

Junda (Joshua) Wu

+1 (646)-515-3105 | La Jolla, CA

juw069@ucsd.edu | github.com/JoshuaWu1997

scholar.google.com/citations?user=_iKeQFwAAAAJ&hl

EDUCATION

University of California San Diego

Ph.D. in Computer Science

Sep 2023 - Now

New York University

M.S. in Computer Engineering

Sep 2021 – May 2023

Chongqing University

B.S. in Statistics, Minor in Computer Science and Engineering

Sep 2016 – Jul 2020

PROFESSIONAL EXPERIENCES

Adobe Research, San Jose

Research Scientist Intern

June 2024 – Nov 2024

- Multimodal Retrieval in AI Assistant for AEP

- Developed a multimodal retriever enhanced by MLLM dense captions and OCR-extracted text to improve visual and textual fusion.
- Developed the Doc-React algorithm, using the proposed multimodal retriever based on user queries, enhancing document interaction.

Adobe Research, San Jose

Research Scientist Intern

May 2023 – Aug 2023

- Knowledge Graph Enhanced Chain-of-Thought Reasoning for Next Prompt Recommendation

- Construct the Adobe AEP specific knowledge graph using LLM-based information extraction methods.
- Built a prompt recommender system to track and explore user intentions on KG and generate next prompt candidates.

Adobe Research, San Jose

Research Scientist Intern

May 2022 – Dec 2022

- Progressive Image Compression and Syncing

- Developed an algorithm to prioritize the most important pixels to be compressed and transmitted, which improves the quality of previews.
- Built a progressive syncing prototype to enable real-time online collaborative image editing feature.

Conference Reviewer:

- ACL Rolling Review: ACL, EMNLP, NAACL, COLING
- Machine Learning: NeurIPS, ICLR, AISTAT, ICML
- Data Mining and Recommendation: KDD, CIKM, SIGIR, WWW, TKDE

PATENTS

- Haoliang Wang, Tong Yu, Sungchul Kim, Ruiyi Zhang, Paiheng Xu, **Junda Wu**, Handong Zhao, and Ani Nenkova. "CONTEXTUAL QUERY GENERATION." U.S. Patent Application 18/339,694, filed December 26, 2024.
- Haoliang Wang, Kaige Xie, Tong Yu, **Junda Wu**, Handong Zhao, Ruiyi Zhang, Kanak Vivek Mahadik, and Ani Nenkova. "Dialogue state aware dialogue summarization." U.S. Patent Application 18/343,389, filed January 2, 2025.
- Tong Yu, Kaige Xie, Haoliang Wang, **Junda Wu**, Handong Zhao, Ruiyi Zhang, Kanak Vivek Mahadik, and Ani Nenkova. "Dialogue skeleton assisted prompt transfer for dialogue summarization." U.S. Patent Application 18/355,901, filed January 23, 2025.
- Mahadik, Kanak, Tong Yu, and **Junda Wu**. "Transferable clustering of contextual bandits for cloud service resource allocation." U.S. Patent Application 18/342,516, filed January 2, 2025.

SELECTED PUBLICATIONS

Under Review and arXiv:

- **Junda Wu**, Rohan Surana, Zhouhang Xie, Yiran Shen, Yu Xia, Tong Yu, Ryan A. Rossi, Prithviraj Ammanabrolu, and Julian McAuley. "In-context Ranking Preference Optimization." *arXiv preprint arXiv:2504.15477* (Under review COLM'25).
- **Junda Wu**, Yuxin Xiong, Xintong Li, Yu Xia, Ruoyu Wang, Yu Wang, Tong Yu et al. "Mitigating Visual Knowledge Forgetting in MLLM Instruction-tuning via Modality-decoupled Gradient Descent." *arXiv preprint arXiv:2502.11740* (Under review SIGKDD'25).
- Wang, Rui, **Junda Wu**, Yu Xia, Tong Yu, Ruiyi Zhang, Ryan Rossi, Lina Yao, and Julian McAuley. "CachePrune: Neural-Based Attribution Defense Against Indirect Prompt Injection Attacks." *arXiv preprint arXiv:2504.21228* (Under review ACL'25).
- Rohan Surana, **Junda Wu**, Zhouhang Xie, Yu Xia, Harald Steck, Dawen Liang, Nathan Kallus, and Julian McAuley. "From Reviews to Dialogues: Active Synthesis for Zero-Shot LLM-based Conversational Recommender System." *arXiv preprint arXiv:2504.15476* (Under review RecSys'25).
- Branislav Kveton, Xintong Li, Julian McAuley, Ryan Rossi, Jingbo Shang, **Junda Wu**, and Tong Yu. (alphabetic order) "Active Learning for Direct Preference Optimization." *arXiv preprint arXiv:2503.01076* (2025).
- **Junda Wu**, Xintong Li, Tong Yu, Yu Wang, Xiang Chen, Jiuxiang Gu, Lina Yao, Jingbo Shang, and Julian McAuley. "Commit: Coordinated instruction tuning for multimodal large language models." *arXiv preprint arXiv:2407.20454* (2024).

2025:

- **Junda Wu**, Xintong Li, Ruoyu Wang, Yu Xia, Yuxin Xiong, Jianing Wang, Tong Yu et al. "OCEAN: Offline Chain-of-thought Evaluation and Alignment in Large Language Models." In *The Thirteenth International Conference on Learning Representations (ICLR'25)*.
- Wang, Yu, Xinshuang Liu, Xiushi Chen, Sean O'Brien, **Junda Wu**, and Julian McAuley. "Self-updatable large language models by integrating context into model parameters." In *The Thirteenth International Conference on Learning Representations (ICLR'25)*.
- Chengkai Huang, **Junda Wu**, Tong Yu, Julian McAuley, Lina Yao. "A Tutorial on Agentic LLM for Recommender Systems" Accepted by 19th ACM Conference on Recommender Systems (RecSys'25)

- **Junda Wu**, Warren Li, Zachary Novack, Amit Namburi, Carol Chen, and Julian McAuley. "CoLLAP: Contrastive Long-form Language-Audio Pretraining with Musical Temporal Structure Augmentation." In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP'25)*.
- **Junda Wu**, Zachary Novack, Amit Namburi, Hao-Wen Dong, Carol Chen, Jiaheng Dai, and Julian McAuley. "FUTGA-MIR: Enhancing Fine-grained and Temporally-aware Music Understanding with Music Information Retrieval." In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP'25)*.
- Songwen Hu, Ryan A. Rossi, Tong Yu, **Junda Wu**, Handong Zhao, Sungchul Kim, and Shuai Li. "Interactive Visualization Recommendation with Hier-SUCB." In *Proceedings of the ACM on Web Conference (WWW'25)*.

2024:

- **Junda Wu**, Eric Chang, Tong Yu, Zhankui He, Jianing Wang, Yupeng Hou, Julian McAuley. CoRAL: Collaborative retrieval-augmented large language models improve long-tail recommendation. In *Proceedings of the Knowledge Discovery and Data Mining (KDD'24)*
- **Junda Wu**, Tong Yu, Xiang Chen, Haoliang Wang, Ryan Rossi, Sungchul Kim, Anup Rao, Julian McAuley. DeCoT: Debiasing chain-of-thought for knowledge-intensive tasks in large language models via causal intervention. In *Proceedings of Annual Meeting of (ACL'24)*.
- Jianing Wang, **Junda Wu**, Yupeng Hou, Yao Liu, Ming Gao, Julian McAuley. InstructGraph: Boosting large language models via graph-centric instruction tuning and preference alignment. In *Proceedings of Findings of (ACL'24)*.
- Yu Xia, **Junda Wu**, Sungchul Kim, Tong Yu, Ryan A. Rossi, Haoliang Wang, and Julian McAuley. "Knowledge-aware query expansion with large language models for textual and relational retrieval." Accepted by the Association for Computational Linguistics (NAACL'24).
- Xiaoyu Liu, Paiheng Xu, Junda Wu, Jiaxin Yuan, Yifan Yang, Yuhang Zhou, Fuxiao Liu et al. "Large Language Models and Causal Inference in Collaboration: A Comprehensive Survey." Accepted by the Association for Computational Linguistics (NAACL'24).
- Zhouhang Xie, **Junda Wu**, Hyunsik Jeon, Zhankui He, Harald Steck, Rahul Jha, Dawen Liang, Nathan Kallus, and Julian McAuley. "Neighborhood-Based Collaborative Filtering for Conversational Recommendation." In *Proceedings of the 18th ACM Conference on Recommender Systems (RecSys'24)*.
- An Yan, Zhengyuan Yang, **Junda Wu**, Wanrong Zhu, Jianwei Yang, Linjie Li, Kevin Lin, Jianfeng Wang, Julian McAuley, Jianfeng Gao, Lijuan Wang. List items one by one: A new data source and learning paradigm for multimodal LLMs. In *Proceedings of the Conference on Language Modeling (COLM'24)*
- Wang, Rui, Tong Yu, Ruiyi Zhang, Sungchul Kim, Ryan Rossi, Handong Zhao, **Junda Wu**, Subrata Mitra, Lina Yao, and Ricardo Henao. "Personalized Federated Learning for Text Classification with Gradient-Free Prompt Tuning." In *Findings of the Association for Computational Linguistics (NAACL'24)*.
- Xu Liu, Tong Yu, Kaige Xie, **Junda Wu**, and Shuai Li. Interact with the Explanations: Causal Debaised Explainable Recommendation System. 2024. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining (WSDM'24)*

2021 - 2023:

- **Junda Wu**, Tong Yu, Rui Wang, Zhao Song, Ruiyi Zhang, Handong Zhao, Chaochao Lu, Shuai Li, Ricardo Henao. 2023. InfoPrompt: Information-Theoretic Soft Prompt Tuning for Natural Language Understanding. 2023. In *Proceedings of the 37th Conference on Neural Information Processing Systems (NeurIPS'23)*
- Yu Xia, **Junda Wu**, Tong Yu, Sungchul Kim, Ryan A Rossi, Shuai Li. Junda Wu User-Regulation Deconfounded Conversational Recommender System with Bandit Feedback. 2023. In *Proceedings of the 29th ACM Knowledge Discovery and Data Mining (SIGKDD'23)*
- Rui Wang, Tong Yu, **Junda Wu**, Handong Zhao, Sungchul Kim, Ruiyi Zhang, Subrata Mitra, and Ricardo Henao. 2023. Federated Domain Adaptation for Named Entity Recognition via Distilling with Heterogeneous Tag Sets. In *Proceedings of the Association for Computational Linguistics (ACL'23 Findings)*.
- **Junda Wu**, Rui Wang, Handong Zhao, Ruiyi Zhang, Chaochao Lu, Shuai Li and Ricardo Henao. 2023. Few-Shot Composition Learning for Image Retrieval with Prompt Tuning. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI'23)*.
- **Junda Wu**, Rui Wang, Tong Yu, Ruiyi Zhang, Handong Zhao, Shuai Li, Ricardo Henao and Ani Nenkova. 2022. Context-aware Information-theoretic Causal De-biasing for Interactive Sequence Labeling. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP'22 Findings)*.
- **Junda Wu**, Zhihui Xie, Tong Yu, Handong Zhao, Ruiyi Zhang, Shuai Li. 2022. Dynamics-Aware Adaptation for Reinforcement Learning Based Cross-Domain Interactive Recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR'22)*.
- Yong Zhuang, Tong Yu, **Junda Wu**, Shiqu Wu, and Shuai Li. 2022. Spatial-Temporal Aligned Multi-Agent Learning for Visual Dialog Systems. In *Proceedings of the 30th ACM International Conference on Multimedia (MM '22)*.
- **Junda Wu**, Canzhe Zhao, Tong Yu, Jingyang Li, and Shuai Li. 2021. Clustering of Conversational Bandits for User Preference Learning and Elicitation. In *Proceedings of the 30th ACM International Conference on Information and Knowledge Management (CIKM '21)*.
- **Junda Wu**, Tong Yu, and Shuai Li. 2021. Deconfounded and Explainable Interactive Vision-Language Retrieval of Complex Scenes. In *Proceedings of the 29th ACM International Conference on Multimedia (MM '21)*.