Tong Mu

Menlo Park, CA | Site: tongmu.github.io | Email: tongm@cs.stanford.edu

PhD candidate with 5 years of experience in reinforcement learning (RL) and artificial intelligence (AI) research, including experience with both TensorFlow and Pytorch deep learning frameworks as well as real-world user testing of algorithms.

Education

PhD Candidate — **Stanford University,** GPA: 3.98 *Sept. 2016 – Present (Expected June 2022)*
Research Advisor: Emma Brunskill
Research Focus: Reinforcement Learning, Bandit, and AI Algorithms for human-centered applications

B.S. in Electrical Engineering — **University of California, Los Angeles (UCLA)**, GPA: 3.98 *Sept. 2012 – June 2016*

Work Experience

Research Intern — **Salesforce**, Remote CA, USA *June 2021 – Oct. 2021*
- Proposed a new deep RL agent model for agent-based simulations that incorporates models of human-decision making from behavioral economics to model complex real-world systems not previously possible. ***(In submission)***

Data Science Lab Research Intern — **Adobe**, Remote CA, USA *June 2020 – Sept. 2020*
- Designed a new RL algorithm which uses domain expertise to learn and reach high performance 2-5 times faster than baselines when validated in multiple simulators fit with real-world data, such as recommendation systems. ***(AAAI2022)***

Visiting Researcher— WarChild Holland, Amsterdam, Netherlands *Oct. 2019 – Dec. 2019*
- Designed an algorithm based on explainable machine learning that accurately predicts interventions to prevent students becoming stuck in a computerized educational program for students in resource-limited regions. ***(EDM2020)***

Visiting Researcher (Contractor) — **Cresta**, San Francisco, CA, USA *June 2019 – Sept.2019*
- Designed off-policy policy evaluation RL algorithms for chat support message recommendations.

Software Intern — **Intel**, Hillsboro OR, USA *June 2015 – Sept. 2015 & June 2016 – Aug. 2016*
- Created web applications with responsive graphics that visualized timing paths in circuits.

Research Experience

Stanford AI4HI (Artificial Intelligence for Human Impact) Lab, Advisor: Emma Brunskill *June 2017 – Present*
- Multiple projects designing and real-world user testing online RL, bandit and AI algorithms for human-centered applications (such as educational activities sequencing).
- Projects studying and designing offline reinforcement learning algorithms when using prior data.
- Work appears in ITS2021, AAMAS2019, Neurips 2018 demonstration, LAS2018 among others.

UCLA Communications Systems Laboratory, Advisor: Richard Wesel *2015-2016*

Teaching and Mentorship

Course Assistant, CS234 Reinforcement Learning — **Stanford University** *Winter 2020*
CURIS Mentor — **Stanford University** *Multiple Academic Quarters*
- Mentored 5 Undergraduate Students through CURIS, the CS Undergraduate Research Program.

Publications (Selected)

Total Publications: 1 in submission, 5 conference, 1 journal, 4 other (workshop, poster, demonstrations, etc)

1. Modeling Bounded Rationality in Multi-Agent Simulations with Rationally Inattentive Reinforcement Learning
   In submission ICLR2022 | **Tong Mu,** Stephan Zheng, Alexander Trott
2. Constraint Sampling Reinforcement Learning: Incorporating Expertise For Faster Learning
   AAAI2022 | **Tong Mu,** Georgios Theocharous, David Arbour, Emma Brunskill
3. ***Best Short Paper Award*** Automatic Adaptive Sequencing in a Foreign Language Game
   ITS 2021 (Intelligent Tutoring Systems) | **Tong Mu**, Shuhan Wang, Erik Andersen, Emma Brunskill

Research Awards

Graduate Fellowship For STEM Diversity (GFSD) *Fall 2017 – Present (Ends Spring 2023)*
Best Short Paper Award *International Conference on Intelligent Tutoring Systems (ITS) 2021*
Don Carlisle Undergraduate Research Award (4 students yearly) *Summer 2014*

Technical Skills

Coding/Deep Learning/Machine Learning: Python, TensorFlow, PyTorch, Scikit-learn, Pandas, Experience in: R, C++, Matlab
Experimentation Skills: Designing user studies, hypothesis testing, data analysis, data visualization