

Ron Yifeng Wang

Address: 353 Jane Stanford Way, Stanford, CA
Email: ronyw AT cs DOT stanford DOT edu

Personal Site: <https://cs.stanford.edu/~ronyw>

Education

Stanford University, Stanford, CA

Sep 2024 - Expected Jun 2026

M.S. in Computer Science, GPA: 3.85/4.00

Coursework: (AI) Deep Reinforcement Learning, Graph Learning, Compound AI Systems, Self-Improving AI Agents; (Sys) OS, Programming Languages, Parallel Computing, Computer Networking

University of California, Berkeley, Berkeley, CA

Aug 2020 - May 2024

B.A. in Computer Science and Data Science, B.A. in Economics, GPA: 3.93/4.00

Coursework: (AI) Machine Learning Theory, Computer Vision, Natural Language Processing, Optimization Models; (Sys) Machine Structures, Database Systems, Computer Security, Software Engineering

Research Experience

Multi-scale Architectures & Systems Team (MAST), Stanford University

Jun 2025 - Jan 2026

Latency-Optimizing CUA Planning and Scheduling

Mentor: Caleb Winston; PI: Prof. Christos Kozyrakis

- Designed, implemented, and evaluated a latency-aware scheduler for Computer-Use Agent (CUA) workloads that leverages parallelism and request hedging.
- Maintained and extended the OSS project with new features, bug fixes, and benchmarks. Open-sourced at [Blast](#).

Sky Computing Lab, UC Berkeley

Dec 2023 - Sep 2025

The Streaming Batch Model for Efficient and Fault-Tolerant Heterogeneous Execution

Mentor: Dr. Frank Luan; PI: Prof. Stephanie Wang

- Evaluated Ray Data's Streaming Batch Model against traditional batch/streaming systems and ML dataloaders, measuring scalability, resource elasticity, and fault tolerance. Open-sourced at [Ray Data](#) and [Ray Data Eval](#).
- Identified and resolved system performance bottlenecks, improving Ray Data image ingestion throughput by 22%. Prototyped an autoscaler that dynamically adjusts executor count to maximize GPU utilization.
- Managed Terraform infrastructure for deploying Ray Data and benchmark baseline clusters at scale.

Stanford Artificial Intelligence Laboratory (SAIL), Stanford University

Jan 2025 - Jun 2025

Containerized Evaluation Infrastructure for AI Benchmarks

Mentor: Andy K. Zhang; PI: Prof. Percy Liang

- Built core benchmarking infrastructure: a secure, isolated execution sandbox with Docker-in-Docker and Nginx-based service orchestration. Open-sourced at [BountyBench](#).
- Authored 7 of 40 evaluation tasks; validated correctness and reproducibility of the automated grading harness across the full benchmark.

Work Experience

Plaid Inc., New York, NY

June 2025 - Sep 2025

Software Engineer Intern, Core Infrastructure (return offer extended)

Mentor: Niko Ramirez

- Designed and implemented a service to automatically detect and mitigate unreachable or failed CI/CD workers, reducing worker-level MTTR from several hours to under 10 minutes. Deployed as a Kubernetes CronJob monitoring hundreds of instances. Owned the spec, security review, implementation, rollout, and knowledge transfer.

Publications (*equal contribution)

Preprints

- [1] **Agent JIT Compilation for Latency-Optimizing Computer-Use Agent Planning and Scheduling.**
Caleb Winston, [Ron Yifeng Wang](#), Azalia Mirhoseini, Christos Kozyrakis.
Under review, 2026

- [2] **The Streaming Batch Model for Efficient and Fault-Tolerant Heterogeneous Execution.**
Frank Sifei Luan*, [Ron Yifeng Wang*](#), Yile Gu, Ziming Mao, Charlotte Lin, Amog Kamsetty, Hao Chen, Cheng Su, Balaji Veeramani, Scott Lee, SangBin Cho, Clark Zinzow, Eric Liang, Ion Stoica, Stephanie Wang.
Under review, 2025 [[Paper](#)].

Conference Proceedings

- [1] **BountyBench: Dollar Impact of AI Agent Attackers and Defenders on Real-World Cybersecurity Systems.**
Andy K Zhang, Joey Ji, Celeste Menders, Riya Dulepet, Thomas Qin, [Ron Yifeng Wang](#), Junrong Wu, Kyleen Liao, Jiliang Li, Jinghan Hu, Sara Hong, ..., Dawn Song, Dan Boneh, Daniel E Ho, Percy Liang.
NeurIPS (Datasets & Benchmarks Track), 2025 [[Paper](#)].

Teaching

TA of Intro to Big Data Systems, Stanford University. Supervisor: <i>Prof. Shiva Shivakumar</i>	Sep 2025 - Dec 2025
TA of Data Engineering, UC Berkeley. Supervisor: <i>Prof. Lisa Yan</i>	Aug 2023 - Dec 2023
TA, Tutor, Lab Assistant of Foundations of Data Science, UC Berkeley	Jun 2021 - Dec 2022

Professional Service

- Artifact Evaluation Committee: FAST 2025

Awards

Microsoft for Startups Founders Hub (25K USD)	Nov 2023
EECS Honors, UC Berkeley	Jul 2023