Nicholas Thomas Walker

Otto-Friedrich University of Bamberg
Gutenbergstraße 13
Bamberg
Germany

nicholas.walker@uni-bamberg.de
www.uni-bamberg.de/ds/team/walker/

1 Research interests

In dialogue research, I am especially interested in using knowledge bases and models of dialogue entities to enhance dialogue system output. My research focuses particularly on how graph-structured knowledge underlying a dialogue system can be created, transformed, and extended to ground dialogue system output in world knowledge. In addition, I am interested in how logical rule-based reasoning can play a useful supplementary role in improving dialogue systems' understanding of dialogue context. Although contemporary large language models exhibit ever more impressive abilities, the problem of model hallucinations and groundedness in dialogue context remains an outstanding challenge in many applications (Zhang et al., 2023) that may be addressed in part by approaches that improve dialogue system access to background knowledge (Dinan et al., 2018).

In particular, I am interested graph-based **dialogue management** for **task-oriented dialogue systems**, particularly the use of **dynamic knowledge-bases**. I am especially interested in representations of knowledge combining information about the world with dialogue or user-specific information, e.g. personal knowledge graphs (Balog and Kenter, 2019), and how dialogue-local information can be effectively combined with other sources of knowledge such as logical rule-based reasoning. In this sense, I am interested in approaches which combine contemporary LLMs with other sources of reasoning, in the vein of neurosymbolic AI (Garcez and Lamb, 2023).

1.1 Previous and Current Work

In my previous work, I have investigated graph-based entity-centric models of dialogue management in the form of *conversational entity graphs* for SDS in **Human-Robot Interaction** (HRI). Such graphs are dynamic knowledge graphs centered on dialogue entities, which can be viewed generally as distinct units of information that may be useful to the system. Physical entities in the world that are important in situated dialogue can be conveniently represented within a graph alongside dialogue-specific, abstract or 'virtual" entities (Ultes et al., 2018) such as calendar events or even conversational intents can also be represented in the graph, with relations describing where an event will take place or which entities an intent refers to. As the dialogue proceeds, the graph of the dialogue state can be updated to represent changes in the entities and relations underlying the dialogue.

My previous work in graph-based dialogue management included a combination of probabilistic rule-based logic programming and neural models. Using ProbLog (De Raedt et al., 2007), the dialogue state can be enriched and extended by applying rules expressing common-sense inferences that can subsequently be verbalized for use by a language model (Walker et al., 2023). Thus, this combination of LLMs with logical rules is a hybrid approach allowing for desired common-sense conclusions and relevant information to be provided explicitly to the LLM through verbalization in the dialogue context.

In current research, I am working towards further extension of methods using the conversational entity graph to enrich the graph state and retrieve relevant information for the dialogue system. Where my previous work has investigated retrieval and verbalization methods combined with common-sense rules, I am now investigating ways to extend previous methods to extract relevant information from larger level graph structures beyond triples, such as meta-paths and node neighborhoods. As datasets often lack labels indicating knowledge graph elements that are relevant for a given dialogue turn, I am developing methods of knowledge extraction that do not rely on large quantities of data for supervised learning.

1.2 Future Work

In my future work, I plan to investigate what knowledge representations are most effective in enabling dialogue systems to converse naturally with users. Building upon my previous work, I intend to investigate new models of dialogue management incorporating abstract knowledge about the dialogue such as user and system goals and their interaction with both turn-level information and the dialogue as a whole. The relation of dialogue entities to overarching goals and objectives of the system has long played a role in classic slot-filling methods of dialogue state-tracking, yet these structures provide only a partial view of the interplay between these elements of dialogue

The 20th Annual Meeting of the Young Researchers' Roundtable on Spoken Dialogue Systems, pages 46–47
September 16–17, 2024.

(Cohen, 2019). I am therefore interested in investigating methods which provide system decision making and output with stronger grounding in dialogue knowledge to provide more consistent and reliable responses. I believe contemporary LLMs provide an opportunity to combine the strong generative capabilities of these models with a more robust understanding of both user intentions, system purpose and goals, and situational knowledge that can enable more naturalistic and robust dialogue systems.

2 Spoken dialogue system (SDS) research

Within SDS research in general, I believe the latest LLMs that accommodate multimodal input will be a central area of investigation. The degree to which such models are able to reason over diverse sources of multimodal information appears likely to be an important topic (Wang et al., 2024), and it may be of use to integrate external modules to refine their input and output in order to assist in areas in which these models fall short. The use of external tools and integration of different modules with LLMs seems likely to be an important area of research in the future due to the difficulties in fine-tuning and retraining such models. It will also likely continue to be important for research to focus on language model explainability and deeper understanding of the transformer architecture.

3 Suggested topics for discussion

- Neurosymbolic AI: How can the latest LLMs (including multimodal models) interact with logic-based decision making? What role can classical rule-based or logical decision making play in the current era of extremely large language models?

- Intentionality: What ways might we impart more intentionality or illocutionary intent to dialogue systems? Would a system that has some facsimile of such intentionality be likely to perform better at tasks or behave more naturally?

- Adapting System Behavior: What approaches are most effective for adapting system behavior to a task or situation when fine-tuning is not an option? Is prompt optimization sufficient to desired induce system behavior in all contexts? If not, which areas remain a challenge?

References

Krisztian Balog and Tom Kenter. 2019. Personal knowledge graphs: A research agenda. In *Proceedings of the ACM SIGIR International Conference on the Theory of Information Retrieval (ICTIR)*.

Philip R Cohen. 2019. Foundations of collaborative task-oriented dialogue: what's in a slot? In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*. pages 198–209.

Luc De Raedt, Angelika Kimmig, and Hannu Toivonen. 2007. Problog: A probabilistic prolog and its application in link discovery. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence IJCAI-07*. Hyderabad, volume 7, pages 2462–2467.

Emily Dinan, Stephen Roller, Kurt Shuster, Angela Fan, Michael Auli, and Jason Weston. 2018. Wizard of wikipedia: Knowledge-powered conversational agents. In *Proceedings of the International Conference on Learning Representations (ICLR)*.

Artur d'Avila Garcez and Luis C Lamb. 2023. Neurosymbolic ai: The 3 rd wave. *Artificial Intelligence Review* pages 1–20.

Stefan Ultes, Paweł Budzianowski, Iñigo Casanueva, Lina M Rojas Barahona, Bo-Hsiang Tseng, Yen-Chen Wu, Steve Young, and Milica Gasic. 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.

Nicholas Thomas Walker, Stefan Ultes, and Pierre Lison. 2023. A retrieval-augmented neural response generation using logical reasoning and relevance scoring. In *Proceedings of the 27th Workshop on the Semantics and Pragmatics of Dialogue, Full Papers*.

Yiqi Wang, Wentao Chen, Xiaotian Han, Xudong Lin, Haiteng Zhao, Yongfei Liu, Bohan Zhai, Jianbo Yuan, Quanzeng You, and Hongxia Yang. 2024. Exploring the reasoning abilities of multimodal large language models (mllms): A comprehensive survey on emerging trends in multimodal reasoning. *arXiv preprint arXiv:2401.06805* .

Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, et al. 2023. Siren's song in the ai ocean: A survey on hallucination in large language models. *arXiv:2309.01219* https://arxiv.org/pdf/2309.01219.pdf.

Biographical sketch

Nick Walker is a researcher at Otto-Friedrich University of Bamberg. He completed his bachelor's degree in linguistics and master's degree in human language technology at the University in Arizona, and will be defending his PhD at the University of Oslo in October of 2024.