YUCHEN ZENG

Bellevue, WA

yuchen.zeng.1998@gmail.com • Homepage • Google Scholar

CURRENT RESEARCH INTEREST

Agents, Reasoning Models, and LLM Post-Training

WORK EXPERIENCE

Senior Researcher, Microsoft AI, USA *08/2026 - Current*
Focus: LLM Post-Training

Senior Researcher, Microsoft Research AI Frontiers, USA *09/2025 - 08/2026*
Mentored by Prof. Dimitris Papailiopoulos

Research Intern, Meta, USA *09/2024 - 12/2024*
Mentored by Dr. Yipin Zhou

Research Intern, Microsoft, USA *06/2024 - 09/2024*
Mentored by Dr. Andreas Mueller

Research Intern, MIT-IBM Watson AI Lab, USA *05/2022 - 09/2022*
Mentored by Dr. Mikhail Yurochkin and Prof. Justin Solomon

EDUCATION

Ph.D in Computer Science, University of Wisconsin-Madison, USA *09/2020 - 06/2025*
Advisor: Prof. Kangwook Lee

M.S. in Statistics, University of Wisconsin-Madison, USA *08/2018 - 05/2020*
Advisor: Prof. Miaoyan Wang

B.S. in Statistics, Nankai University, China *09/2015 - 06/2019*
Advisor: Prof. Zhaojun Wang

SELECTED PUBLICATIONS (⋆ EQUAL CONTRIBUTION; † EQUAL ADVISING)

Yuchen Zeng, Giorgio Severi, Blake Bullwinkel, Eric Horvitz, Dimitris Papailiopoulos (2026). "Lazarus, Come Forth: Reprogramming Dead Neurons to Backdoor CLI Agents". *Work in progress.*

Wonjun Kang, Kevin Galim, Seunghyuk Oh, Minjun Kang, Sanghyun Park, Donghoon Kim, Minjae Lee, Minseo Kim, Rishabh Tiwari, **Yuchen Zeng**, Hyung Il Koo, Kangwook Lee (2026). "AsyncOPD: How Stale Can On-Policy Distillation Be?". In *Neural Information Processing Systems (NeurIPS 2026)*.

Yuchen Zeng, Dimitris Papailiopoulos (2026). "You Don't Need to Run Every Eval". *Under Review (ArXiv).*

Vasilis Kontonis, **Yuchen Zeng**, Shivam Garg, Lingjiao Chen, Hao Tang, Ziyan Wang, Ahmed Awadallah, Eric Horvitz, John Langford, Dimitris Papailiopoulos (2026). "Memento: Teaching LLMs to Manage Their Own Context". In *Conference on Language Modeling (COLM 2026)*.

Xuhan Tong, **Yuchen Zeng**†, Jiawei Zhang† (2026). "Demonstrations, CoT, and Prompting: A Theoretical Analysis of ICL". *Under Review (ArXiv).*

Yuchen Zeng⋆, Shuibai Zhang⋆, Wonjun Kang⋆, Shutong Wu, Lynnix Zou, Ying Fan, Heeju Kim, Ziqian Lin, Jungtaek Kim, Hyung Il Koo, Dimitris Papailiopoulos, Kangwook Lee (2026). "ReJump: A Tree-Jump Representation for Analyzing and Improving LLM Reasoning". In *International Conference on Machine Learning (ICML 2026)*.

Wonjun Kang*, Kevin Galim*, Seunghyuk Oh*, Minjae Lee, **Yuchen Zeng**, Shuibai Zhang, Coleman Hooper, Yuezhou Hu, Hyung Il Koo, Nam Ik Cho, Kangwook Lee (2026). "ParallelBench: Understanding the Trade-offs of Parallel Decoding in Diffusion LLMs". In *International Conference on Learning Representations (ICLR 2026)*.

Tong Xie*, Yuwei Wan*, Yixuan Liu, **Yuchen Zeng**, Wenjie Zhang, Chunyu Kit, Dongzhan Zhou, Bram Hoexter (2025). "DARWIN 1.5: Large Language Models as Materials Science Adapted Learners". In *AI for Science 2025 (AI4X 2025)*.

Wonjun Kang*, Kevin Galim*, **Yuchen Zeng***, Minjae Lee, Hyung Il Koo, Nam Ik Cho (2025). "State-offset Tuning: State-based Parameter-Efficient Fine-Tuning for State Space Models". In Proceedings of the 63rd Annual Meeting of the *Association for Computational Linguistics (ACL 2025)*. (Volume 2: Short Papers), pages 474–487.

Yuchen Zeng*, Tuan Dinh*, Wonjun Kang, Andreas Mueller (2025). "TabFlex: Scaling Tabular Learning to Millions with Linear Attention". Proceedings of the 42nd *International Conference on Machine Learning (ICML 2025)*.

Kevin Galim*, Wonjun Kang*, **Yuchen Zeng***, Hyung Il Koo, Kangwook Lee (2025). "Parameter-Efficient Fine-Tuning of State Space Models". Proceedings of the 42nd *International Conference on Machine Learning (ICML 2025)*.

Yuchen Zeng*, Wonjun Kang*, Yicong Chen, Hyung Il Koo, Kangwook Lee (2024). "Can MLLMs Perform Text-to-Image In-Context Learning?". In *First Conference on Language Modeling (COLM 2024)*.

Yuchen Zeng, Kangwook Lee (2024). "The Expressive Power of Low-Rank Adaptation". In *International Conference on Learning Representations (ICLR 2024)*.

Yuchen Zeng, Hongxu Chen, Kangwook Lee (2023). "Federated Learning with Local Fairness Constraints". 2023 IEEE *International Symposium on Information Theory (ISIT 2023)*, pp. 1937-1942.

Ozgur Guldogan*, **Yuchen Zeng***, Jy-yong Sohn, Ramtin Pedarsani, Kangwook Lee (2023). "Equal Improvability: A New Fairness Notion Considering the Long-term Impact". In *International Conference on Learning Representations (ICLR 2023)*.

Tuan Dinh*, **Yuchen Zeng***, Ruisu Zhang, Ziqian Lin, Michael Gira, Shashank Rajput, Jy-yong Sohn, Dimitris Papailiopoulos, Kangwook Lee (2022). " LIFT: Language-Interfaced FineTuning for Non-Language Machine Learning Tasks". Advances in *Neural Information Processing Systems 36 (NeurIPS 2022)*.

Miaoyan Wang, **Yuchen Zeng** (2019). "Multiway clustering via tensor block models". Advances in *Neural Information Processing Systems 32 (NeurIPS 2019)*.

INDEPENDENT PROJECT

Persistent Personal Assistant (2026). Designed and built a cross-domain personal AI chief of staff that combines persistent personal context, relationship-aware communication, domain-specific judgment, and tool routing across career, health, finance, research, and daily planning. [Project Page] [Live Demo]

SERVICES & AWARDS

COLM 2026 Oral Spotlight for Memento (24 of 856 accepted papers)
ICML 2025 Spotlight for TabFlex (top 2.6% of submissions)
Recipient of 2024 COLM DEI Travel Scholarship
Selected to participate in Future Leaders Summit 2024
Organizer of MLOPT Idea Seminar
Conference reviewer: NeurIPS, ICML, ICLR, TMLR, ICCV, ACM MM
NeurIPS Travel Award, 2019