

Reilly Raab

✉ raabrp@gmail.com | 🏠 reillyraab.com | 🌐 github.com/raabrp

SUMMARY

PhD in computer science focused on machine learning and multi-agent systems research, a strong mathematical background in physics, and extensive programming experience implementing highly parallel numerical simulations, novel optimization techniques, and signal processing algorithms.

SKILLS

Machine Learning: Constraints & Guarantees, Nonstationary Environments, Information Theory.
Optimization: Convex Optimization, Evolutionary Game Theory, Approximate Newton Methods.
Programming: GNU/Linux, Python (incl. NumPy, SciPy, Jax, Taichi, Gym), C, Lisp, JavaScript, Open-Source.

EDUCATION

PhD, Computer Science and Engineering Sept 2019 – Mar 2024
University of California, Santa Cruz Santa Cruz, CA
• ARCS Scholar • Dean's Fellow • Regents Fellow • Dissertation Year Fellow • Advancement with Honors

BSc, Physics Sept 2011 – June 2015
University of California, Santa Barbara (College of Creative Studies) Santa Barbara, CA
• High Honors • Distinction in the Major • Education Abroad Scholarships

RESEARCH & INTERESTS

My doctoral research focused on questions of safety and alignment in machine learning and dynamical systems formed by multiple learning agents. My work in this domain has pioneered new models for dynamical systems affected by data-driven policy [3], introduced new algorithms for constrained optimization [7], and uncovered fundamental connections between evolutionary game theory and natural gradient descent [5].

I strongly suspect that the next generation of artificial intelligence will not be focused on single models, but on multi-agent systems understood in terms of interacting units of control. I am interested in conducting research towards this end.

HONORS

Paper Awards

Best Paper Runner-Up [6]	RTML* Workshop at ICLR [†]	2023
Highlighted Paper [6]	RTML Workshop at ICLR	2023
Spotlight Paper [3]	NeurIPS [‡]	2021

Academic Honors

Advancement with Honors (to PhD candidacy)	UC Santa Cruz	2021
High Honors (BSc)	UC Santa Barbara	2015
Distinction in the Major (Physics)	UC Santa Barbara	2015

Scholarships & Fellowships

ARCS Scholar	ARCS Foundation, Inc., Northern California Chapter	2022
Regents Fellowship	UC Santa Cruz	2019
Dean's Fellowship	UC Santa Cruz	2019

PUBLICATIONS

[7] Fair Participation via Sequential Policies.

Reilly Raab, Ross Boczar, Maryam Fazel, and Yang Liu. *AAAI*, 2024.

*Trustworthy and Reliable Large-Scale Machine Learning Models (RTML)

[†]International Conference on Learning Representations (ICLR)

[‡]Conference on Neural Information Processing Systems (NeurIPS)

- [6] Long-Term Fairness with Unknown Dynamics.
Tongxin Yin[§], **Reilly Raab**[§], Mingyan Liu, and Yang Liu. *NeurIPS*, 2023.
- [5] Conjugate Natural Selection.
Reilly Raab, Luca de Alfaro, and Yang Liu. *arXiv Preprint*, 2023.
- [4] Fairness Transferability Subject to Bounded Distribution Shift.
Yatong Chen[§], **Reilly Raab**[§], and Yang Liu. *NeurIPS*, 2022.
- [3] Unintended Selection: Persistent Qualification Rate Disparities and Interventions.
Reilly Raab and Yang Liu. *NeurIPS (Spotlight Paper)*, 2021.
- [2] Single-Gate Error for Superconducting Qubits Imposed by Sideband Products of IQ Mixing.
Reilly P. Raab. *UC Santa Barbara Physics Department Website*, 2015.
- [1] Systematic Study of Exciton Diffusion Length in Organic Semiconductors by Six Experimental Methods.
Jason Lin **et al.** *Materials Horizons*, 2014.

EXPERIENCE

Graduate Student Researcher

Sept 2019 – Mar 2024

Human-Centered Machine Learning, UC Santa Cruz

Santa Cruz, CA

The Human-Centered Machine Learning Group at UC Santa Cruz researches the real-world, human consequences of deployed machine learning (ML) systems. My role in this group involves proposing original research questions, deriving theoretical results, designing numerical experiments, authoring computational simulations, writing papers, and presenting our research in multiple top ML conferences in talks and posters.

Accomplishments:

- Described dynamics of systems of mutual learners using evolutionary game theory [3].
- Established adversarial bounds for fairness violation due to distribution shift [4].
- Discovered exact, least-squares correspondence between replicator dynamics and natural gradient descent [5].
- Adapted online reinforcement learning methods to systems of mutually interacting learners [6].
- Mapped machine learning with policy-induced distribution shift to a novel constrained optimization algorithm [7].
- Wrote multi-agent simulations in Python using JAX and Taichi for GPU acceleration.

Software Developer

Oct 2016 – Aug 2018

Breadware, Inc.

Reno, NV

As a startup company, Breadware, Inc. offered consulting and rapid prototyping for internet-of-things (IoT) products. My role at the company was to build proprietary tools to modularize device development and automate associated engineering tasks.

Accomplishments:

- Wrote proprietary software to automate electronic design tasks, such as PCB layout (Python).
- Mapped abstract hardware APIs to I2C bus protocols for modular embedded devices (Python, C).
- Implemented web-based testing of user-logic for embedded devices in simulated environments (JavaScript).

Teaching Assistant and Residential Mentor

Summer 2015 | Summer 2016

The Summer Science Program

Socorro, NM | Boulder, CO

The Summer Science Program is a non-profit organization, run by its alumni (of which I am one), that, since 1959, has hosted advanced high school students at university campuses to take accelerated college-level coursework in celestial mechanics, programming, and mathematics and conduct research in observational astronomy. I was a Teaching Assistant and Residential Mentor for the program in the summers of 2015 (at New Mexico Tech) and 2016 (at CU Boulder), where I helped supervise a cohort of students in this research, graded homework, designed supplementary challenges, and gave supplementary lectures.

[§]Equal contribution