

Yutai Zhou

Ph.D. Student at the University of Southern California

✉ yutaizho@usc.edu ·  LinkedIn ·  Google Scholar

📍 Los Angeles, California, USA

RESEARCH INTEREST

I am broadly interested in model evaluation, data attribution, and learning from human feedback (e.g. intervention, gaze) for generative models, with applications predominantly in decision-making and robot learning settings. My work spans interactive imitation learning, preference-based reward modeling, offline reinforcement learning, and Bayesian deep learning.

EDUCATION

- **University of Southern California** Los Angeles, CA
Ph.D. in Computer Science; GPA: 3.83 / 4.00 *Aug. 2023 – Aug. 2028 (Expected)*
Co-advisors: Erdem Biyik and Stephen Tu
- **Stanford University** Online
Stanford Center for Professional Development; GPA: 4.00 / 4.00 *Aug. 2021-May. 2023*
Took three graduate courses as part of a non-degree program while working full time.
- **University of Florida** Gainesville, FL
B.S. in Computer Science; GPA: 3.73 / 4.00 (Cum Laude); Honors Program *Jun. 2014 - Dec. 2019*

EXPERIENCE

- **University of Southern California** Los Angeles, CA
Graduate Research Assistant *Aug. 2023 – Present*
Co-Advisor: Erdem Biyik and Stephen Tu
 - **Disentangled preference learning:** Apply disentangled representation learning techniques to decompose preference representations into semantically meaningful axes.
 - **Bayesian active learning:** Developed efficient Bayesian neural network methods for training reward models from preference feedback. (Work under submission)
 - **Learning from human feedback:** Developing interactive imitation learning methods to learn from scalable human supervision signals such as gaze and intervention. (IROS 2025, ICRA 2026)
 - **Offline policy evaluation:** Developing scalable algorithms for evaluating robot policies in a safe manner, using techniques such as model-free and model-based RL, and diffusion-based generative models. (Co-PI: Professor Stephen Tu)
- **MIT Lincoln Laboratory** Lexington, MA
Research Engineer *Jan. 2020 - Jun. 2023*
 - **Human-AI teaming:** Studied trust in human-AI teams in the game of Hanabi. (NeurIPS 2021)
 - **Multiagent RL:** Studied coordination and emergent communication in MARL, in collaboration with Stanford University (PI: Professor Mykel Kochenderfer). (ICRA 2022 & AAMAS 2021)
 - **Anomaly detection:** Leveraged self-supervised learning for anomaly detection in air traffic data.
 - **Teaching:** Created and taught an RL course for the U.S. Air Force Academy cadets. Hosted a tutorial on the same topic for “Recent Advances in AI for National Security Workshop (RAAINS 2020)”

Undergraduate Research Intern *May 2019 - Aug. 2019*

 - **Video object detection:** Investigated integration of multimodal data and varying model architectures for a human-in-the-loop video object detection system. (Finalist for R&D 100 (2020))

- **Machine Learning and Sensing Lab**

Undergraduate Research Assistant

Principal Investigator: Professor Alina Zare

Gainesville, FL

Aug. 2018 - Apr. 2019

- **Image processing:** Developed a Python library for hyperspectral imaging for plant phenotyping, in collaboration with Donald Danforth Plant Science Center. (hsi_toolkit_py, manuscript in progress)

- **Sandia National Laboratories**

Undergraduate Software Research Intern

Albuquerque, NM

Jan. 2018 - Aug. 2018

- **Video object tracking:** Develop an automated pipeline for data simulation, feature engineering, and neural network training. With application for video anomaly detection.
- **Backend engineering:** Improved runtime performance and UX of a satellite remote sensing system by upgrading various UI and backend components in ExtJS.

PROFESSIONAL SERVICES

- **Reviewer for Conferences**

- International Conference on Machine Learning (ICML): 2026
- International Conference on Learning Representations (ICLR): 2026
- IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS): 2025
- Conference on Neural Information Processing Systems (NeurIPS): 2025
- International Conference on Learning Representations (ICLR): 2026
- IEEE International Conference on Robotics and Automation (ICRA): 2026

- **Reviewer for Workshops**

- Workshop on Human-in-the-Loop Robot Learning: Teaching, Correcting, and Adapting (RSS 2025)
- Eval&Deploy@CoRL2025: Evaluation and Deployment Across the Robot Learning Lifecycle (CoRL 2025)

TEACHING EXPERIENCE

- **University of Southern California**

Graduate Teaching Assistant

Los Angeles, CA

- DSCI 552: Machine Learning for Data Science (Summer 2024, Fall 2024)
- CSCI 599: Autonomous Decision-Making (Spring 2024)

- **University of Florida**

Undergraduate Teaching Assistant

Gainesville, FL

- COP 2271: MATLAB Programming for Engineers (Spring 2017, Fall 2017)
- EGM 3520: Mechanics of Materials (Honors) (Summer 2016)

OUTREACH & MENTORING

- **Student Mentorship During Ph.D.**

- Abhinav Srivatsa (2025)
- Nitya Kashyap (2025)
- Meredith Li (2025)
- Arpita Tomar (2025)
- Shreya Ramanujam (2024)

Next: 2nd year undergrad at USC

Next: 4th year undergrad at BITS Pilani

Next: MS at Stanford University

- **USC Thomas Lord Department of Computer Science**

Los Angeles, CA

- **CURVE Research Mentor** (2025)
- **IUSSTF Research Mentor** (2024, 2025)

- **Undergraduate Mentoring Program Organizer** (2024, 2025): Matched Ph.D. student mentors with undergraduate students from under-represented communities.
- **University of Florida** Gainesville, FL
 - **Gator Powerlifting Club Vice President** (2017): Established club training sessions and merchandise discounts with local powerlifting gym.
 - **STEPUP XXI Program Mentor** (2015): Resident assistant and teaching assistant for freshmen engineering preparation program.

HONORS & AWARDS

- **USC Capital One Fellowship (2025)**: Center for Responsible AI and Decision Making in Finance (CREDIF)
- **R&D 100 Finalist (2020)**: For work done as an undergraduate intern at MIT Lincoln Laboratory.
- **IEEE HPEC 2020 Outstanding Paper Award**: Awarded for the paper “Towards a Distributed Framework for Multi-Agent Reinforcement Learning Research”.
- **Florida Bright Future Academic Scholarship**: State-wide scholarship for undergraduate studies.
- **TeamWork USA Education Foundation Scholarship (2014)**: Private scholarship for undergraduate studies.
- **Palm Beach County Gator Club Scholarship (2014)**: Private scholarship for undergraduate studies.
- **Charles R. O’Melia Scholarship (2014)**: Private scholarship for undergraduate studies.

CONFERENCE PUBLICATIONS

- [C1] A. Liang, P. Czempin, M. Hong, **Y. Zhou**, E. Biyik, and S. Tu, “CLAM: Continuous Latent Action Models for Robot Learning from Unlabeled Demonstrations,” in IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2026)
- [C2] L. Gong, F. Bahrani, **Y. Zhou**, A. Banayeeanzade, J. Li, and E. Biyik, “AutoFocus-IL: VLM-based saliency maps for data-efficient visual imitation learning without extra human annotations,” in International Conference on Robotics and Automation (ICRA 2026)
- [C3] A. Banayeeanzade, F. Bahrani, **Y. Zhou**, and E. Biyik, “GABRIL: Gaze-Based Regularization for Mitigating Causal Confusion in Imitation Learning,” in 2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2025)
- [C4] H. C. Siu, J. Pena, **Y. Zhou**, E. Chen, V. J. Lopez, K. Palko, K. C. Chang, R. E. Allen, “Evaluation of Human-AI Teams for Learned and Rule-Based Agents in Hanabi,” in Advances in Neural Information Processing Systems, Curran Associates, Inc., 2021, pp. 16183–16195.
- [C5] S. Li, **Y. Zhou**, R. Allen, and M. J. Kochenderfer, “Learning Emergent Discrete Message Communication for Cooperative Reinforcement Learning,” in 2022 International Conference on Robotics and Automation (ICRA), May 2022, pp. 5511–5517. doi: 10.1109/ICRA46639.2022.9812285.
- [C6] **Y. Zhou**, S. Manuel, P. Morales, S. Li, J. Pena, and R. Allen, “Towards a Distributed Framework for Multi-Agent Reinforcement Learning Research,” in 2020 IEEE High Performance Extreme Computing Conference (HPEC), Sept. 2020, pp. 1–9. doi: 10/gj7hzv.

WORKSHOP PAPERS

- [W1] A. Liang, P. Czempin, **Y. Zhou**, S. Tu, and E. Biyik, “In-Context Generalization to New Tasks From Unlabeled Observation Data,” presented at the ICML 2024 Workshop on In-Context Learning, June 2024.
- [W2] R. E. Allen, J. K. Gupta, J. Pena, **Y. Zhou**, J. W. Bear, and M. J. Kochenderfer, “Health-Informed Policy Gradients for Multi-Agent Reinforcement Learning,” The 12th Workshop on Optimization and Learning in Multiagent Systems at AAMAS 2021, Jan. 2021.

PREPRINTS

[P1] H. C. Siu, J. D. Peña, **Y. Zhou**, and R. E. Allen, “In Pursuit of Predictive Models of Human Preferences Toward AI Teammates,” Jan. 31, 2025, arXiv: arXiv:2503.15516. doi: 10.48550/arXiv.2503.15516.

MISCELLANEOUS

- **Citizenship:** United States
- **Spoken Languages:** English, Mandarin Chinese (both native fluency)
- **Interests:** Powerlifting, bodybuilding, swimming, traveling