

Boxuan Li

boxuanli@microsoft.com | 412-251-1251 | <https://li-boxuan.com>

PUBLICATIONS

- [1] Frank F. Xu*, Yufan Song*, **Boxuan Li***, Yuxuan Tang, Kritanjali Jain, Mengxue Bao, Zora Z. 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 (2024). TheAgentCompany: Benchmarking LLM Agents on Consequential Real World Tasks. Under Review at ICML 2025. * stands for equal contribution.
- [2] Xingyao Wang, **Boxuan Li**, Yufan Song, Frank F. Xu, Xiangru Tang, Mingchen Zhuge, Jiayi Pan, Yueqi Song, Bowen Li, Jaskirat Singh, Hoang H. Tran, Fuqiang Li, Ren Ma, Mingzhang Zheng, Bill Qian, Yanjun Shao, Niklas Muennighoff, Yizhe Zhang, Binyuan Hui, Junyang Lin, Robert Brennan, Hao Peng, Heng Ji, Graham Neubig (2024). OpenHands: An Open Platform for AI Software Developers as Generalist Agents. The Thirteenth International Conference on Learning Representations (ICLR 2025).
- [3] **Boxuan Li**, Reynold Cheng, Jiafeng Hu, Yixiang Fang, Min Ou, Ruibang Luo, Kevin Chen-Chuan Chang, Xuemin Lin (2020). MC-Explorer: Analyzing and Visualizing Motif-Cliques on Large Networks. 36th IEEE International Conference on Data Engineering (ICDE 2020) Demo Track.
- [4] **Boxuan Li**, Martin Carrington, Peter Marbach (2020). Stable community structures and social exclusion. International Conference on Social Informatics.
- [5] Cheng Wang, Xusheng Chen, Weiwei Jia, **Boxuan Li**, Haoran Qiu, Shixiong Zhao, Heming Cui (2018). PLOVER: Fast, Multi-core Scalable Virtual Machine Fault-tolerance. 15th USENIX Symposium on Networked Systems Design and Implementation (NSDI 18).

RESEARCH EXPERIENCE

The Agent Company | Cofounder Aug 2024 – Present
TheAgentCompany is the first benchmark that examines AI's ability to complete real-world consequential tasks. I designed a reproducible, extensible evaluation framework from scratch, led a team of 10+ software engineers to complete the coding part, and co-authored the paper. It is under review at ICML 2025.

OpenHands | Cofounder Apr 2024 – Present
OpenHands is the most popular open-source coding agent. I am one of the early cofounders and an active maintainer of OpenHands. It has been downloaded by more than 800k times, reported by multiple media, and cited by 118+ academic papers.

JanusGraph | Maintainer Sep 2019 – Present
Graph databases are a fundamental building block of AI applications to leverage the private domain knowledge. I served as Technical Steering Committee and led the development of JanusGraph, the most popular open-source distributed graph database. It is widely used by many companies including Microsoft, Apple, Netflix, Uber, RedHat, eBay, and Goldman Sachs. It has been downloaded by more than 500k times, and cited by 27+ academic papers.

WORK EXPERIENCE

Microsoft Redmond, WA
Senior Software Engineer Feb 2023 – Present
I work in [Scope](#) query optimizer team. Scope is the language for Microsoft's internal big data platform. Scope carries the company, processing Exabytes of data and nearly a million jobs every day.
I lead the performance improvement efforts, including building sophisticated statistics model, tuning TPC-DS benchmarks, and developing state-of-the-art algorithms for join operations. My work has led to millions of dollars of savings per year for the company.

Goldman Sachs Hong Kong
Software Engineer Jul 2019 – Jul 2021
My team builds a massive live knowledge graph that enables company-wide business intelligence, machine learning, monitoring, and data governance. My contribution includes:

- Graph query optimizations that reduce average system latency by 50%.
- An Infrastructure as Code (IAC) solution that reduces data governance labor by 90%.
- A Spark streaming pipeline to ingest ~10 million production telemetries daily.

EDUCATION

Carnegie Mellon University, School of Computer Science Pittsburgh, PA
Master of Computational Data Science (Systems track), GPA: 4.0/4.0 Jul 2021 - Dec 2022

- Teaching Assistant: Distributed Systems (15-640), Fall 2022

The University of Hong Kong Hong Kong, HK
Bachelor of Engineering in Computer Science, GPA: 3.87/4.30 (Top 5%) Sep 2015 - May 2019

- Research Assistant: Systems & Database Group, 2018-2020