

Ziyan Jiang

Ph.D. Student in Computer Science, UC Santa Barbara | Palo Alto, US | Updated October 26, 2026

✉ ziyanjiang528@gmail.com | 🏠 [xmhzz2018.github.io](https://github.com/xmhzz2018) | 🎓 [Google Scholar](#) | 🔄 [XMHZZ2018](#) | 🐦 [@Ernestzyj](#)

Education

University of California, Santa Barbara

Ph.D. in Computer Science

Advisor: Prof. Shiyu Chang

Santa Barbara, CA

Sep 2025 – Present

Johns Hopkins University

M.S. in Electrical and Computer Engineering

Baltimore, MD

Sep 2018 – May 2020

Nanjing University

B.S. in Acoustics and Minor in Physics

Nanjing, China

Sep 2014 – Jul 2018

Research & Industry Experience

Research Intern, Salesforce AI Research

Mentors: Semih Yavuz and Shafiq Joty

Research on self-evolving agents that learn from their work, transfer useful lessons to unfamiliar tasks and environments, and improve as experience accumulates.

Palo Alto, CA

Jun 2026 – Present

Applied Scientist II, Amazon Alexa AI / AGI

Seattle / Bellevue, WA

Jul 2020 – Aug 2024

Developed machine learning systems to learn customers' implicit affinities, preferences, and habits from interactions with Amazon products. Developed Alexa query rewriting systems to identify and reduce user defects and improve natural language understanding.

Research Assistant, Johns Hopkins University (CLSP)

Advisors: Shinji Watanabe and Dan Povey

Worked on end-to-end speech recognition and contributed to ESPnet, achieving state-of-the-art results on more than 10 datasets spanning automatic speech recognition, speech translation, and text-to-speech. Studied adversarial examples to improve automatic speech recognition.

Baltimore, MD

Jan 2019 – Jan 2020

Publications

* denotes equal contribution where marked. [Google Scholar](#).

2026

- [1] **WorldAuditBench: Interactive 3D World Auditing with Multimodal Agents**
Ziyan Jiang*, Jingbo Yang*, Jiabao Ji*, Yujian Liu, Qiucheng Wu, Tommi Jaakkola, Yang Zhang, Shiyu Chang
[arXiv preprint](#) | [Paper](#) | [Website](#) | [Code](#)
- [2] **VisualSkill: Multimodal Skills for Computer-Use Agents**
Ziyan Jiang*, Li An*, Yujian Liu, Jiabao Ji, Qiucheng Wu, Jacob Andreas, Yang Zhang, Shiyu Chang
[arXiv preprint](#) | [Paper](#) | [Code](#)
- [3] **MMEB-V3: Measuring the Performance Gaps of Omni-Modality Embedding Models**
Haohang Huang, Xuan Lu, Mingyi Su, Xuan Zhang, Ziyan Jiang, Ping Nie, Kai Zou, Tomas Pfister, Wenhu Chen, Wei Zhang, Xiaoyu Shen, Rui Meng
[COLM 2026](#) | [Paper](#) | [Code](#)
- [4] **VLM2Vec-V2: Advancing Multimodal Embedding for Videos, Images, and Visual Documents**
Rui Meng*, Ziyan Jiang*, Ye Liu, Mingyi Su, Xinyi Yang, Yuepeng Fu, Can Qin, Raghuveer Thirukovalluru, Xuan Zhang, Zeyuan Chen, Ran Xu, Caiming Xiong, Yingbo Zhou, Wenhu Chen, Semih Yavuz
[TMLR 2026](#) | [Paper](#) | [Website](#)

2025

- [5] **Universal Retrieval for Multimodal Trajectory Modeling**
Xuan Zhang, Ziyan Jiang, Rui Meng, Yifei Leng, Zhenbang Xiao, Zora Zhiruo Wang, Yanyi Shang, Dehan Kong
[Computer-use Agents @ ICML 2025 \(Workshop\)](#) | [Paper](#)
- [6] **StructEval: Benchmarking LLMs' Capabilities to Generate Structural Outputs**
Jialin Yang, Dongfu Jiang, Lipeng He, Sherman Siu, Yuxuan Zhang, Disen Liao, Zhuofeng Li, Huaye Zeng, Yiming Jia, Haozhe Wang, Benjamin Schneider, Chi Ruan, Wentao Ma, Zhiheng Lyu, Yifei Wang, Yi Lu, Quy Duc Do, Ziyan Jiang, Ping Nie, Wenhu Chen
[TMLR 2025](#) | [Paper](#) | [Website](#)
- [7] **VLM2Vec: Training Vision-Language Models for Massive Multimodal Embedding Tasks**
Ziyan Jiang, Rui Meng, Xinyi Yang, Semih Yavuz, Yingbo Zhou, Wenhu Chen
[ICLR 2025](#) | [Paper](#) | [Website](#)
- [8] **MEGA-Bench: Scaling Multimodal Evaluation to over 500 Real-World Tasks**
Jiacheng Chen, Tianhao Liang, Sherman Siu, Zhengqing Wang, Kai Wang, Yubo Wang, Yuansheng Ni, Wang Zhu, Ziyan Jiang, Bohan Lyu, Dongfu Jiang, Xuan He, Yuan Liu, Hexiang Hu, Xiang Yue, Wenhu Chen
[ICLR 2025](#) | [Paper](#) | [Website](#)
- [9] **Knowledge Enhanced Multi-Domain Recommendations in an AI Assistant Application**
Elan Markowitz, Ziyan Jiang, Fan Yang, Xing Fan, Tony Chen, Greg Ver Steeg, Aram Galstyan
[ICASSP 2025](#) | [Paper](#)

2024

- [10] **LongRAG: Enhancing Retrieval-Augmented Generation with Long-context LLMs**
Ziyan Jiang, Xueguang Ma, Wenhui Chen
[arXiv preprint](#) | [Paper](#)
- [11] **MMLU-Pro: A More Robust and Challenging Multi-Task Language Understanding Benchmark**
Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyan Jiang, Tianle Li, Max Ku, Kai Wang, Alex Zhuang, Rongqi Fan, Xiang Yue, Wenhui Chen
[NeurIPS 2024](#) | [Paper](#) | [Code](#)
- [12] **VideoScore: Building Automatic Metrics to Simulate Fine-grained Human Feedback for Video Generation**
Xuan He, Dongfu Jiang, Ge Zhang, Max Ku, Achint Soni, Sherman Siu, Haonan Chen, Abhranil Chandra, Ziyan Jiang, Aaran Arulraj, Kai Wang, Quy Duc Do, Yuansheng Ni, Bohan Lyu, Yaswanth Narsupalli, Rongqi Fan, Zhiheng Lyu, Yuchen Lin, Wenhui Chen
[EMNLP 2024](#) | [Paper](#) | [Website](#)
- [13] **Semi-Supervised Reward Modeling via Iterative Self-Training**
Yifei He, Haoxiang Wang, Ziyan Jiang, Alexandros Papangelis, Han Zhao
[EMNLP 2024 \(Findings\)](#) | [Paper](#) | [Code](#)
- [14] **Position-Aware Parameter Efficient Fine-Tuning Approach for Reducing Positional Bias in LLMs**
Zheng Zhang, Fan Yang, Ziyan Jiang, Zheng Chen, Zhengyang Zhao, Chengyuan Ma, Liang Zhao, Yang Liu
[arXiv preprint](#) | [Paper](#)
- [15] **RecMind: Large Language Model Powered Agent For Recommendation**
Yancheng Wang, Ziyan Jiang, Zheng Chen, Fan Yang, Yingxue Zhou, Eunah Cho, Xing Fan, Xiaojiang Huang, Yanbin Lu, Yingzhen Yang
[NAACL 2024 \(Findings\)](#) | [Paper](#)

2023

- [16] **Graph Meets LLM: A Novel Approach to Collaborative Filtering for Robust Conversational Understanding**
Zheng Chen*, Ziyan Jiang*, Fan Yang*, Eunah Cho, Xing Fan, Xiaojiang Huang, Yanbin Lu, Aram Galstyan
[EMNLP 2023 \(Industry\)](#) | [Paper](#)
- [17] **PALR: Personalization Aware LLMs for Recommendation**
Fan Yang*, Zheng Chen*, Ziyan Jiang*, Eunah Cho, Xiaojiang Huang, Yanbin Lu
[Gen-IR @ SIGIR 2023 \(Workshop\)](#) | [Paper](#)
- [18] **KG-ECO: Knowledge Graph Enhanced Entity Correction for Query Rewriting**
Jinglun Cai, Mingda Li, Ziyan Jiang, Eunah Cho, Zheng Chen, Yang Liu, Xing Fan, Chenlei Guo
[ICASSP 2023](#) | [Paper](#)

2022

- [19] **PENTATRON: Personalized coNText-Aware Transformer for Retrieval-based coNversational uNderstanding**
Niranjan Uma Naresh*, Ziyan Jiang*, Ankit*, Sungjin Lee, Jie Hao, Xing Fan, Chenlei Guo
[EMNLP 2022 \(Industry\)](#) | [Paper](#)

2021

- [20] **Personalized Search-Based Query Rewrite System for Conversational AI**
Eunah Cho, Ziyan Jiang, Jie Hao, Zheng Chen, Saurabh Gupta, Xing Fan, Chenlei Guo
[NLP4ConvAI @ EMNLP 2021 \(Workshop\)](#) | [Paper](#)
- [21] **#HowYouTagTweets: Learning User Hashtagging Preferences via Personalized Topic Attention**
Yuji Zhang, Yubo Zhang, Chunpu Xu, Jing Li, Ziyan Jiang, Baolin Peng
[EMNLP 2021](#) | [Paper](#) | [Code](#)

2020

- [22] **Black-box Attacks on Spoofing Countermeasures Using Transferability of Adversarial Examples**
Yuekai Zhang, Ziyan Jiang, Jesús Villalba, Najim Dehak
[Interspeech 2020](#) | [Paper](#) | [Demo](#)

2019

- [23] **A Comparative Study on Transformer vs RNN in Speech Applications**
Shigeki Karita, Nanxin Chen, Tomoki Hayashi, Takaaki Hori, Hirofumi Inaguma, Ziyan Jiang, Masao Someki, Nelson Enrique Yalta Soplin, Ryuichi Yamamoto, Xiaofei Wang, Shinji Watanabe, Takenori Yoshimura, Wangyou Zhang
[ASRU 2019](#) | [Paper](#)

Service

Conference Reviewer: NeurIPS (2024–present), ICLR (2024–present), ICML (2026–present), ACL ARR (2023–present), ICASSP (2024–present), EACL (2023).

Journal Reviewer: IEEE Access, TOIS, TMLR.

Workshop Organizer: The First Workshop on Personalized Generative AI (CIKM 2023); The Second Workshop on Generative Information Retrieval (SIGIR 2024).