Laetitia Hilgendorf

Natural Language Generation and Dialogue Systems Group
Otto-Friedrich-University of Bamberg
96047 Bamberg
Bavaria, Germany

laetitia-mina.hilgendorf@stud.uni-bamberg.de

1 Research interests

My research focuses on developing methods to **guide response generation in LLM-based dialogue systems**. One approach involves learning dialogue policies through reinforcement learning, where reward signals from a user simulator influence the system's output behavior. Another approach investigates how dialogue flows represented as graph structures can be used to explicitly constrain system output. The system's behavior can be structured along predefined paths of system and user actions, with the goal to improve reliability.

Beyond dialogue control, I also work on **Retrieval-Augmented Generation (RAG)** (Lewis et al., 2021) with **knowledge graphs** to ground responses. I have evaluated how GraphRAG (Edge et al., 2025) performs on dialogue data and contributed to a method for retrieving compact and relevant subgraphs to improve answer quality (Walker et al., 2025).

Finally, I explore how **emotion-aware dialogue systems** can be used to adapt their behavior based on the user's emotional state. I examined how using detected user emotion to influence response generation affects user perception in human-robot interaction, and compared different approaches to emotion classification in dialogue.

1.1 Controllability of Dialogue Generation

Large Language Models (LLMs) offer powerful capabilities for dialogue systems but often lack controllability in how responses are generated. I am investigating ways to introduce explicit control through a reinforcement learning approach, where a dialogue policy is trained based on reward signals from a user simulator. In parallel, within the research group, we are working on modeling dialogue flow as graph structures that encode possible system and user actions. These graphs are external to the LLM and help constrain its outputs to follow coherent and goal-oriented interaction patterns.

1.2 Retrieval-Augmented Generation with Knowledge Graphs

To address the limitations of unguided LLM responses, I explore RAG techniques that incorporate structured background knowledge. Specifically, I have evaluated the GraphRAG (Edge et al., 2025) framework on dialogue data to assess how well it supports contextually appropriate generation. We also developed subgraph retrieval strategies tailored to dialogue settings, aiming to maximize relevance while minimizing extraneous content (Walker et al., 2025). These approaches enhance grounding and allow the system to draw on precise and verifiable information during generation.

1.3 Emotion in Dialogue Systems

My research explores how emotion awareness can enhance interaction quality in dialogue systems. In a human-robot interaction study, it was examined how responses influenced by user emotion—automatically detected via the robot's perception system—affect user perception. I also investigate emotion recognition in human-human dialogue using both prompt-based and transformer-based methods. These insights can help build more adaptive and emotionally responsive dialogue agents that are sensitive to user state and context.

2 Spoken dialogue system (SDS) research

Dialogue Systems, especially those powered by Large Language Models, are increasingly becoming part of everyday workflows and decision-making processes. Their natural language interface lowers the barrier to access, allowing people of all backgrounds to interact with complex systems. However, this naturalness can also be misleading: when LLMs produce fluent but incorrect outputs, users may struggle to recognize inaccuracies (Ji et al., 2023). This highlights the need for better control mechanisms—such as dialogue policies, retrieval modules, or knowledge graph integration—to guide system behavior.

Researchers today are well-equipped to address these challenges. With both the technical understanding to improve these systems and the awareness to do so responsibly, we can design Dialogue Systems that are safe, fair, and useful in practice.

One central challenge is to reliably infer during the interaction what users want, especially in open-ended or

The 21th Annual Meeting of the Young Researchers' Roundtable on Spoken Dialogue Systems, pages 18–19
Aug 28–29, 2025.

ambiguous scenarios. As Dialogue Systems take on the role of interfaces to systems that used to require expert knowledge, they must be robust enough to handle diverse users and use cases. Addressing this challenge is essential for building dialogue systems that are inclusive, trustworthy, and effective.

3 Suggested topics for discussion

1. **Controlling LLM-based Dialogue Systems: How far can structure guide generation?** As LLMs replace traditional pipelines, how can we reintroduce controllability? What are effective methods—dialogue policies, graph structures, dialogue flow induction, or other approaches—that can be practically integrated into LLM-based Dialogue Systems?

2. **Retrieval-Augmented Dialogue Systems: How to retrieve meaningfully?** RAG approaches are powerful, but retrieving useful context can remain a key challenge. How can we improve retrieval strategies—such as using knowledge graphs or a subgraph selection—to ensure the context is both relevant and informative, especially in task-oriented settings?

3. **Emotion in Dialogue: How can emotional cues guide system behavior?** What role should emotion recognition play in dialogue systems? Beyond detecting emotion, how can it meaningfully guide system responses—and what are the ethical and design implications in human-robot and human-agent interactions?

References

Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, Dasha Metropolitansky, Robert Osazuwa Ness, and Jonathan Larson. 2025. From Local to Global: A Graph RAG Approach to Query-Focused Summarization. https://doi.org/10.48550/arXiv.2404.16130.

Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Yejin Bang, Delong Chen, Wenliang Dai, Ho Shu Chan, Andrea Madotto, and Pascale Fung. 2023. Survey of Hallucination in Natural Language Generation. *ACM Computing Surveys* 55(12):1–38. https://doi.org/10.1145/3571730.

Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2021. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. https://doi.org/10.48550/arXiv.2005.11401.

Nicolas Wagner and Stefan Ultes. 2024. On the Controllability of Large Language Models for Dialogue Interaction. In Tatsuya Kawahara, Vera Demberg, Stefan Ultes, Koji Inoue, Shikib Mehri, David Howcroft, and Kazunori Komatani, editors, *Proceedings of the 25th Annual Meeting of the Special Interest Group on Discourse and Dialogue*. Association for Computational Linguistics, Kyoto, Japan, pages 216–221. https://doi.org/10.18653/v1/2024.sigdial-1.19.

Nicholas Thomas Walker, Pierre Lison, Laetitia Hilgendorf, Nicolas Wagner, and Stefan Ultes. 2025. Retrieving relevant knowledge subgraphs for task-oriented dialogue. In *Proceedings of the 26th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*. Association for Computational Linguistics, Avignon, France.

Biographical sketch

Laetitia Hilgendorf is a Master's student in Applied Computer Science at the Otto-Friedrich-University of Bamberg, where she also completed her Bachelor's degree in 2024. Her academic focus lies in language technology, particularly in natural language understanding, natural language generation, and dialogue systems. Since 2023, she has been part of the university's Chair of Natural Language Generation and Dialogue Systems. In 2026, she will undertake a study stay at the Institute of Formal and Applied Linguistics in Prague to deepen her expertise, before beginning her PhD in Bamberg.