, its content is split by an LLM into semantic units and mentioned conversational entities. The prompts for these tasks are inspired by the prompts for extracting atomic facts and key elements in the GraphReader approach to document processing (Li et al., 2024; Bratanic, 2024).

Finally, conversational entities can be grounded to real-world entities in a domain knowledge graph, which represents the state of the world (conceptual structure of the world combined with immediate information from external sensors). Of course, con-

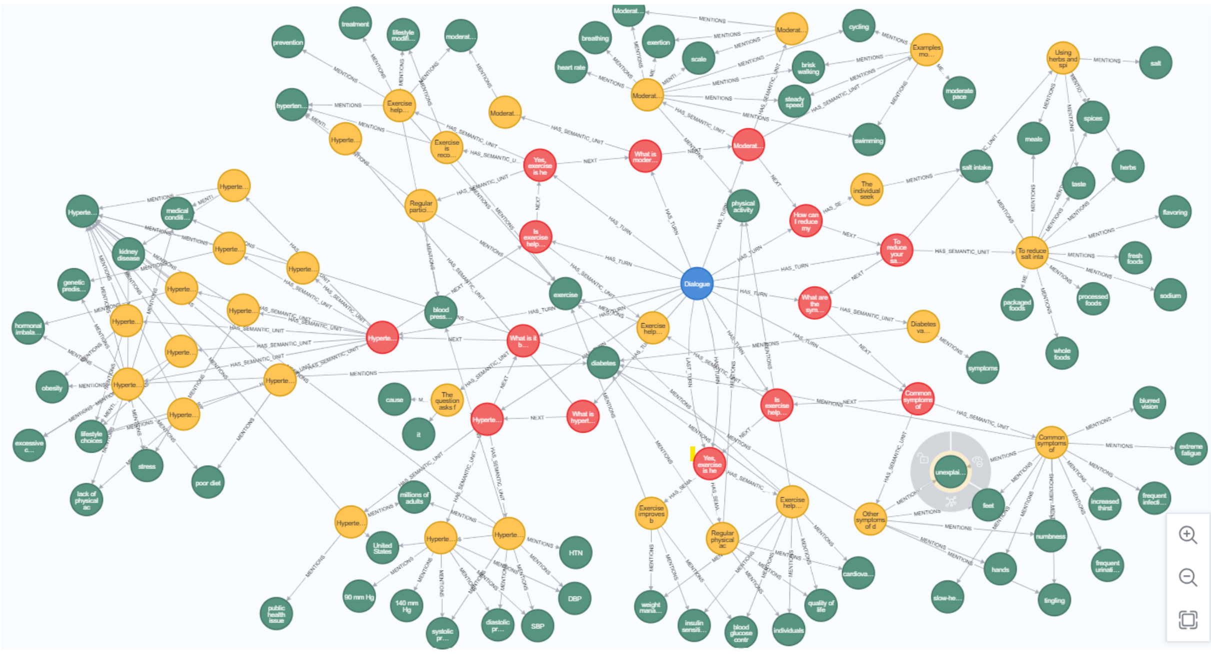

Figure 3: An overview of the dialogue history graph for the full dialogue in Table 1.

versational entities might not exist in the knowledge graph (it's fine to talk about unicorns for example, although they do not exist in the real world), but it is possible to build a knowledge graph and ontology for imaginary worlds within which such imaginary entities can be grounded.

4 Evaluation

This paper describes ongoing work for which user evaluation has not yet been possible. Results for a qualitative evaluation are shown in Appendices.

Appendix A lists all semantic units extracted from response turns. The turns are correctly split into semantic units, and the pronouns are correctly disambiguated. The semantic units correspond to conceptual propositions expressed in the surface turn, which can be further used as a basis for search, translation, or summarization.

Appendix B shows examples of the mentioned conversational entities and linked semantic units. The entities are in general correctly extracted from the dialogue turns that mention them, but they are not always linked to the correct semantic unit within the turn. This requires improvements to the prompts that guide the LLM in this task.

5 Conclusion and Future Work

Most existing approaches to dialogue histories and mentioned entities in spoken dialogue systems keep them separate. While acknowledging the different logics behind them, the paper describes an approach that integrates them based on recent advances in knowledge graphs and generative AI.

Such integration has several benefits for dialogue managent and we expect dialogue graphs also to facilitate research on conversational grounding, where even modest improvements could reduce misunderstandings and lead to more reliable spoken dialogue interactions with AI agents and robots. For instance, recent studies have investigated the use of knowledge graphs to bridge information gaps in dialogues with grounded exchanges (Schneider et al., 2024), and the capacity of LLMs to build shared knowledge by classifying grounding dialogue acts and by extracting mutually grounded information (Jokinen et al., 2024).

Future work will include improvements to the prompts to help the LLM to more carefully link conversational entities extracted from the turns to the specific semantic units that mention them.

Finally, it is crucial to enhance work in automatic building of knowledge graphs. As discussed in (Zhong et al., 2023), conversational entities can be used to link to entities in a domain graph, or to build new domain graphs, with verbs representing relationships. Given the potential of learning about the world through interaction, future research will also focus on how conversational grounding can facilitate acquiring knowledge by interaction and help to build effective knowledgebases.

References

Xuefeng Bai, Yulong Chen, Linfeng Song, and Yue Zhang. 2021. Semantic representation for dialogue modeling. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4430–4445, Online. Association for Computational Linguistics.

Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. 2013. Abstract Meaning Representation for sembanking. In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 178–186, Sofia, Bulgaria. Association for Computational Linguistics.

Tom Bocklisch, Joey Faulkner, Nick Pawlowski, and Alan Nichol. 2017. Rasa: Open source language understanding and dialogue management. ArXiv:1712.05181.

Dan Bohus and Alexander I. Rudnicky. 2009. The ravenclaw dialog management framework: Architecture and systems. *Comput. Speech Lang.*, 23(3):332–361.

Claire Bonial, Lucia Donatelli, Mitchell Abrams, Stephanie M. Lukin, Stephen Tratz, Matthew Marge, Ron Artstein, David Traum, and Clare Voss. 2020. Dialogue-AMR: Abstract Meaning Representation for dialogue. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 684–695, Marseille, France. European Language Resources Association.

Tomaz Bratanic. 2024. Implementing GraphReader with Neo4j and LangGraph. `https://github.com/tomasonjo/blogs/blob/master/graphreader/graphreader_langgraph.ipynb`.

Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gašić. 2018. MultiWOZ - a large-scale multi-domain Wizard-of-Oz dataset for task-oriented dialogue modelling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5016–5026, Brussels, Belgium. Association for Computational Linguistics.

Tanja Bunk, Daksh Varshneya, Vladimir Vlasov, and Alan Nichol. 2020. DIET: Lightweight language understanding for dialogue systems. ArXiv:2004.09936.

Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, and Jonathan Larson. 2024. From local to global: A Graph RAG approach to query-focused summarization. *Preprint*, arXiv:2404.16130.

Kristiina Jokinen, Phillip Schneider, and Taiga Mori. 2024. Towards harnessing large language models for comprehension of conversational grounding. In *14th International Workshop on Spoken Dialogue Systems Technology (IWSDS 2024)*, Sapporo, Japan.

Kristiina Jokinen and Graham Wilcock. 2025. Towards domain graphs and dialogue graphs for conversational grounding in HRI. In *Proceedings of the 2025 ACM/IEEE International Conference on Human-Robot Interaction*, HRI '25, page 1373–1377. IEEE Press.

LangChain. 2023. How to add message history. `https://python.langchain.com/v0.2/docs/how_to/message_history/`.

LangChain. 2025. Memory. `https://langchain-ai.github.io/langgraph/concepts/memory/`.

Shilong Li, Yancheng He, Hangyu Guo, Xingyuan Bu, Ge Bai, Jie Liu, Jiaheng Liu, Xingwei Qu, Yangguang Li, Wanli Ouyang, Wenbo Su, and Bo Zheng. 2024. GraphReader: Building graph-based agent to enhance long-context abilities of large language models. *Preprint*, arXiv:2406.14550.

Antoine Raux, Brian Langner, Dan Bohus, Alan W Black, and Maxine Eskenazi. 2005. Let's go public! taking a spoken dialog system to the real world. In *9th European Conference on Speech Communication and Technology, Lisbon, Portugal*.

Phillip Schneider, Nektarios Machner, Kristiina Jokinen, and Florian Matthes. 2024. Bridging information gaps in dialogues with grounded exchanges using knowledge graphs. In *Proceedings of the 25th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 110–120, Kyoto, Japan. Association for Computational Linguistics.

Stefan Ultes, Paweł Budzianowski, Iñigo Casanueva, Lina M. Rojas-Barahona, Bo-Hsiang Tseng, Yen-Chen Wu, Steve Young, and Milica Gašić. 2018. Addressing objects and their relations: The conversational entity dialogue model. In *Proceedings of the 19th Annual SIGdial Meeting on Discourse and Dialogue*, pages 273–283, Melbourne, Australia. Association for Computational Linguistics.

Stefan Ultes, Lina M. Rojas-Barahona, Pei-Hao Su, David Vandyke, Dongho Kim, Iñigo Casanueva, Paweł Budzianowski, Nikola Mrkšić, Tsung-Hsien Wen, Milica Gašić, and Steve Young. 2017. PyDial: A multi-domain statistical dialogue system toolkit. In *Proceedings of ACL 2017, System Demonstrations*, pages 73–78, Vancouver, Canada. Association for Computational Linguistics.

Alain Vázquez, Asier López-Zorrilla, Javier Mikel Olaso, and María Inés Torres. 2023. Dialogue management and language generation for a robust conversational virtual coach: Validation and user study. *Sensors*, 23(3):1423.

Nicholas Walker, Torbjørn Dahl, and Pierre Lison. 2022a. Dialogue management as graph transformations. In *Conversational AI for Natural Human-Centric Interaction*, pages 219–227, Singapore. Springer Nature Singapore.

Nicholas Walker, Stefan Ultes, and Pierre Lison. 2022b. GraphWOZ: Dialogue management with conversational knowledge graphs. ArXiv:2211.12852.

Nicholas Walker, Stefan Ultes, and Pierre Lison. 2023. A graph-to-text approach to knowledge-grounded response generation in human-robot interaction. ArXiv:2311.16137.

Graham Wilcock and Kristiina Jokinen. 2022a. Conversational AI and knowledge graphs for social robot interaction. In *ACM/IEEE International Conference on Human-Robot Interaction (HRI 2022)*, pages 1090–1094, Sapporo, Japan. Association for Computing Machinery.

Graham Wilcock and Kristiina Jokinen. 2022b. Cooperative and uncooperative behaviour in task-oriented dialogues with social robots. In *31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN 2022)*, pages 763–768, Napoli, Italy.

Jason Williams, Antoine Raux, Deepak Ramachandran, and Alan Black. 2013. The dialog state tracking challenge. In *Proceedings of the SIGDIAL 2013 Conference*, pages 404–413, Metz, France. Association for Computational Linguistics.

Koichiro Yoshino, Yun-Nung Chen, Paul Crook, Satwik Kottur, Jinchao Li, Behnam Hedayatnia, Seungwhan Moon, Zhengcong Fei, Zekang Li, Jinchao Zhang, Yang Feng, Jie Zhou, Seokhwan Kim, Yang Liu, Di Jin, Alexandros Papangelis, Karthik Gopalakrishnan, Dilek Hakkani-Tur, Babak Damavandi, and 9 others. 2024. Overview of the tenth dialog system technology challenge: Dstc10. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 32:765–778.

Lingfeng Zhong, Jia Wu, Qian Li, Hao Peng, and Xindong Wu. 2023. A comprehensive survey on automatic knowledge graph construction. *Preprint*, arXiv:2302.05019.

A Splitting Dialogue Turns into Semantic Units

Dialogue Turns	Semantic Units
Hypertension (HTN) is a medical condition characterized by consistently elevated blood pressure, defined as a systolic blood pressure (SBP) of 140 mm Hg or higher and/or a diastolic blood pressure (DBP) of 90 mm Hg or higher. It remains a significant public health issue, affecting millions of adults in the United States.	Hypertension (HTN) is a medical condition characterized by consistently elevated blood pressure. Hypertension is defined as a systolic blood pressure (SBP) of 140 mm Hg or higher and/or a diastolic blood pressure (DBP) of 90 mm Hg or higher. Hypertension remains a significant public health issue, affecting millions of adults in the United States.
Hypertension can be caused by a variety of factors, including genetic predisposition, lifestyle choices such as poor diet, lack of physical activity, obesity, and excessive alcohol consumption. Other contributing factors may include chronic stress, certain medical conditions, and the use of specific medications.	Hypertension can be caused by a variety of factors. Genetic predisposition can cause hypertension. Lifestyle choices can cause hypertension. A poor diet can cause hypertension. Lack of physical activity can cause hypertension. Obesity can cause hypertension. Excessive alcohol consumption can cause hypertension. Chronic stress can contribute to hypertension. Certain medical conditions can contribute to hypertension. The use of specific medications can contribute to hypertension.
Yes, exercise is helpful in managing hypertension, as regular moderate-intensity endurance exercise can lower blood pressure both acutely and chronically. It is recommended for individuals with hypertension to engage in regular physical activity to maximize health benefits.	Exercise is helpful in managing hypertension. Regular moderate-intensity endurance exercise can lower blood pressure both acutely and chronically. It is recommended for individuals with hypertension to engage in regular physical activity to maximize health benefits.
Moderate-intensity endurance exercise is physical activity that raises your heart rate and breathing but still allows you to talk comfortably, typically rated as 5 to 6 on a scale of 0 to 10 for exertion. Examples include brisk walking, cycling at a leisurely pace, or swimming at a moderate speed.	Moderate-intensity endurance exercise is physical activity that raises heart rate and breathing but still allows for comfortable talking. Moderate-intensity endurance exercise is typically rated as 5 to 6 on a scale of 0 to 10 for exertion. Examples of moderate-intensity endurance exercise include brisk walking, cycling at a leisurely pace, and swimming at a moderate speed.
To reduce your salt intake, you can start by cooking at home using fresh ingredients, avoiding processed foods, and reading food labels to choose low-sodium options. Additionally, using herbs and spices for flavor instead of salt can help you cut back on sodium.	To reduce salt intake, one can cook at home using fresh ingredients, avoid processed foods, and read food labels to choose low-sodium options. Using herbs and spices for flavor instead of salt can help cut back on sodium.
Common symptoms of diabetes include increased thirst, frequent urination, extreme fatigue, blurred vision, and slow-healing sores or frequent infections. Other symptoms may include unexplained weight loss and tingling or numbness in the hands or feet.	Common symptoms of diabetes include increased thirst, frequent urination, extreme fatigue, blurred vision, and slow-healing sores or frequent infections. Other symptoms of diabetes may include unexplained weight loss and tingling or numbness in the hands or feet.
Yes, exercise is helpful for managing diabetes, as it improves blood glucose control, enhances insulin sensitivity, and can lower the risk of complications. Regular physical activity is essential for optimal health in individuals with diabetes.	Exercise is helpful for managing diabetes. Exercise improves blood glucose control. Exercise enhances insulin sensitivity. Exercise can lower the risk of complications. Regular physical activity is essential for optimal health in individuals with diabetes.

Table 2: How dialogue turns from the dialogue in Table 1 are split into semantic units by the LLM.

B Examples of linking Conversational Entities to Semantic Units

Semantic Units	**Conversational Entities**
Using herbs and spices for flavor instead of salt can help cut back on sodium.	herbs spices flavor salt sodium
To reduce salt intake, one can start by cooking at home using fresh ingredients.	salt intake cooking home fresh ingredients processed foods food labels sodium low-sodium options herbs spices
Avoiding processed foods and reading food labels can help in choosing low-sodium options.	processed foods food labels low-sodium options
Moderate-intensity endurance exercise is physical activity that raises heart rate and breathing but still allows for comfortable talking.	Moderate-intensity endurance exercise heart rate breathing scale brisk walking exertion cycling swimming steady speed
Examples of moderate-intensity endurance exercise include brisk walking, cycling at a moderate pace, or swimming at a steady speed.	brisk walking cycling moderate pace swimming steady speed
Moderate-intensity endurance exercise is typically rated as 5 to 6 on a scale of 0 to 10 for exertion.	Moderate-intensity endurance exercise exertion scale 0 to 10 5 to 6

Table 3: How conversational entitiesfrom the dialogue in Table 1 are linked to semantic units by the LLM.