Overview of My PhD Research

Zhu Liu

Tsinghua University
liuzhu22@mails.tsinghua.edu.cn

This single page gives a short yet global picture of my first-author publications during my Ph.D journey since September 2022. My research focuses on lexical semantics from a computational perspective. In lexical semantics, I am particularly interested in the one-to-many correspondence between linguistic forms and meanings, such as polysemy, ambiguity, vagueness, multifunctionality. From a computational perspective, I explore how learning-based language models and/or hard-coded algorithms represent such correspondences. The representation involves two aspects: model-external evaluation and model-internal interpretation.

For model-external evaluation, our paper (Liu and Liu 2023), published in *Findings of ACL 2023*, investigates how neural network-based language models capture the uncertainty involved in Word Sense Disambiguation (WSD). We design multiple scenarios to simulate both data and model uncertainty, and employ quantitative metrics to evaluate the degrees of uncertainty under these settings. Afterwards, we analyze the multi-dimensional lexical properties, which may influence the uncertainty. We apply similar methods to a shared task (Liu, Hu, and Liu 2025) in CoMedi, a workshop at *Coling 2025*, which aims to estimate the annotator disagreement and the averaged sense labeling in determining to what extent the same word shares meaning across two distinct contexts. Our system formulates the two subtasks as parameter estimation problems under a Gaussian distribution, and achieves substantial improvement in the official post-evaluation stage.

Besides content words examined in previous studies, I am also interested in function words which often exhibit more flexible and diverse meanings in context, a phenomenon known as multifunctionality. The semantic map model is an effective tool widely used in the area of linguistic typology. We design a top-down graph-based algorithm (Liu et al. 2025b), published in *NAACL 2025*, and also develop a corresponding visualization tool (Liu et al. 2025a). Several case studies have empirically validated the effectiveness of our algorithm. We are currently working to expand the scale of our data and incorporate learning-based models in our future work.

For model-internal interpretation, I am curious about how the interior representations of language models with different architectures encode and reflect lexical semantic understanding. In our paper (Liu et al. 2024), published in *Findings of ACL 2024*, we compare decoder-only large language models (LLMs) with traditional encoder-only models (such as BERT) on a binary WSD task. We conclude that the lower layers of LLMs are primarily responsible for semantic understanding, whereas the higher layers focus more on prediction. In another paper (Liu et al. 2025c) currently under review, we explore the graph structures constructed from word embeddings across LLMs of different parameter scales, by analyzing their global statistics and local topological patterns. We find that these graphs consistently exhibit the small-world effect, although larger-scale models tend to have longer associative paths and stronger small-word characteristics among token nodes. This indicates that larger-scale models capture more complex relations among concepts.

References

Liu, Zhu, Zhen Hu, Lei Dai, and Ying Liu. 2025a. Xism: an exploratory and interactive graph tool to visualize and evaluate semantic map models. https://arxiv.org/abs/2507.04070.

Liu, Zhu, Zhen Hu, and Ying Liu. 2025. JuniperLiu at CoMeDi shared task: Models as annotators in lexical semantics disagreements. In *Proceedings of Context and Meaning: Navigating Disagreements in NLP Annotation*, pages 103–112, International Committee on Computational Linguistics, Abu Dhabi, UAE.

Liu, Zhu, Cunliang Kong, Ying Liu, and Maosong Sun. 2024. Fantastic semantics and where to find them: Investigating which layers of generative LLMs reflect lexical semantics. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 14551–14558, Association for Computational Linguistics, Bangkok, Thailand.

Liu, Zhu, Cunliang Kong, Ying Liu, and Maosong Sun. 2025b. A top-down graph-based tool for modeling classical semantic maps: A case study of supplementary adverbs. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 4567–4576, Association for Computational Linguistics, Albuquerque, New Mexico.

Liu, Zhu and Ying Liu. 2023. Ambiguity meets uncertainty: Investigating uncertainty estimation for word sense disambiguation. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 3963–3977, Association for Computational Linguistics, Toronto, Canada.

Liu, Zhu, Ying Liu, KangYang Luo, Cunliang Kong, and Maosong Sun. 2025c. From the new world of word embeddings: A comparative study of small-world lexico-semantic networks in llms. https://arxiv.org/abs/2502.11380.