

Caroline Wang

+1 (716) 531-2370 | caroline.l.wang@utexas.edu | carolinewang01.github.io

Research Interests: Human-AI coordination, multi-agent systems, ad hoc teamwork/zero-shot coordination, open-ended learning, reinforcement learning, imitation learning.

Education

University of Texas at Austin (GPA 4.0/4.0)

Final year PhD candidate supervised by Prof. Peter Stone

Austin, TX

August 2020 – Present

Duke University (GPA 3.93/4.00)

B.S. in Computer Science and Mathematics

Graduation with Highest Distinction.

Durham, NC

August 2016 – May 2020

Selected Publications

- **C. Wang**, D. Kasenberg, K. Stachenfeld, P. S. Castro. Discovering Differences in Strategic Behavior Between Humans and LLMs. Under submission / **HC-AIR at ICLR 2026**.
- **C. Wang***, A. Rahman*, J. Cui, Y. Sung, P. Stone. ROTATE: Regret-Driven Open-ended Training for Ad Hoc Teamwork. Under submission / **CoCoMARRL Workshop at RLC 2025**.
- **C. Wang**, A. Rahman, I. Durugkar, E. Liebman, P. Stone. N-Agent Ad Hoc Teamwork. **NeurIPS 2024**.
- Z. Wang*, **C. Wang***, X. Xiao, Y. Zhu, P. Stone. Building Minimal and Reusable Causal State Abstractions for Reinforcement Learning. **AAAI 2024**.
- **C. Wang***, I. Durugkar*, E. Liebman*, P. Stone. DM²: Distributed multi-agent reinforcement learning via distribution matching. **AAAI 2023**.
- **C. Wang**, G. Warnell, P. Stone. D-Shape: Demonstration-shaped reinforcement learning via goal-conditioning. **AAMAS 2023**.
- **C. Wang***, B. Han*, B. Patel, F. Mohideen, C. Rudin. In pursuit of interpretable, fair and accurate machine learning for recidivism prediction. **Quantitative Criminology, 2021**.
- C. Rudin, **C. Wang**, B. Coker. The age of secrecy and unfairness in recidivism prediction. **Harvard Data Science Review, 2020**.

Internships

Google DeepMind

Host: Pablo S. Castro

Project: LLM-powered scientific discovery for human behavior in strategic scenarios.

- Characterized differences in strategic behavior between LLMs and humans using AlphaEvolve, an LLM-powered evolutionary framework for scientific discovery.

June 2025 – Feb 2026

Sony AI

Host: Varun Kompella

Project: Automatically weighing event tables to improve reinforcement learning sample efficiency.

- Developed experience-replay based algorithms to improve sample efficiency of reinforcement learning on Sony game environments.

Summer 2023

Data Science for Social Good

Project: Quantifying traffic dynamics in Greater London to improve air pollution models

- Created open-source, end-to-end system that analyzes London traffic camera feeds using computer vision, and outputs junction-level, near real-time traffic statistics for emissions models.
- Published historical dataset of London traffic statistics. Dataset can be applied to improve emissions models, and/or by local transport authorities to optimize traffic flow/assess effects of roadworks.

Summer 2019

Software

JaxAHT

- Led team of 9 in developing and releasing the first open-source, JAX-based library for ad hoc teamwork algorithms and evaluation (<https://larg.github.io/jax-aht/>).

2025

Caroline Wang

+1 (716) 531-2370 | caroline.l.wang@utexas.edu | carolinewang01.github.io

Research Experience

Research Assistant

Fall 2020 – Present

Supervised by Prof. Peter Stone, Department of Computer Science, UT Austin

Project: Regret-driven Open-ended Training for Ad Hoc Teamwork

- Reformulated ad hoc teamwork as an open-ended learning problem and proposed a regret-driven algorithm that achieves state-of-the-art results against out-of-distribution teammates.

Project: Language Model Agents in Social Dilemmas

- Led team in participating in the Concordia Challenge at NeurIPS 2024, placing 6th/25.
- Constructed LLM-based agent for dealing with a variety of social scenarios (e.g. negotiation, collective bargaining, alliances) using prompt engineering and chain-of-thought reasoning.

Project: N-Agent Ad Hoc Teamwork

- Proposed and formulated novel problem of N-agent ad hoc teamwork.
- Developed agent modeling based MARL algorithm for solving N-agent ad hoc teamwork.

Project: Building Minimal and Reusable Causal State Abstractions for Reinforcement Learning

- Developed SAC-based algorithm that learns causal dynamics and reward relationships to generate a state abstraction that is reusable across environments with the same dynamics.

Project: Distributed multi-agent reinforcement learning via distribution matching

- Developed distribution matching algorithm for multi-agent cooperation using GAIL and IPPO.

Project: Deep reinforcement learning for learning a sprint in RoboCup 3DSim

- Parallelized PPO algorithm on a Condor cluster to learn a humanoid robot sprint.
- UTAustinVilla team placed 1st at RoboCup 3dSim League Competition in 2021; 3rd in 2022/23.

Project: Combining reinforcement learning with imitation from observation

- Developed novel, model-free reinforcement learning + imitation from observation algorithm.

Leadership

MAL-GAI Workshop at ICLR 2026

Fall 2025 – Spring 2026

- Organized workshop on generative multi-agent learning with international committee.

Graduate Women and Gender Minorities in Computing (GWGMC)

Spring 2022 – Spring 2025

- Re-started and co-led the GWGMC student organization at UT Austin.
- Organizer of Women and Gender Minorities in Computing Research Symposium.

Adaptive Learning Agents (ALA) Workshop at AAMAS 2023 / 24

Fall 2022 – Spring 2024

- Organized workshops on multi-agent systems in collaboration with international committee.

Awards and Honors

UT Austin Graduate School Fellowship

Summer 2024

- Fellowship support from the Graduate School at UT Austin.

CRA Outstanding Undergraduate Researcher Award Finalist

2020

- Computing Research Association (CRA) award program recognizes undergraduate students who show outstanding potential in computing research; 23 finalists selected from North America.

Barry M. Goldwater Scholarship

2019

- Highly prestigious undergraduate STEM research scholarship. 2019 Goldwater Scholars selected from an estimated pool of over 5,000 college students and 443 academic institutions in U.S.

Skills and Interests

- **Languages:** Python (proficient), C++ (intermediate), R.
- **Tools/Frameworks:** JAX, PyTorch, NumPy, Git, Bash/Unix, distributed computing clusters.
- **Knowledge:** Multi-agent systems, reinforcement learning, large language models, machine learning, data science, statistics.