

Zora Wang Curriculum Vitae

Carnegie Mellon University
Gates Hillman Complex, 4902 Forbes Ave
Pittsburgh, PA 15213

zorazrw.github.io/
zhiruow@cs.cmu.edu

Education

08/2023 – **Carnegie Mellon University**

06/2027 Ph.D. in Computer Science

Language Technologies Institute

Supervisor: Graham Neubig, Daniel Fried

Thesis: Progressive Adaptation of AI Agents

Committee: Daniel Fried, Graham Neubig, Diyi Yang, Tom Mitchell

08/2021 – **Carnegie Mellon University**

08/2023 M.S. in Language Technologies

Language Technologies Institute

GPA: 4.02/4.3

09/2016 – **Beijing Normal University**

06/2020 B.S. in Mathematics & Applied Mathematics

Major GPA: 3.98/4.0

2019 *Chinese University of Hong Kong*, Exchange/Visiting program under full scholarship

2017 *The College of William & Mary*, Summer School under Liyun student scholarship

Employment

05/2025 – **Stanford University**

08/2025 Visiting Student Researcher, Hosted by Diyi Yang

Computer Science Department & Stanford NLP Group

- Study the impact of LM-based agents on human workforce

08/2023 – **Microsoft, GitHub Next Team**

12/2023 Consultant, Hosted by Eddie Aftandilian

- Consult on code language model training and robust open-domain evaluation

05/2022 – **Amazon Alexa AI, Conversational Modeling Team**

08/2022 Research Intern, Hosted by Dilek Hakkani-Tur

- Proposed an adaptive recipe for chatbot training for knowledge-intensive tasks

02/2020 – **Microsoft Research Asia, Data Knowledge Intelligence Team**
08/2021 Research Intern → Assistant Researcher, Hosted by Dongmei Zhang
◦ Designed table foundation model TUTA and shaped into Microsoft Word and Excel

Selected Awards and Fellowship

2025, 2026 **Google PhD Fellowship, Machine Learning and ML Foundations**
2023 **Teng Family Presidential Fellowship at Carnegie Mellon University**
2019 **China National Scholarship for Undergraduate Students**

Publications

* indicates equal contribution, [†] indicate mentoring.

Google Scholar: https://scholar.google.com/citations?hl=en&user=K2hRQgUAAAAJ&view_op=list_works

Preprints

- [C39] Xuehang Guo, **Zora Zhiruo Wang**, Qingyun Wang, Graham Neubig, Xingyao Wang. Enhancing Software Engineering Through Closed-Loop Memory Optimization. 2026. [pdf]
- [C38] Vishakh Padmakumar, Lujain Ibrahim, **Zora Zhiruo Wang**, Jennifer Wang, Q Vera Liao, Diyi Yang. Offloading Score: Measuring AI Reliance Through Counterfactual Workflows. 2026. [pdf]
- [C37] **Zora Zhiruo Wang**, Sanidhya Vijayvargiya, Aspen Chen, Hanmo Zhang, Venu Arvind Arangarajan, Jett Chen, Valerie Chen, Diyi Yang, Daniel Fried, Graham Neubig. How Well Does Agent Development Reflect Real-World Work?. 2026. [pdf]
- [C36] Faria Huq*, **Zora Zhiruo Wang***[†], Zhanqiu Guo, Venu Arvind Arangarajan, Tianyue Ou, Frank Xu, Shuyan Zhou, Graham Neubig, Jeffery P. Bigham. Modeling Distinct Human Interaction in Web Agents. 2026. [pdf]

- [C35] **Zora Zhiruo Wang***, John Yang*, Kilian Lieret*, Alexa Tartaglino, Valerie Chen, Yuxiang Wei, Zijian Wang, Lingming Zhang, Karthik Marasimhan, Ludwig Schmidt, Graham Neubig, Daniel Fried, Diyi Yang. Position: Humans are Missing from AI Coding Agent Research. 2026. [pdf]

Peer-Reviewed Conference Publications

- [C34] Yijia Shao, **Zora Zhiruo Wang**, Neel Ahuja, Yicheng Wang, Bowen Liu, Diyi Yang. CollabSkill: Evaluating Human-Agent Collaboration on Real-World Tasks. In *Proceedings of the Conference of Language Modeling (CoLM)*. 2026. [pdf]
- [C33] Shaofeng Yin, Jiaxin Ge, **Zora Zhiruo Wang**, Chenyang Wang, Xiuyu Li, Michael J. Black, Trevor Darrell, Angjoo Kanazawa, Haiwen Feng. Vision-as-Inverse-Graphics Agent via Interleaved Multimodal Reasoning. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 2026. [pdf]
- [C32] Xuhui Zhou, Valerie Chen, **Zora Zhiruo Wang**, Graham Neubig, Maarten Sap, Xingyao Wang. TOM-SWE: User Mental Modeling For Software Engineering Agents. In *Proceedings of the International Conference on Machine Learning (ICML)*. 2026. [pdf]
- [C31] Sanidhya Vijayvargiya, Aditya Bharat Soni, Xuhui Zhou, **Zora Zhiruo Wang**, Nouha Dziri, Graham Neubig, Maarten Sap. OpenAgentSafety: A Comprehensive Framework for Evaluating Real-World AI Agent Safety. In *Proceedings of the Eleventh International Conference on Learning Representations*. 2026. [pdf]
- [C30] Frank Fangzheng Xu, Yufan Song, Boxuan Li, Yuxuan Tang, Kritanjali Jain, Mengxue Bao, **Zora Zhiruo Wang**, Xuhui Zhou, Zhitong Guo, Murong Cao, Mingyang Yang, Hao Yang Lu, Amaad Martin, Zhe Su, Leander Maben, Raj Mehta, Wayne Chi, Lawrence Jang, Yiqing Xie, Shuyan Zhou, Graham Neubig. TheAgentCompany: Benchmarking LLM Agents on Consequential Real World Tasks. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*. 2025. [pdf]
- [C29] **Zora Zhiruo Wang**, Yijia Shao, Omar Shaikh, Daniel Fried, Graham Neubig, Diyi Yang. How Do AI Agents Do Human Work? Comparing AI and Human Workflows Across Diverse Occupations. 2025. [pdf]
- [C28] Violet Yinuo Han, Jesse T. Gonzalez, Christina Yang, **Zora Zhiruo Wang**, Scott E. Hudson, Alexandra Ion. Towards Unobtrusive Physical AI: Augmenting Everyday Objects with Intelligence and Robotic Movement for Proactive Assistance. In *Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology (UIST)*. 2025. [pdf]

- [C27] Yilin Zhang, Xinran Zhao, **Zora Zhiruo Wang**[†], Chenyang Yang, Jiayi Wei, Tongshuang Wu. cAST: Enhancing Code Retrieval-Augmented Generation with Structural Chunking via Abstract Syntax Tree. *In Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 2025. [pdf]
- [C26] Vishwa Shah, Vishruth Veerendranath, Graham Neubig, Daniel Fried, **Zora Zhiruo Wang**[†]. Exploring the Pre-Conditions for Memory-Learning Agents. *In Proceedings of the Scaling Self-Improving Foundation Models Without Human Supervision Workshop at ICLR*. 2025. [pdf]
- [C25] Boyuan Zheng, Michael Y. Fatemi, Xiaolong Jin, **Zora Zhiruo Wang**, Apurva Gandhi, Yueqi Song, Yu Gu, Jayanth Srinivasa, Gaowen Liu, Graham Neubig, Yu Su. SkillWeaver: Web Agents can Self-Improve by Discovering and Honing Skills. 2025. [pdf]
- [C24] **Zora Zhiruo Wang**, Apurva Gandhi, Graham Neubig. Inducing Programmatic Skills for Agentic Tasks. *In Proceedings of the Conference of Language Modeling (CoLM)*. 2025. [pdf]
- [C23] Eduardo Treviño, Hugo Contant, James Ngai, Graham Neubig, **Zora Zhiruo Wang**[†]. Benchmarking Failures in Tool-Augmented Language Models. *In Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics*. 2025. [pdf]
- [C22] Faria Huq, **Zora Zhiruo Wang**[†], Frank F. Xu, Tianyue Ou, Shuyan Zhou, Jeffery P. Bigham, Graham Neubig. CowPilot: A Framework for Autonomous and Human-Agent Collaborative Web Navigation. *In Proceedings of the Annual Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics (Demo Track)*. 2025. [pdf]
- [C21] Jiaxin Ge*, **Zora Zhiruo Wang**^{*†}, Xuhui Zhou, Yi-Hao Peng, Sanjay Subramanian, Qinyue Tan, Maarten Sap, Alane Suhr, Daniel Fried, Graham Neubig, Trevor Darrell. AutoPresent: Designing Structured Visuals from Scratch. *In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2025. [pdf]
- [C20] **Zora Zhiruo Wang**, Jiayuan Mao, Daniel Fried, Graham Neubig. Agent Workflow Memory. *In Proceedings of the International Conference on Machine Learning (ICML)*. 2025. [pdf]
- [C19] Siddhant Waghjale, Vishruth Veerendranath, **Zora Zhiruo Wang**[†], Daniel Fried. ECCO: Can We Improve Model-Generated Code Efficiency Without Sacrificing Functional Correctness. *In Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 2024. [pdf]

- [C18] **Zora Zhiruo Wang***, Akari Asai*, Xinyan Yu, Frank F. Xu, Yiqing Xie, Graham Neubig, Daniel Fried. CodeRAG-Bench: Can Retrieval Augment Code Generation? In *Proceedings of the Annual Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics (NAACL)*. 2024. [pdf]
- [C17] **Zora Zhiruo Wang**, Zhoujun Cheng, Hao Zhu, Daniel Fried, Graham Neubig. What Are Tools Anyway? A Survey from the Language Model Perspective. In *Proceedings of the 1st Conference on Language Modeling (CoLM)*. 2024. [pdf]
- [C16] Jennifer Hsia, Afreen Shaikh, **Zora Zhiruo Wang[†]**, Graham Neubig. RAGGED: Towards Informed Design of Retrieval Augmented Generation Systems. *NeurIPS Adaptive Foundation Models (AFM) Workshop*. 2024. [pdf]
- [C15] **Zora Zhiruo Wang**, Daniel Fried, Graham Neubig. TroVE: Inducing Verifiable and Efficient Toolboxes for Solving Programmatic Tasks. *International Conference on Machine Learning (ICML)*. 2024. [pdf]
- [C14] **Zora Zhiruo Wang**, Jun Araki, Zhengbao Jiang, Md Rizwan Parvez, Graham Neubig. Learning to Filter Context for Retrieval-Augmented Generation. 2023. [pdf]
- [C13] Yihao Cao, Shuyi Chen, Ryan Liu, **Zora Zhiruo Wang[†]**, Daniel Fried. API-Assisted Code Generation for Question Answering on Varied Table Structures. In *Proceedings of the Empirical Methods in Natural Language Processing*. 2023. [pdf]
- [C12] I-Chun Chern, **Zora Zhiruo Wang[†]**, Sanjan Das, Bhavuk Sharma, Pengfei Liu, Graham Neubig. Improving Factuality of Abstractive Summarization via Contrastive Reward Learning. In *Proceedings of the 3rd Workshop on Trustworthy Natural Language Processing*. 2023. [pdf]
- [C11] Lijun Yu, Yong Cheng, **Zora Zhiruo Wang**, Vivek Kumar, Wolfgang Macherey, Yanping Huang, David A Ross, Irfan Essa, Yonatan Bisk, Ming-Hsuan Yang, Kevin Murphy, Alexander G Hauptmann, Lu Jiang. SPAE: Semantic Pyramid AutoEncoder for Multimodal Generation with Frozen LLMs. In *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*. 2023. [pdf]

- [C10] Raymond Li, Loubna Ben Allal, Yangtian Zi, Niklas Muennighoff, Denis Kocetkov, Chenghao Mou, Marc Marone, Christopher Akiki, Jia Li, Jenny Chim, Qian Liu, Evgenii Zheltonozhskii, Terry Yue Zhuo, Thomas Wang, Olivier Dehaene, Mishig Davaadorj, Joel Lamy-Poirier, João Monteiro, Oleh Shliazhko, Nicolas Gontier, Nicholas Meade, Armel Zebaze, Ming-Ho Yee, Logesh Kumar Umapathi, Jian Zhu, Benjamin Lipkin, Muhtasham Oblokulov, **Zora Zhiruo Wang**, ZRudra Murthy, Jason Stillerman, Siva Sankalp Patel, Dmitry Abulkhanov, Marco Zocca, Manan Dey, Zhihan Zhang, Nour Fahmy, Urvashi Bhattacharyya, Wenhao Yu, Swayam Singh, Sasha Luccioni, Paulo Villegas, Maxim Kunakov, Fedor Zhdanov, Manuel Romero, Tony Lee, Nadav Timor, Jennifer Ding, Claire Schlesinger, Hailey Schoelkopf, Jan Ebert, Tri Dao, Mayank Mishra, Alex Gu, Jennifer Robinson, Carolyn Jane Anderson, Brendan Dolan-Gavitt, Danish Contractor, Siva Reddy, Daniel Fried, Dzmitry Bahdanau, Jernite. StarCoder: May the Source be with You! *In Proceedings of Transactions on Machine Learning Research (TMLR)*. 2023. [pdf]
- [C9] Shuyan Zhou, Uri Alon, Frank F. Xu, **Zora Zhiruo Wang**, Zhengbao Jiang, Graham Neubig. DocPrompting: Generating Code by Retrieving the Docs. *In Proceedings of the Eleventh International Conference on Learning Representations (ICLR)*. 2022. [pdf]
- [C8] **Zora Zhiruo Wang***, Grace Cuenca*, Shuyan Zhou, Frank F. Xu, Graham Neubig. MCoNaLa: A Benchmark for Code Generation from Multiple Natural Languages. *In Findings of the Association for Computational Linguistics (EACL)*. 2023. [pdf]
- [C7] **Zora Zhiruo Wang**, Shuyan Zhou, Daniel Fried, Graham Neubig. Execution-Based Evaluation for Open-Domain Code Generation. *In Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 2022. [pdf]
- [C6] **Zora Zhiruo Wang**, Zhengbao Jiang, Eric Nyberg, Graham Neubig. Table Retrieval May Not Necessitate Table-Specific Model Design. *In Proceedings of the Workshop on Structured and Unstructured Knowledge Integration (SUKI)*. 2022. [pdf]
- [C5] Zhengbao Jiang, Luyu Gao, Jun Araki, Haibo Ding, **Zora Zhiruo Wang**, Jamie Callen, Graham Neubig. Retrieval As Attention: End-to-End Learning of Retrieval and Reading within a Single Transformer. *In Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 2022. [pdf]
- [C4] Zhoujun Cheng*, Haoyu Dong*, **Zora Zhiruo Wang***, Ran Jia, Jiaqi Guo, Yan Gao, Shi Han, Jian-Guang Lou, Dongmei Zhang. HiTab: A Hierarchical Table Dataset for Question Answering and Natural Language Generation. *In Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (ACL)*. 2021. [pdf]

- [C3] **Zora Zhiruo Wang***, Haoyu Dong*, Ran Jia, Jia Li, Zhiyi Fu, Shi Han, Dongmei Zhang. TUTA: Tree-Based Transformers for Generally Structured Table Pre-Training. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 2021. [pdf]
- [C2] Weijie Liu, Peng Zhou, Zhe Zhao, **Zora Zhiruo Wang**, Haotang Deng, Qi Ju. FastBERT: A Self-Distilling BERT with Adaptive Inference Time. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*. 2020. [pdf]
- [C1] Weijie Liu, Peng Zhou, Zhe Zhao, **Zora Zhiruo Wang**, Qi Ju, Haotang Deng, Ping Wang. K-BERT: Enabling Language Representation with Knowledge Graph. In *Proceedings of the AAAI Conference on Artificial Intelligence*. 2020. [pdf]

Workshop Publications

- [W3] Terry Yue Zhao, **Zora Zhiruo Wang**, Wen-Ding Li, Giovanni Zappella, Qian Liu, Zijian Wang. Deep Learning for Code: Towards Human-Centered Coding Agents. *International Conference on Machine Learning (ICML)*. 2026. [pdf]
- [W2] Zijian Wang, Giovanni Zappella, Qian Liu, **Zora Zhiruo Wang**, Wen-Ding Li, Wasi Uddin Admad, Binyuan Hui. Deep Learning for Code in the Agentic Era. In *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*. 2025. [pdf]
- [W1] Zijian Wang, Ying Sheng, Giovanni Zappella, Qian Liu, Devjeet Raj Roy, Gabriel Orlanski, **Zora Zhiruo Wang**, Wen-Ding Li. Emergent Possibilities and Challenges in Deep Learning for Code. In *Proceedings of the International Conference on Learning Representations (ICLR)*. 2025. [pdf]

Tutorials

- [T2] **Zora Zhiruo Wang**, Yijia Shao, David Nguyen, Diyi Yang. Future of Work in the Age of LLMs. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (ACL)*. 2026. [pdf]
- [T1] Haoyu Dong, **Zora Zhiruo Wang**. Large Language Models for Tabular Data: Progress and Future Directions. In *Proceedings of the 47th International ACM SIGIR*. 2024. [pdf]

Teaching

TA at CMU CS Courses

- Spring 2024* **11-891 Neural Code Generation**
Language Technologies Institute, School of Computer Science, Carnegie Mellon University
Instructor: Daniel Fried, Sean Welleck
- Fall 2023* **11-711 Advanced Natural Language Processing**
Language Technologies Institute, School of Computer Science, Carnegie Mellon University
Instructor: Daniel Fried
- Spring 2023* **11-711 Advanced Natural Language Processing**
Language Technologies Institute, School of Computer Science, Carnegie Mellon University
Instructor: Graham Neubig
- Fall 2021* **11-631 Data Science Seminar**
Language Technologies Institute, School of Computer Science, Carnegie Mellon University
Instructor: Eric Nyberg

Invited Lecturer at Courses

- Spring 2026* **11-711 Advanced Natural Language Processing**
“Language Model-Based Agents”
Language Technologies Institute, School of Computer Science, Carnegie Mellon University
Instructor: Sean Welleck
- Spring 2026* **11-766 Large Language Model Applications**
“AI Coding Agents”
Language Technologies Institute, School of Computer Science, Carnegie Mellon University
Instructor: Fernando Diaz, Daphne Ippolito
- Fall 2025* **49-594 Human-Centered Natural Language Processing**
“Language Model-Based Agents”
Human Computer Interaction Institute, School of Computer Science, Carnegie Mellon University
Instructor: Sherry Tongshuang Wu

Spring **42-698 Clinical Translation of AI**

2024 "Language Model-Based Agents"

Department of Biomedical Engineering, Carnegie Mellon University

Instructor: Newell Washburn

Spring **11-891 Neural Code Generation**

2024 "Evaluation & Benchmarks for Code Generation"

Language Technologies Institute, School of Computer Science, Carnegie Mellon University

Instructor: Daniel Fried, Sean Welleck

Fall 2023 **11-711 Advanced Natural Language Processing**

"Evaluation & Benchmarks for Code Generation"

Language Technologies Institute, School of Computer Science, Carnegie Mellon University

Instructor: Daniel Fried

Spring **11-711 Advanced Natural Language Processing**

2023 "Tool Use and Language Agents"

Language Technologies Institute, School of Computer Science, Carnegie Mellon University

Instructor: Graham Neubig

Mentoring

01/2024 – **Public Office Hour**

Present Virtual & In-person office hours open to the public (four 15min 1-on-1 sessions every week)

More than 50 students joined to discuss research and Ph.D. applications

PhD Students

08/2024 – **Faria Huq**

Present CMU HCII Ph.D. Student (advised by Hirokazu Shirado)

11/2025 – **Yucheng Jiang**

06/2026 Stanford NLP Ph.D. Student (advised by Monica Lam)

09/2025 – **Yijia Shao**

05/2025 Stanford NLP Ph.D. Student (advised by Diyi Yang)

01/2025 – **Jiaxin Ge**

09/2025 Berkeley BAIR Ph.D. Student (advised by Trevor Darrell)

Predctoral Students

- 08/2025 – **Aspen Chen**
Present CMU MIIS Master's Student
- 08/2025 – **Jett Chen**
Present CMU Undergraduate Student
- 08/2025 – **Sanidhya Vijayvargiya**
03/2026 CMU Master's Student → Research Engineer at Samsung Research
- 08/2025 – **Venu Arvind Arangarajan**
12/2025 CMU Master's Student → Machine Learning Engineer at JPMorgan
- 09/2024 – **Xinyue Liu**
05/2025 CMU Master's Student → Stony Brook Ph.D.
- 08/2024 – **Xueying Jia**
05/2025 CMU Master's Student → Google Deepmind
- 01/2024 – **Yibo Peng**
01/2025 CMU Master's Student → Research Assistant at UC Berkeley
- 08/2023 – **Vishruth Veerendranath**
05/2025 CMU Master's Student → Snowflake AI
- 05/2024 – **Eduardo Treviño**
10/2024 CMU Master's Student → Member of Technical Staff at xAI
- 05/2024 – **Hugo Contant**
10/2024 CMU Undergraduate Student → Research Assistant at MLD, CMU
- 08/2023 – **Siddhant Waghjale**
05/2024 CMU Master's Student → Applied Assistant at Mistral AI
- 08/2023 – **Afreen Shaikh**
05/2024 CMU Master's Student → AI/ML Applied Scientist at SAP Labs
- 01/2023 – **Shuyi Chen**
05/2023 CMU Master's Student → CMU ML&OR Ph.D.
- 01/2023 – **Ryan Liu**
05/2023 CMU Master's Student → Princeton Ph.D.
- 01/2023 – **Ethan Chen**
05/2023 CMU Master's Student → Shanghai Jiaotong University (SJTU) Ph.D.
- 01/2023 – **Yihan Cao**
05/2023 CMU Master's Student → Machine Learning Engineer at LinkedIn

Service

Academic Service within the University

2024 – Ph.D. Admission Committee Member

2026 Reviewed Ph.D. application materials for applicants to the Language Technologies Institute at CMU

2024 – Faculty Hiring Committee, Student Member

2025 Provided student feedback on faculty candidate job talks and 1-on-1 chats

2024 – Belonging and Engagement Committee (formerly DEIB), Student Representative

2025 Represented student perspectives at weekly DEI meetings of the department

2023 – Faculty Hiring Committee, Student Chair

2024 Managed student reviewing process of faculty candidate talks and chats

05/2024 Head organizer of the first CMU Agent Workshop

Invited keynote speakers, managed the paper submission, reviewing, and presentation processes

08/2022 Organizer of the LTI Student Research Symposium

Organized the paper reviewing and presentation of student research at the Language Technologies Institute

Research Community

07/2026 Workshop Organizer

Deep Learning for Code: Towards Human-Centered Coding Agents (ICML 2026)

07/2026 Tutorial Instructor

Future of Work in the Age of LLMs (ACL 2026)

12/2025 Workshop Organizer

Deep Learning for Code in the Agentic Era (NeurIPS 2025)

04/2025 Tutorial Instructor

Language Model-Based Agents (CMU Agent Workshop 2025)

04/2025 Workshop Organizer

Emergent Possibilities and Challenges in Deep Learning for Code (ICLR 2025)

- 07/2024 **Tutorial Instructor**
Large Language Models for Tabular Data: Progress and Future Directions (SIGIR 2024)
- 10/2025 **Area Chair**
NeurIPS 2025, Multi-Turn LLM Agent Workshop
- 2020 – **Reviewer**
Present ACL Rolling Review (ARR) / ACL / EACL / EMNLP, 2022-present
NeurIPS 2025–2026, NeurIPS Dataset and Benchmark Track 2025–2026
ICML 2025, ICML Position Paper Track 2026
CoLM 2025–2026
CVPR 2026, ECCV 2026
CHI 2026
AAAI 2023–2026
Workshops: TRL at NeurIPS 2023, Student Research Workshop at ACL 2023, TLLM 2023 at SIGDIAL 2023, AGI at ICLR 2024, LLM for Agents at ICLR 2024

Invited Talks

- 05/2026 Berkeley Artificial Intelligence Research (BAIR) Lab (hosted by Alane Suhr)
- 05/2026 Augmented Mind Podcast, Stanford University (hosted by Yijia Shao)
- 05/2026 University of Washington NLP Group (hosted by Yulia Tsvetkov)
- 05/2026 Kungfu.AI (hosted by Steve Kramer)
- 04/2026 New York University NLP Group (hosted by Greg Durrett)
- 03/2026 Schmidt Sciences Trustworthy AI Convening
- 03/2026 Toronto Data Workshop (hosted by Rohan Alexander)
- 03/2026 Lightning Research Talk at LTI Open House
- 03/2026 Snorkel AI (hosted by Ramya Ramakrishnan)
- 12/2025 AI-TIME Youth Talk (hosted by Yicheng He)
- 10/2025 WInE Lab at HCI Institute at CMU (hosted by Sherry Tongshuang Wu)
- 10/2025 Machine Learning Department at CMU (hosted by Ruslan Salakhutdinov)
- 08/2025 Stanford NLP Seminar (hosted by Diyi Yang)

07/2025 Renmin University of China (RUC) (hosted by Yunhai Wang)

06/2025 Pareto AI (hosted by Wenjia Hu)

06/2025 Orby AI (hosted by Peng Qi)

04/2025 Fujitsu Research

03/2025 Lightning Research Talk at LTI Open House

03/2025 AAAI Workshop on Advancing LLM-Based Multi-Agent Collaboration (hosted by Raphael Shu)

03/2025 AAAI Web Agent Revolution Workshop (hosted by Avi Yaeli)

12/2024 Tsinghua University, Institute for Interdisciplinary Information Sciences (hosted by Yi Wu)

11/2024 University of Illinois Urbana-Champaign, ISE Seminar (hosted by Lingming Zhang)

09/2024 CAMEL.AI (hosted by Guohao Li)

05/2024 Foundation and LAnguage Model (FLAME) Seminar (hosted by Zhiqing Sun)

02/2024 LLM as Agent Seminar at CMU Catalyst's Group (hosted by Kevin Kuo)

04/2023 JetBrains Research Team (hosted by Misha Evtikhiev)