

Jia-Chen Gu

gujc@ucla.edu | (310) 254-6013 | <https://JasonForJoy.github.io/> | Google Scholar

EXPERIENCE AND EDUCATION

University of California, Los Angeles <i>Postdoctoral Researcher, hosted by Prof. Nanyun (Violet) Peng and Prof. Kai-Wei Chang</i>	Los Angeles, USA Nov 2023 - Present
University of Science and Technology of China <i>Postdoctoral Researcher, hosted by Prof. Zhen-Hua Ling</i>	Anhui, China Jun 2022 - Oct 2023
Microsoft <i>Research Intern, hosted by Dr. Chongyang Tao and Dr. Daxin Jiang</i>	Beijing, China Sep 2020 - Jun 2021
Queen's University <i>Visiting Student, hosted by Prof. Xiaodan Zhu</i>	Kingston, Canada Dec 2019 - May 2020
University of Science and Technology of China <i>Ph.D. in Signal and Information Processing, supervised by Prof. Zhen-Hua Ling</i>	Anhui, China Sep 2017 - Jun 2022
Hohai University <i>B.E. in Communication Engineering</i>	Jiangsu, China Sep 2013 - Jun 2017

SELECTED HONORS

Chancellor's Award for Postdoctoral Research Nomination Award, UCLA	2025
Best Paper Honorable Mention Award, ACL	2023
Best Paper Award, DialDoc@ACL	2022
Outstanding Doctoral Dissertation Nomination Award, Chinese Information Processing Society of China	2022
Presidential Scholarship, Chinese Academy of Sciences (Top 1%)	2022
Outstanding Graduate Award, USTC	2022
China National Scholarship, Ministry of Education	2021
Huayu Scholarship, NERC-SLIP (Top 2 graduates)	2021
Presidential Yan Kai Scholarship, Hohai University (Highest honor Top 0.1%)	2015

PUBLICATIONS

2.3k+ Google Scholar citations in total. Published at ACL, EMNLP, ICML, ICLR, SIGIR, CIKM, IJCAI, AAAI, TASLP, etc.
*: Equal contribution. #: Corresponding author.

- **Jia-Chen Gu***, Junyi Zhang*, Di Wu, Yuankai Li, Kai-Wei Chang, Nanyun Peng. BRIEF-Pro: Universal Context Compression with Short-to-Long Synthesis for Fast and Accurate Multi-Hop Reasoning. In *Findings of the Association for Computational Linguistics: ACL 2026*. 2026.
- Qingyuan Liu*, **Jia-Chen Gu***, Yunzhi Yao, Hong Wang, Nanyun Peng. Energy-Regularized Sequential Model Editing on Hyperspheres. In *Proceedings of the Fourteenth International Conference on Learning Representations (ICLR)*. 2026.
- Hao-Xiang Xu, Ziqi Peng, Jun-Yu Ma, Yuhao Sun, Zhen-Hua Ling, **Jia-Chen Gu#**. TangleScore: Tangle-Guided Purge and Imprint for Unstructured Knowledge Editing. In *Proceedings of the Fourteenth International Conference on Learning Representations (ICLR)*. 2026.
- Hao-Xiang Xu, Jun-Yu Ma, Ziqi Peng, Yuhao Sun, Zhen-Hua Ling, **Jia-Chen Gu#**. Multiplicative Orthogonal Sequential Editing for Language Models. In *Proceedings of the 40th AAAI Conference on Artificial Intelligence*. 2026.
- Jun-Yu Ma, Hong Wang, Hao-Xiang Xu, Zhen-Hua Ling, **Jia-Chen Gu#**. Perturbation-Restrained Sequential Model Editing. In *Proceedings of the Thirteenth International Conference on Learning Representations (ICLR)*. 2025.
- Yuankai Li*, **Jia-Chen Gu***, Di Wu, Kai-Wei Chang, Nanyun Peng. BRIEF: Bridging Retrieval and Inference for Multi-hop Reasoning via Compression. In *Findings of the Association for Computational Linguistics: NAACL 2025*. 2025.

- **Jia-Chen Gu***, Hao-Xiang Xu*, Jun-Yu Ma, Pan Lu, Zhen-Hua Ling, Kai-Wei Chang, Nanyun Peng. Model Editing Harms General Abilities of Large Language Models: Regularization to the Rescue. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 2024. **Best Paper Nomination <2%**.
- Jun-Yu Ma, Zhen-Hua Ling, Ningyu Zhang, **Jia-Chen Gu#**. Neighboring Perturbations of Knowledge Editing on Large Language Models. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*. 2024.
- **Jia-Chen Gu**, Zhen-Hua Ling, Quan Liu, Cong Liu, Guoping Hu. GIFT: Graph-Induced Fine-Tuning for Multi-Party Conversation Understanding. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL)*. 2023. **Best Paper Honorable Mention Award**.
- **Jia-Chen Gu***, Chao-Hong Tan*, Caiyuan Chu, Zhen-Hua Ling, Chongyang Tao, Quan Liu, Cong Liu. MADNet: Maximizing Addressee Deduction Expectation for Multi-Party Conversation Generation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 2023.
- **Jia-Chen Gu***, Chao-Hong Tan*, Chongyang Tao, Zhen-Hua Ling, Huang Hu, Xiubo Geng, Daxin Jiang. HeterMPC: A Heterogeneous Graph Neural Network for Response Generation in Multi-Party Conversations. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (ACL)*. 2022.
- **Jia-Chen Gu**, Chongyang Tao, Zhen-Hua Ling. Who Says What to Whom: A Survey of Multi-Party Conversations. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*. 2022.
- **Jia-Chen Gu**, Chongyang Tao, Zhen-Hua Ling, Can Xu, Xiubo Geng, Daxin Jiang. MPC-BERT: A Pre-Trained Language Model for Multi-Party Conversation Understanding. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics (ACL)*. 2021.
- **Jia-Chen Gu**, Zhen-Hua Ling, Yu Wu, Quan Liu, Zhigang Chen, Xiaodan Zhu. Detecting Speaker Personas from Conversational Texts. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 2021.
- **Jia-Chen Gu**, Hui Liu, Zhen-Hua Ling, Quan Liu, Zhigang Chen, Xiaodan Zhu. Partner Matters! An Empirical Study on Fusing Personas for Personalized Response Selection in Retrieval-Based Chatbots. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*. 2021.
- **Jia-Chen Gu**, Tianda Li, Zhen-Hua Ling, Quan Liu, Zhiming Su, Yu-Ping Ruan, Xiaodan Zhu. Deep Contextualized Utterance Representations for Response Selection and Dialogue Analysis. *IEEE/ACM Transactions on Audio, Speech, and Language Processing (TASLP)*. 2021.
- **Jia-Chen Gu**, Zhen-Hua Ling, Quan Liu, Zhigang Chen, Xiaodan Zhu. Filtering before Iteratively Referring for Knowledge-Grounded Response Selection in Retrieval-Based Chatbots. In *Findings of the Association for Computational Linguistics: EMNLP 2020*. 2020.
- **Jia-Chen Gu**, Tianda Li, Quan Liu, Zhen-Hua Ling, Zhiming Su, Si Wei, Xiaodan Zhu. Speaker-Aware BERT for Multi-Turn Response Selection in Retrieval-Based Chatbots. In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management (CIKM)*. 2020.
- **Jia-Chen Gu**, Zhen-Hua Ling, Quan Liu. Utterance-to-Utterance Interactive Matching Network for Multi-Turn Response Selection in Retrieval-Based Chatbots. *IEEE/ACM Transactions on Audio, Speech, and Language Processing (TASLP)*. 2020.
- **Jia-Chen Gu**, Tianda Li, Quan Liu, Xiaodan Zhu, Zhen-Hua Ling, Yu-Ping Ruan. Pre-Trained and Attention-Based Neural Networks for Building Noetic Task-Oriented Dialogue Systems. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence Workshop on the 8th Dialog System Technology Challenge (DSTC)*. 2020.
- **Jia-Chen Gu**, Zhen-Hua Ling, Xiaodan Zhu, Quan Liu. Dually Interactive Matching Network for Personalized Response Selection in Retrieval-Based Chatbots. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 2019.
- **Jia-Chen Gu**, Zhen-Hua Ling, Quan Liu. Interactive Matching Network for Multi-Turn Response Selection in Retrieval-Based Chatbots. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management (CIKM)*. 2019.
- **Jia-Chen Gu**, Zhen-Hua Ling, Yu-Ping Ruan, Quan Liu. Building Sequential Inference Models for End-to-End Response Selection. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence Workshop on the 7th Dialog System Technology Challenge (DSTC)*. 2019.
- **Jia-Chen Gu**, Zhen-Hua Ling, Nitin Indurkha. A Study on Improving End-to-End Neural Coreference Resolution. In *Proceedings of the 17th China National Conference on Computational Linguistics (CCL)*. 2018.
- Xiaojie Gu, Ziyang Huang, **Jia-Chen Gu**, Kai Zhang. UltraEdit: Training-, Subject-, and Memory-Free Lifelong Editing in Language Models. *Transactions on Machine Learning Research (TMLR)*. 2026.

- Canyu Chen, Baixiang Huang, Zekun Li, Zhaorun Chen, Shiyang Lai, Xiong Xiao Xu, **Jia-Chen Gu**, Jindong Gu, Huaxiu Yao, Chaowei Xiao, Xifeng Yan, William Yang Wang, Philip Torr, Dawn Song, Kai Shu. Can Editing LLMs Inject Harm? In *Proceedings of the 40th AAAI Conference on Artificial Intelligence*. 2026.
- Wenbo Hu, **Jia-Chen Gu**, Zi-Yi Dou, Mohsen Fayyaz, Pan Lu, Kai-Wei Chang, Nanyun Peng. MRAG-Bench: Vision-Centric Evaluation for Retrieval-Augmented Multimodal Models. In *Proceedings of the Thirteenth International Conference on Learning Representations (ICLR)*. 2025.
- Yun Zhu, **Jia-Chen Gu**, Caitlin Sikora, Ho Ko, Yinxiao Liu, Chu-Cheng Lin, Lei Shu, Liangchen Luo, Lei Meng, Bang Liu, Jindong Chen. Accelerating Inference of Retrieval-Augmented Generation via Sparse Context Selection. In *Proceedings of the Thirteenth International Conference on Learning Representations (ICLR)*. 2025.
- Yunzhi Yao, Jizhan Fang, **Jia-Chen Gu**, Ningyu Zhang, Shumin Deng, Huajun Chen, Nanyun Peng. CaKE: Circuit-aware Editing Enables Generalizable Knowledge Learners. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 2025.
- Hang Chen, Chao-Han Huck Yang, **Jia-Chen Gu**, Sabato Marco Siniscalchi, Jun Du. MISP-Meeting: A Real-World Dataset with Multimodal Cues for Long-form Meeting Transcription and Summarization. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (ACL)*. 2025.
- Hao-Xiang Xu, Jun-Yu Ma, Zhen-Hua Ling, Ningyu Zhang, **Jia-Chen Gu**. Constraining Sequential Model Editing with Editing Anchor Compression. In *Findings of the Association for Computational Linguistics: NAACL 2025*. 2025.
- Di Wu, **Jia-Chen Gu**, Fan Yin, Nanyun Peng, Kai-Wei Chang. Synchronous Faithfulness Monitoring for Trustworthy Retrieval-Augmented Generation. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 2024.
- Zhen Li, Xiaohan Xu, Tao Shen, Can Xu, **Jia-Chen Gu**, Yuxuan Lai, Chongyang Tao, Shuai Ma. Leveraging Large Language Models for NLG Evaluation: Advances and Challenges. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 2024.
- Mengru Wang, Yunzhi Yao, Ziwen Xu, Shuofei Qiao, Shumin Deng, Peng Wang, Xiang Chen, **Jia-Chen Gu**, Yong Jiang, Pengjun Xie, Fei Huang, Huajun Chen, Ningyu Zhang. Knowledge Mechanisms in Large Language Models: A Survey and Perspective. In *Findings of the Association for Computational Linguistics: EMNLP 2024*. 2024.
- Qian Wang, **Jia-Chen Gu**, Zhen-Hua Ling. X-ACE: Explainable and Multi-factor Audio Captioning Evaluation. In *Findings of the Association for Computational Linguistics: ACL 2024*. 2024.
- Qian Wang, **Jia-Chen Gu**, Zhen-Hua Ling. Multiscale Matching Driven by Cross-Modal Similarity Consistency for Audio-Text Retrieval. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 2024.
- Jun-Yu Ma, **Jia-Chen Gu**, Zhen-Hua Ling, Quan Liu, Cong Liu, Guoping Hu. Syntax-augmented Hierarchical Interactive Encoder for Zero-shot Cross-lingual Information Extraction. *IEEE/ACM Transactions on Audio, Speech, and Language Processing (TASLP)*. 2024.
- Chao-Hong Tan, **Jia-Chen Gu**, Zhen-Hua Ling. Is ChatGPT a Good Multi-Party Conversation Solver?. In *Findings of the Association for Computational Linguistics: EMNLP 2023*. 2023.
- Yue Chen, Tian-Wei He, Hong-Bin Zhou, **Jia-Chen Gu**, Heng Lu, Zhen-Hua Ling. Symbolization, Prompt, and Classification: A Framework for Implicit Speaker Identification in Novels. In *Findings of the Association for Computational Linguistics: EMNLP 2023*. 2023.
- Jun-Yu Ma, **Jia-Chen Gu**, Jiajun Qi, Zhen-Hua Ling, Quan Liu, Xiaoyi Zhao. USTC-NELSLIP at SemEval-2023 Task 2: Statistical Construction and Dual Adaptation of Gazetteer for Multilingual Complex NER. In *Proceedings of the 17th International Workshop on Semantic Evaluation*. 2023.
- Caiyuan Chu, Ya Li, Yifan Liu, **Jia-Chen Gu**, Quan Liu, Yongxin Ge, Guoping Hu. Multi-Stage Coarse-to-Fine Contrastive Learning for Conversation Intent Induction. In *Proceedings of Workshop on the 11th Dialog System Technology Challenge*. 2023.
- Jun-Yu Ma, Beiduo Chen, **Jia-Chen Gu**, Zhen-Hua Ling, Wu Guo, Quan Liu, Zhigang Chen, Cong Liu. Wider & Closer: Mixture of Short-channel Distillers for Zero-shot Cross-lingual Named Entity Recognition. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. 2022.
- Chao-Hong Tan, **Jia-Chen Gu**, Chongyang Tao, Zhen-Hua Ling, Can Xu, Huang Hu, Xiubo Geng, Daxin Jiang. TegTok: Augmenting Text Generation via Task-specific and Open-world Knowledge. In *Findings of the Association for Computational Linguistics: ACL 2022*. 2022.

- Tianda Li, **Jia-Chen Gu**, Zhen-Hua Ling, Quan Liu. Conversation- and Tree-Structure Losses for Dialogue Disentanglement. In *Proceedings of the Second DialDoc Workshop on Document-grounded Dialogue and Conversational Question Answering*. 2022. **Best Paper Award**.
- Tianda Li, **Jia-Chen Gu**, Hui Liu, Quan Liu, Zhen-Hua Ling, Zhiming Su, Xiaodan Zhu. Have You Made A Decision? Where? A Pilot Study on Interpretability of Polarity Analysis Based on Advising Problem. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 2021.
- Chao-Hong Tan, Xiaoyu Yang, Zi'ou Zheng, Tianda Li, Yufei Feng, **Jia-Chen Gu**, Quan Liu, Dan Liu, Zhen-Hua Ling, Xiaodan Zhu. Learning to Retrieve Entity-Aware Knowledge and Generate Responses with Copy Mechanism for Task-Oriented Dialogue Systems. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence Workshop on the 9th Dialog System Technology Challenge (DSTC)*. 2021.
- Hui Liu, Zhan Shi, **Jia-Chen Gu**, Quan Liu, Si Wei, Xiaodan Zhu. End-to-End Transition-Based Online Dialogue Disentanglement. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*. 2020.
- Yu-Ping Ruan, Zhen-Hua Ling, **Jia-Chen Gu**, Quan Liu. Fine-Tuning BERT for Schema-Guided Zero-Shot Dialogue State Tracking. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence Workshop on the 8th Dialog System Technology Challenge (DSTC)*. 2020.
- Yu-Ping Ruan, Zhen-Hua Ling, Xiaodan Zhu, Quan Liu, **Jia-Chen Gu**. Generating Diverse Conversation Responses by Creating and Ranking Multiple Candidates. *Computer Speech and Language*. 2020.
- Yu-Ping Ruan, Zhen-Hua Ling, Quan Liu, **Jia-Chen Gu**, Xiaodan Zhu. Promoting Diversity for End-to-End Conversation Response Generation. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence Workshop on the 7th Dialog System Technology Challenge (DSTC)*. 2019.

PREPRINTS

- Di Wu*, **Jia-Chen Gu***, Kai-Wei Chang, Nanyun Peng. Self-Routing RAG: Binding Selective Retrieval with Knowledge Verbalization. *Preprint*. 2025.
- Junyi Zhang, **Jia-Chen Gu**, Wenbo Hu, Yu Zhou, Robinson Piramuthu, Nanyun Peng. TemMed-Bench: Evaluating Temporal Medical Image Reasoning in Vision-Language Models. *Preprint*. 2025.
- Rui Li, **Jia-Chen Gu**, Po-Nien Kung, Heming Xia, Junfeng Liu, Xiangwen Kong, Zhifang Sui, Nanyun Peng. LLM-REVal: Can We Trust LLM Reviewers Yet?. *Preprint*. 2025.
- Shi-Qi Yan*, **Jia-Chen Gu***, Yun Zhu, Zhen-Hua Ling. Corrective Retrieval Augmented Generation. *Preprint*. 2024.
- Jun-Yu Ma*, **Jia-Chen Gu***, Zhen-Hua Ling, Quan Liu, Cong Liu. Untying the Reversal Curse via Bidirectional Language Model Editing. *Preprint*. 2023.

TUTORIALS

- Ningyu Zhang, **Jia-Chen Gu**, Yunzhi Yao, Mengru Wang, Xiang Chen, Shumin Deng. "Knowledge Editing for Large Language Models" at IJCAI 2024.
- Ningyu Zhang, **Jia-Chen Gu**, Yunzhi Yao, Zhen Bi, Shumin Deng. "Knowledge Editing for Large Language Models" at AAAI 2024.
- **Jia-Chen Gu**, Zhuosheng Zhang, Zhen-Hua Ling. "Learning WHO Saying WHAT to WHOM in Multi-Party Conversations" at IJCNLP-AAACL 2023.
- **Jia-Chen Gu**. "Learning WHO Saying WHAT to WHOM in Multi-Party Conversations" at CCKS 2023.

LECTURES

- 2025.10: Guest lecture "Linear Regression & Multi-Class Classification" at UCLA CM146.
- 2025.04: Guest lecture "Towards Trustworthy and Efficient Retrieval-Augmented Generation" at CUHK-Shenzhen.
- 2024.07: Instructor of "Generative AI" at UCLA Summer Institute CS97.
- 2024.02: Guest lecture "Pre-trained Large Language Models" at UCLA CS162.
- 2023.12: Guest lecture "Multi-Party Conversational AI" at Zhejiang University.

INVITED TALKS

- 2025.04: "Self-Routing RAG: Binding Selective Retrieval with Knowledge Verbalization" at MLNLP
- 2024.08: "Measuring and Mitigating Perturbations of Knowledge Editing on Large Language Models" at CAS-ICT.
- 2023.09: "Learning Conversation Structures for Multi-Party Conversations" at MLNLP 2023.
- 2023.05: "Learning Conversation Structures for Multi-Party Conversations" at YSSNLP 2023.
- 2022.11: "Multi-Party Conversation Understanding and Generation" at CCF and CAS-ICT.
- 2022.07: "Multi-Party Conversation Understanding and Generation" at MLNLP

SHARED TASKS

- 2022, Rank first on both subtasks of Track 2 in the 11th Dialog System Technology Challenges (DSTC 11).
- 2020, Rank second under objective metrics and fourth under subjective metrics on Track 1 in the 9th Dialog System Technology Challenges (DSTC 9), as the team leader.
- 2019, Rank first, first, third and fourth on subtasks of Track 2 in the 8th Dialog System Technology Challenges (DSTC 8), as the first author.
- 2018, Rank second on a subtask of Track 1 in the 7th Dialog System Technology Challenges (DSTC 7), as the first author.

SERVICE

- Area Chair for Conferences: ACL, EMNLP, NAACL, NeurIPS, ICML.
- Reviewer for Conferences: ACL, EMNLP, NAACL, EACL, ACL Rolling Review, ICLR, COLM, AAAI, IJCAI, NLPCC, CCKS.
- Reviewer for Journals: TASLP, TOIS, TALLIP, CIM, IPM, CL, SPL.