

Zhehao Zhang

☎ (+1) 640-240-4027 • ✉ zhehao.zhang.gr@dartmouth.edu • **in** LinkedIn Profile
🐙 GitHub Profile • 🏠 Google Scholar • 🏠 Home Page • 🐦 @Zhehao_Zhang123

Education

Dartmouth College

Master of Science in Computer Science,

Hanover, NH

Sep 2023-Jun 2025 (Expected)

Shanghai Jiao Tong University (SJTU)

Bachelor of Engineering in Artificial Intelligence Honor Class,

Shanghai, China

Sep 2019-Jun 2023

Publications

1. **Zhehao Zhang**, Jiaao Chen, Diyi Yang, DARG: Dynamic Evaluation of Large Language Models via Adaptive Reasoning Graph, Arxiv [PDF Code](#)
2. **Zhehao Zhang**, Yan Gao, Jian-Guang Lou, E^5 : Zero-shot Hierarchical Table Analysis using Augmented LLMs via Explain, Extract, Execute, Exhibit and Extrapolate, Proceedings of the 2024 Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL) [PDF Code](#)
3. **Zhehao Zhang**, Xitao Li, Yan Gao, Jian-Guang Lou, CRT-QA: A Dataset of Complex Reasoning Question Answering over Tabular Data, the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP) [PDF Code](#)
4. **Zhehao Zhang**, Jiaao Chen, Diyi Yang, Mitigating Biases in Hate Speech Detection from A Causal Perspective, Findings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP) [PDF Code](#)
5. **Zhehao Zhang**, Tong Yu, Handong Zhao, Kaige Xie, Lina Yao, Shuai Li, "Exploring Soft Prompt Initialization Strategy for Few-shot Continual Text Classification," 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) [PDF](#)
6. Ziems Caleb, William Held, Omar Shaikh, Jiaao Chen, **Zhehao Zhang**, and Diyi Yang. "Can Large Language Models Transform Computational Social Science?." Computational Linguistics (2023). [PDF Code](#)

Industry Experience

Microsoft Research Asia

Research Intern, Data, Knowledge, and Intelligence Lab, Mentor: Dr. Yan Gao

Beijing, China

Dec 2022 - Aug 2023

- Explored Large Language Models'(e.g., GPT-4 etc.) reasoning ability on structured data. Constructed the first large-scale table question-answering dataset which requires the model to have multi-step complex reasoning capability with a detailed reasoning taxonomy. Comprehensively investigate LLMs' ability on different reasoning types. This work has been accepted by EMNLP 2023.
- Built a table analysis system for large hierarchical tables in a zero-shot manner using LangChain, which avoided hand-crafted in-context exemplars and considerably decreased the token usage in calling LLMs. This approach makes it possible for models with limited context length to analyze large-scale tabular data and achieve state-of-the-art performances. This work has been accepted by NAACL 2024.

Research Experience

Stanford University

Visiting Research Intern, Social and Language Technologies (SALT) Lab, Advisor: Prof.Diyi Yang

Stanford, CA

Jun 2022 - Present

- Searched for biased grammar patterns on hate speech detection datasets. Analyzed the spuriousness of different biases using causal interference, and then proposed a method to mitigate such biases based on several confounders. Validated the effectiveness of the method by running experiments across nine hate speech detection datasets with an out-of-domain challenge set to reach positive conclusions on its use for reducing hate speech bias. This work has been accepted by EMNLP 2023.
- Participated in constructing a road map for using LLMs as computational social science (CSS) tools and contributed a set of prompting best practices and an extensive evaluation pipeline to measure the zero-shot performance of 13 language models on 24 representative CSS benchmarks. Responsible for building and analyzing various baseline models (e.g., T-5, Roberta etc.) on all CSS datasets. [PDF](#)

Skills

Programming Languages: Python, C/C++, MATLAB, JavaScript

Tools and Frameworks: Retrieval-Augmented Generation (RAG), Parameter-efficient-tuning (PEFT), LangChain, Gradio, LitGPT, Git, GitHub, L^AT_EX, PyTorch, Huggingface, Numpy, Scikit-learn, Pandas

Spoken Language: English, Mandarin

Service

NeurIPS 2023: Reviewer, Thirty-seventh Conference on Neural Information Processing Systems Reviewer

EMNLP 2023: Reviewer, Volunteer, Conference on Empirical Methods in Natural Language Processing

ACL 2024: Reviewer, The 62nd Annual Meeting of the Association for Computational Linguistics

COLM 2024: Reviewer, The 1st Conference on Language Modeling

CIKM 2024: Reviewer, The 33rd ACM International Conference on Information and Knowledge Management