

Alyssa Allen

The Ohio State University
Oxley Hall
1712 Neil Avenue
Columbus, OH 43210

allen.2334@osu.edu
linguistics.osu.edu/people/allen.2334

1 Research interests

My research interests focus on **natural language generation** (NLG) regarding how to make system outputs more intuitive and comprehensible for the human-user and **conversational entrainment and alignment** from the perspective of how dialogue systems could or should personalize its responses to the human user. As it relates to NLG, my current work focuses on training a system to **auto-generate comments** for SQL queries produced by a Text-to-SQL parser. The goal is to make the connection between technical SQL language and the user's question more transparent. My linguistic training lies primarily at the intersection of computational and **soci-linguistics**. As such, my curiosities in conversational entrainment and alignment focus on the extent to which conversational agents can or should adjust their language based on human characteristics such as age, race, gender, etc.

1.1 Natural Language Generation (NLG)

My work in natural language generation has revolves around SQL query explainability. Users asking a question to a database may see the query as parsed by the system along with the results. Without notable amounts of training though, SQL commands can be difficult to understand, especially for complex queries. Additionally the output could yield unexpected or misleading results due to an incorrect parse of the user's initial question. The user may not be able to easily identify the mistake if only given the query as explanation.

Past research in making database language more comprehensible to humans have largely taken the route of summary comments or template language (Narechania et al., 2021; Kokkalis et al., 2012; Eleftherakis et al., 2021). These approaches do offer clarity on how the query answers the user's question. That said, templates or summaries can still heavily rely on database terminology (e.g. tables, columns, results, etc) that are not intuitive to the average human. The approaches also do not prioritize infusing user language into the comments or templates, leaving the cognitive load of making these connections to the human.

My research takes the approach of training a system to generate line-by-line comments for each SQL command, avoid database terminology, and leverage the user's language where appropriate. Line-by-line comments can directly state what information is being found in by each SQL command as a step toward answering the user question. One main benefit of this approach is that line-by-line comments can make errors in the SQL query more obvious to the human user.

My ongoing work has been focused on developing training data for such a model. I have manually annotated a small set of user questions and SQL queries. These hand-written examples of ideal comments are being used in few-shot prompting to ChatGPT where the model is tasked with generating comments for unseen queries. This set of ChatGPT-generated comments will become training and dev sets for fine-tuning an open-source LLM.

One challenge to generating natural-sounding comments for incorrect SQL queries is the balance between staying faithful to the query while integrating user language. Current findings suggest ChatGPT favors aligning with user language when the SQL command diverges from the user question, leaving errors hidden. If this bias exists in the ChatGPT-generated comments, it is likely the bias will persist in the fine-tuned open-source LLM as well. Further exploration of how to manage this problem is in progress.

This work will also explore improving the quality of the training and dev sets through filtering and comment refinement strategies. Open-sourced LLMs will be fine-tuned on each version of the training and dev sets to try improving comment generation capabilities. The ultimate goal is to have an open-source LLM be able to auto-generate comments for any unseen query.

1.2 Conversational Entrainment and Alignment

As mentioned in the previous section, leveraging the user's language can help alleviate points of confusion when explaining a rigid structure such as a database. Taking the user's language into consideration can also be beneficial in more flexible scenarios with non-database-oriented conversational agents.

As a future course of research, I am interested in approaching the question of conversational entrainment and alignment with a sociolinguistic lens for task-oriented dialogue systems.

Factors such as age, spoken dialect, relationship, power dynamic, gender, etc all impact human-human conversations. Humans take in this information subconsciously (or at times consciously) and may mirror or distance their language based on judgements made about the second interlocutor. Research in the field of human-machine interactions has started to better understand how humans align with the machine they are speaking to, such as Amazon’s Alexa or Apple’s Siri (Cohn et al., 2020), but I am primarily interested in the ways a machine can adjust their speech in order to improve the dialogue experience for humans. Leveraging human-human sociolinguistic findings (e.g. features of child directed speech (Nicola Dawson, 2021), perceptions of humor based on gender (Crawford, 2003), or variation patterns within different communities (Kiesling, 1998; DeCapua et al., 2006; Beebe and Takahashi, 1989) as a basis for future research questions, we can begin to investigate about how those preferences shift or remain intact for human-machine conversations.

My interests in this area build on the Computers Are Social Actors paradigm (Nass et al., 1994). If humans view computers as active interlocutors, then sociolinguistic insights from human-human conversations should provide some guidance in better developing and assessing human-machine conversations.

2 Spoken dialogue system (SDS) research

SDS research will continue to dive into questions around interpretability, explainability, and controllability as LLM capabilities progress. Questions around ethical usage of using conversational agents and improving robot social intelligence will also continue to be major considerations in the field. I think a growing interest will be in the dynamics occurring between the human and the dialogue system during a conversation.

As conversational agents become even more commonplace in the coming years, understanding the machine as an active interlocutor will be necessary in order to create more advanced conversational experiences for humans. I am curious about what socially-focused considerations could result in improved levels of personalization versus what adjustments could lead to toxic or harmful speech. For example, one could imagine a dialogue system trained to speak to all women in one way and all men in another would lead to harmful stereotyping. Improved levels of personalization though could take shape as a system picking up on a human user’s dialect and leveraging words from that dialect to appear more familiar.

When considering future applications of SDS, it will

be crucial for research in academic and industry spaces to be discussed and available between the groups. As products are being developed and launched in industry, findings from academia can be helpful in improving the way these systems are designed. I think this collaboration is where questions around what should a system be capable of in addition to what the system can do will be most effectively addressed.

3 Suggested topics for discussion

- What are the potential tradeoffs between developing conversational agents that rely on templated language versus leveraging LLMs to generate more flexible and dynamic dialogue?
- How can sociolinguistic theory impact research on conversational agent entrainment and alignment in order to improve dialogue personalization? How do elements such as the perceived gender of the machine voice impact effectiveness of the dialogue system.
- How can a multi-modal approach reduce cases of conversational ambiguity and improve the human-user experience?

References

- Leslie M. Beebe and Tomoko Takahashi. 1989. *Sociolinguistic Variation in Face-Threatening Speech Acts*, Springer US, Boston, MA, pages 199–218. https://doi.org/10.1007/978-1-4899-0900-8_13.
- Michelle Cohn, Patrik Jonell, Taylor Kim, Jonas Beskow, and Georgia Zellou. 2020. Embodiment and gender interact in alignment to tts voices.
- Mary Crawford. 2003. Gender and humor in social context. *Journal of Pragmatics* 35(9):1413–1430. *The Pragmatics of Humor*. [https://doi.org/https://doi.org/10.1016/S0378-2166\(02\)00183-2](https://doi.org/https://doi.org/10.1016/S0378-2166(02)00183-2).
- Andrea DeCapua, Diana Berkowitz, and Diana Boxer. 2006. Women talk revisited: Personal disclosures and alignment development 25(4):393–412. <https://doi.org/doi:10.1515/MULTI.2006.021>.
- Stavroula Eleftherakis, Orest Gkini, and Georgia Koutrika. 2021. Let the database talk back: Natural language explanations for sql. In *SEA-Data@VLDB*.
- Scott F. Kiesling. 1998. Men’s identities and sociolinguistic variation: The case of fraternity men. *Journal of Sociolinguistics* 2:69–99.
- Andreas Kokkalis, Panagiotis Vagenas, Alexandros Zervakis, Alkis Simitsis, Georgia Koutrika, and Yannis Ioannidis. 2012. Logos: A system for translating

queries into narratives. In *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data*. Association for Computing Machinery, New York, NY, USA, SIGMOD '12, pages 673–676. <https://doi.org/10.1145/2213836.2213929>.

Arpit Narechania, Adam Fourney, Bongshin Lee, and Gonzalo Ramos. 2021. Diy: Assessing the correctness of natural language to sql systems. In *26th International Conference on Intelligent User Interfaces*. Association for Computing Machinery, New York, NY, USA, IUI '21, pages 597–607. <https://doi.org/10.1145/3397481.3450667>.

Clifford Nass, Jonathan Steuer, and Ellen Siminoff. 1994. Computer are social actors. page 204. <https://doi.org/10.1145/259963.260288>.

Yaling Hsiao Alvin Wei Ming Tan Nilanjana Banerji Kate Nation Nicola Dawson. 2021. Features of lexical richness in children’s books: Comparisons with child-directed speech. *Language Development Research, Carnegie Mellon University Library Publishing Service* 1(1).

Biographical sketch



Alyssa Allen is a second-year PHD student at The Ohio State University. Her work focuses on developing conversational agents that leverage human-user language to generate more natural-sounding conversation turns. She is also interested in pursuing how sociolinguistic insights can improve human-machine communication.

Alyssa has a Master’s Degree in Linguistics from Eastern Michigan University and completed her Bachelor’s Degree from the University of Michigan. Prior to attending Eastern Michigan University, Alyssa worked in public relations with a focus on helping technology-focused companies promote and explain their AI-driven products to their respective audiences.