Yuxiao Wen

New York, NY · yuxiaowen@nyu.edu · (347) 833-8643 · yw99.github.io

EDUCATION

NEW YORK UNIVERSITY **New York, NY**

Ph.D. in Computer Science · GPA: 3.95/4.0 Expected 2027

Advisors: Yanjun Han and Zhengyuan Zhou · MacCracken fellowship · NYU AI Safety fellowship

M.Phil. in Computer Science Jan 2026

B.A. in Honors Mathematics; B.A. in Computer Science · GPA: 3.88/4.0 May 2021

Magna cum laude · NYU CS Outstanding Performance Award (2 recipients/year)

RESEARCH PROFILE

- Reinforcement learning and sequential decision-making, with research spanning **policy evaluation and improvement**, **contextual bandits**, **causal inference**, **and LLM post-training**.
- Develop statistically principled, sample-efficient learning and experimentation methods, with applications to **online advertising**, **recommendation**, **pricing**, **and agentic systems**.

WORK EXPERIENCE

Amazon Web Services — Applied Scientist Intern May 2025 - Aug 2025 | Seattle, WA

- Developed reinforcement-learning for dynamic floor pricing in online advertising for **>1M** requests/min.
- Improved publisher's revenue by **>3%** through adaptive floor optimization.

Spotify — Machine Learning Engineer Intern June 2024 - Aug 2024 | New York, NY

- Designed and implemented an adaptive exploration strategy serving **>5M** users/day, improving user experience while preserving exploration data quality.
- Incorporated a contextual-bandit algorithm, achieving **~20%** relative increase in click-through rate.

Amazon Web Services — Software Development Engineer Intern June 2021 - Aug 2021 | Seattle, WA

- Built **5** full-stack search interfaces and a Java parser supporting **20+** search features.
- Translated structured user inputs to SQL queries for large-scale data retrieval.

Amazon Web Services — Software Development Engineer Intern June 2020 - Aug 2020 | Seattle, WA

- Automated cleaning and migration for **~10K** data files and CLI search with AWS Athena and Lambda.

ENGINEERING PROJECTS

PaperTalk — Open-source Codex Skill for Research Paper Analysis

- Built an evidence-grounded multi-agent system with Author, Reviewer, and interactive Panel roles.
- Designed role-specific **distillation/reconstruction** and persistent sessions for critique, explanation, and research extensions.

LLM-Emerald — Long Horizon LLM Agent for Pokemon Emerald

- Built a Voyager-style multimodal agent with **hierarchical planning, reusable skills, and in-game states** for long horizon interactions and across-run learning.
- Implemented a **skill-memory and verification pipeline** that converts model-generated plans into reusable executable behaviors and evaluates successful skills before reuse.

TECHNICAL SKILLS

Languages: Python, Java, SQL

Research: Reinforcement learning, Contextual bandits, Online learning, Causal inference, Experiment Design

Tools: PyTorch, JAX, NumPy, Git, AWS toolkits

SELECTED RESEARCH & PUBLICATIONS

[7] Y. Wen, Y. Pang, Y. Han, Z. Deng. **Trustworthy Recursive Self-improvement under Operational Constraints**, *Working draft*, 2026.
- Constraint-aware data reweighting and post-training audits for reliable self-improvement under operational constraints.

[6] Y. Wen. **Fast A/B/n Testing: Exact Multi-Policy Comparison via Tree-Coupled Feedback Sharing**, submitted to *ICLR*, 2026. (arXiv: 2608.12831)
- Exact feedback sharing across adaptive policies, saving 20-40% evaluation costs in experiments.

[5] Y. Wen, Y. Han, Z. Zhou. **Joint Value Estimation and Bidding in Repeated First-price Auctions**, submitted to *Management Science*, 2026. (POMS-HK 2026 Best Student Paper Finalist, POMS-China 2026 Best Student Paper Third Price, arXiv: 2502.17292)
- Combined online causal-effect estimation with adaptive bidding with minimal randomized experiments.

[4] Q. Bai*, Y. Han*, Y. Wen*, W. Xu, Z. Zhou. **Functional Critics Are Essential for Actor-Critic: From Off-Policy Stability to Efficient Exploration**, submitted to *ICLR*, 2026. (arXiv: 2509.22964).
- Functional critics for stable, sample-efficient off-policy actor-critic learning and exploration.

[3] Y. Wen, Z. Hu, Y. Han, Y. Yao, Z. Zhou. **The (Marginal) Value of a Search Ad: An Online Causal Framework for Repeated Second-price Auctions**, *ICML*, 2026.
- Online bidding methods that estimate the incremental causal value of paid exposure.

[2] Y. Wen, Y. Han, Z. Zhou. **Optimal Arm Elimination Algorithms for Combinatorial Bandits**, *AISTATS*, 2026.
- Efficient exploration-exploitation framework for combinatorial (multi-selection) bandits.

[1] Y. Wen. **Adversarial Combinatorial Semi-bandits with Graph Feedback**, *ICML*, 2025.
- Near-optimal algorithms for combinatorial bandits with side observations.

Additional works: ICLR 2025/2026; NeurIPS 2024, ICML 2024, *Physica D* 2024; *SIAM J. Scientific Computing* 2023.

Full list: yw99.github.io/publications.

AWARDS

[3] POMS-China Best Student Paper — Third Prize, 2026
[2] POMS-HK Best Student Paper — Finalist, 2026
[1] ICML Silver Reviewer Award, 2026

SERVICE

Service: Reviewer, ICML 2024-2026, NeurIPS 2024-2025, ICLR 2025-2027, AISTATS 2026; Journal of Machine Learning Research, Management Science, Math of Operations Research.

Teaching: TA, for Applied Statistics (Fall 2024/2023); Math for Econ III (Spring 2024).