

Letian (Zac) Chen

✉ zac.letian.chen@gmail.com

🌐 www.letianchen.me/

☎ +1 (404) 247-9580

📷 [Zac Chen](#)

EDUCATION

Georgia Institute of Technology, Atlanta GA

May 2025

Doctor of Philosophy in Computer Science, School of Interactive Computing. GPA: 4.00/4.00.

- Thesis: “Personalized, Safe, and Interactive Robot Programming via Human Demonstrations”.
- Advisor: Dr. Matthew Gombolay.

Georgia Institute of Technology, Atlanta GA

May 2020

Master of Science in Computer Science, College of Computing. GPA: 4.00/4.00.

- Thesis: “Robot Learning from Heterogeneous Demonstration”.

Peking University, Beijing China

July 2018

Bachelor of Science in Computer Science, School of Electronics Engineering and Computer Science.

Bachelor of Science in Psychology, School of Psychological and Cognitive Sciences.

INDUSTRY EXPERIENCE

Waymo, Research Scientist

June 2025 – Current

Waymo, Research Intern

May 2024 – Dec 2024

- Designed input-output representations for fine-tuning large Vision-Language-Models (VLMs) for the vehicle planning task.
- Proposed a novel Reinforcement Learning algorithm for fine-tuning VLMs towards planning metrics to replace the usual token matching objectives in LM training.
- Developed training and evaluation pipeline infrastructure of the large VLM. Experiments show significant improvements in target behavior metrics via the proposed method.

Toyota Research Institute, Research Intern

May 2023 – Aug 2023

- Implemented DIAYN to generate diverse driving policies for autonomous racing.
- Proposed a novel algorithm, Learn Thy Enemy, to model and leverage opponent information in multi-car racing.
- Deployed DIAYN and LTE policies on motion simulator hardware and demonstrated qualitatively and quantitatively high performance.

iRobot Corporation, Reinforcement Learning Research Intern

May 2021 – Aug 2021

- Identified real-world challenges of Offline Policy Evaluation (OPE) algorithms.
- Created a benchmark dataset for OPE with real-world challenges.
- Proposed an ad-hoc OPE algorithm selection method to decide which OPE algorithm is reliable with the task and available offline data.

PUBLICATIONS

ICML '25

L. Chen, N. Moorman, and M. Gombolay, “Elemental: Interactive learning from demonstrations and vision-language models for reward design in robotics,” in *Proceedings of International Conference on Machine Learning (ICML)*, 2025

RA-L '25

C. Xue, **L. Chen**, and M. Gombolay, “Better than diverse demonstrators: Reward decomposition from suboptimal and heterogeneous demonstrations,” *IEEE Robotics and Automation Letters*, 2025

- HSMR '25 M. Natarajan, Z. Li, **L. Chen**, Z. Wu, P. Ogara, P. Borges, G. Rance, R. Srey, R. Harari, S. Mendu, M. Zenati, R. Dias, and M. Gombolay, "Ai-based decision support for perfusionists during cardiopulmonary bypass," in *Hamlyn Symposium on Medical Robotics*, 2025
- RA-L '24 B. Han, M. Cho, **L. Chen**, R. Paleja, Z. Wu, S. Ye, E. Seraj, D. Sidoti, and M. Gombolay, "Learning multi-agent coordination for replenishment at sea," *IEEE Robotics and Automation Letters*, 2024
- HSMR '24 Z. Li, M. Natarajan, **L. Chen**, Z. Wu, P. Ogara, P. Borges, G. Rance, R. Srey, R. Harari, M. Zenati, R. Dias, and M. Gombolay, "Using ml for perfusionists' decision prediction for robotic-assisted cardiopulmonary bypass in cardiac surgery," in *Hamlyn Symposium on Medical Robotics workshop on Hybrid Human-Machine Interaction in Surgery*, 2024
- HSMR '24 R. Dias, R. Harari, M. Zenati, G. Rance, R. Srey, **L. Chen**, and M. Gombolay, "A clinician-centered explainable artificial intelligence framework for decision support in the operating theatre," in *Hamlyn Symposium on Medical Robotics*, 2024
- Neuromodulation '24 E. Cole, M. Schrum, E. Opri, **L. Chen**, A. Nascimento, P. Testini, A. Wang, F. Isbaine, R. Gross, M. Gombolay, and S. Miocinovic, "Automated deep brain stimulation programming via meta-active learning of evoked potentials," in *North American Neuromodulation Society Conference*, 2024
- HRI '24 **L. Chen***, Y. Yang*, Z. Zaidi*, S. van Waveren, A. Krishna, and M. Gombolay, "Enhancing safety in learning from demonstration algorithms via control barrier function shielding," in *Proceedings of International Conference on Human-Robot Interaction (HRI)*, 2024
- ML4AD '23 L. Chen, S. Manuel, J. Delgado, S. John, and T. Paul, "Learn thy enemy: Online, task-aware opponent modeling in autonomous racing," in *Machine Learning for Autonomous Driving Symposium*, 2024
- CoRL '23
OOD Workshop V. Sreeramdas, R. R. Paleja, **L. Chen**, S. van Waveren, and M. Gombolay, "Generalized behavior learning from diverse demonstrations," in *First Workshop on Out-of-Distribution Generalization in Robotics at CoRL 2023*, 2023
- CoRL '23 **L. Chen***, S. Jayanthi*, N. Balabanska, V. Duong, E. Scarlatescu, E. Ameperosa, Z. H. Zaidi, D. Martin, T. D. Matto, M. Ono, and M. Gombolay, "Droid: Learning from offline heterogeneous demonstrations via reward-policy distillation," in *Proceedings of Conference on Robot Learning (CoRL)*, 2023
- Current Robotics Reports **L. Chen***, M. Natarajan*, E. Seraj*, B. Altundas*, R. Paleja*, S. Ye*, R. Jensen, K. C. Chang, and M. Gombolay, "Human-robot teaming: Grand challenges," *Current Robotics Reports*, pp. 1–20, 2023
- IROS '23 S. Ye, M. Natarajan, Z. Wu, R. Paleja, **L. Chen**, and M. C. Gombolay, "Learning models of adversarial agent behavior under partial observability," in *Proceedings of International Conference on Intelligent Robots and Systems (IROS)*, 2023

- MRS '23 Z. Wu, S. Ye, M. Natarajan, **L. Chen**, R. Paleja, and M. C. Gombolay, "Adversarial search and track with multiagent reinforcement learning in sparsely observable environment," in *Proceedings of International Symposium on Multi-Robot & Multi-Agent Systems (MRS)*, 2023
- IEEE RA-L Z. Zaidi, D. Martin, N. Belles, V. Zakharov, A. Krishna, K. M. Lee, P. Wagstaff, S. Naik, M. Sklar, S. Choi, Y. Kakehi, R. Patil, D. Mallemadugula, F. Pesce, P. Wilson, W. Hom, M. Diamond, B. Zhao, N. Moorman, R. Paleja, **L. Chen**, E. Seraj, and M. Gombolay, "Athletic mobile manipulator system for robotic wheelchair tennis," *IEEE Robotics and Automation Letters*, vol. 8, no. 4, pp. 2245–2252, 2023
- HRI '23 K. M. Lee*, A. Krishna*, Z. Zaidi, R. Paleja, **L. Chen**, E. Hedlund, M. Schrum, and M. Gombolay, "The effect of robot skill level and communication in rapid, proximate human-robot collaboration," in *Proceedings of International Conference on Human-Robot Interaction (HRI)*, 2023
- CoRL '22
Agility Workshop Y. Yang*, **L. Chen***, and M. Gombolay, "Safe inverse reinforcement learning via control barrier function," in *Proceedings of CoRL Agility Workshop*, 2022
- CoRL '22 **L. Chen***, S. Jayanthi*, R. R. Paleja, D. Martin, V. Zakharov, and M. Gombolay, "Fast lifelong adaptive inverse reinforcement learning from demonstrations," in *Proceedings of Conference on Robot Learning (CoRL)*, 2022
- AAAI '22
IML Workshop **L. Chen***, S. Jayanthi*, and M. Gombolay, "Strategy discovery and mixture in lifelong learning from heterogeneous demonstration," in *Proceedings of AAAI Interactive Machine Learning workshop*, 2022
- IEEE T-RO E. Seraj, **L. Chen**, and M. Gombolay, "A hierarchical coordination framework for joint perception-action tasks in composite robot teams," *IEEE Transactions on Robotics*, 2021
- AAAI '21
AI-HRI Symposium **L. Chen**, R. Paleja, and M. Gombolay, "Towards sample-efficient apprenticeship learning from suboptimal demonstration," in *Artificial Intelligence for Human-Robot Interaction (AI-HRI)*, AAAI Fall Symposium Series, 2021
- MICCAI '21 R. D. Dias, M. A. Zenati, G. Rance, R. Srey, D. Arney, **L. Chen**, R. Paleja, L. R. Kennedy-Metz, and M. Gombolay, "Using machine learning to predict perfusionists' critical decision-making during cardiac surgery," *Computer Methods in Biomechanics and Biomedical Engineering: Imaging & Visualization*, vol. 10, no. 3, pp. 308–312, 2022
- CoRL '20 **[Best Paper Finalist][Plenary Talk] L. Chen**, R. Paleja, and M. Gombolay, "Learning from suboptimal demonstration via self-supervised reward regression," in *Proceedings of Conference on Robot Learning (CoRL)*, 2020
- Master's Thesis **L. Chen**, "Robot learning from heterogeneous demonstration," *Master Thesis*, 2020
- Neurips '20 R. Paleja, A. Silva, **L. Chen**, and G. Matthew, "Interpretable and personalized apprenticeship scheduling: Learning interpretable scheduling policies from heterogeneous user demonstrations," in *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, 2020

HRI '20

L. Chen, R. Paleja, M. Ghuy, and M. Gombolay, "Joint goal and strategy inference across heterogeneous demonstrators via reward network distillation," in *Proceedings of International Conference on Human-Robot Interaction (HRI)*, 2020

Undergrad Thesis

L. Chen, "Model-free vs model-based algorithms in human sequential decision making," *Undergraduate Thesis*, 2018

AWARDS & HONORS

- Amazon Science Scholarship for AAAI 2022 Feb 2022
- Best Paper Finalist for Conference on Robot Learning (CoRL 2020) Nov 2020
- First Place in Brainhack ATL 2019 Track 2 Nov 2019
- Beijing Merit Graduate Jul 2018
- Excellent Graduate, Peking University Jul 2018
- Zhang Wenjin Scholarship Dec 2017
- Scholarship for Undergraduate research Sep 2017
- First-Prize Winner of National Olympiad in Informatics (Zhejiang Province) 2008, 2009, 2010, 2011

SELECTED RESEARCH PROJECTS

Leveraging Visual Language Models (VLMs) in Robot Learning for Demonstration 2024

- Developed a novel method to leverage VLMs to extract semantically meaningful features from human demonstrations.
- Proposed a novel IRL algorithm to learn reward functions based on VLM extracted features and user demonstrations to achieve value alignment.
- Experimenting the algorithms with simulated domains and real user studies.

Safe Learning from Demonstration 2022-2023

- Created a new modality for users to specify safe vs. unsafe states for robots via demonstrations.
- Proposed a novel shielding algorithm, SECURE, that can be applied on policies to enforce customized safety bounds (defined by users), via a combination of data-driven control-barrier function and task-aware safe action search.
- Tested SECURE on two simulated robotic control tasks and a real robot kitchen cutting task where the robot is equipped with a knife; showed SECURE successfully prevent all unsafe executions, such as human-hand entering robot cutting space.

Learning Interpretable Tree-based Control Policies for Autonomous Driving 2022-2023

- Developed interpretable, tree-based continuous-control models that allow gradient updates.
- Demonstrated the strong qualitative and quantitative performance of the proposed model in comparison with black-box neural networks in 10+ driving scenarios.
- Verified interpretability with user-studies to show the proposed model is easier and faster to interpret than neural networks and other interpretable models.

Learning from Offline Heterogeneous Demonstrations 2022-2023

- Analyzed real Mars rover driving data and identified heterogeneity among rover drivers.
- Proposed a novel IRL framework, DROID, to accommodate the offline learning required by the application while allowing learning from heterogeneous demonstrations via dual reward and policy distillation.
- Applied DROID on two simulated robotic control tasks and the real Mars rover path-planning problem; achieved better learning and generalization to unseen conditions in all three domains.

Fast Lifelong Adaptive Learning from Demonstrations 2021

- Analyzed the personalization problem in lifelong learning from demonstration process where large number of heterogeneous demonstrations arrive sequentially by federation among users.
- Proposed a novel IRL framework, FLAIR, to provide efficient personalization and scalability by constructing *policy mixtures* with a concise set of prototypical strategy policies.
- Applied FLAIR on three virtual robotic control tasks and a real robot table-tennis task; achieved better personalization with significantly higher sample efficiency.

Learning from Suboptimal Demonstrations

2020

- Characterized policy performance degradation from noise injection by a sigmoid function.
- Proposed a novel IRL framework, SSRR, to learn policies that are better than suboptimal demonstrations by inferring the idealized reward function (i.e., the latent intent of the demonstrator).
- Applied algorithm on three virtual robotic tasks and a real robot table-tennis task; achieved accurate recovery of the demonstrator intention and a better-than-best-demonstration policy.

Learning from Heterogeneous Demonstrations

2019

- Modeled humans' latent objective via shared task reward and individual strategy reward.
- Proposed a novel IRL framework, MSRD, to jointly infer task reward and strategy reward to gain a better estimation of both.
- Applied algorithm on two virtual robot control tasks and one real robot table-tennis task; achieved better learning of task reward than SOTA AIRL, extracted precise strategic rewards, and optimized versatile policies that resemble the heterogeneous demonstrations.

TEACHING & LEADERSHIP EXPERIENCE

Advising & Mentorship

- Yue Yang, M.Sc. Student at Georgia Tech, now Ph.D. Student at UNC
- Sravan Jayanthi, B.Sc. and M.Sc. Student at Georgia Tech, now at C3 AI
- Daniel Martin, M.Sc. Student at Georgia Tech, now at Amazon Robotics
- Steve Zakharov, M.Sc. Student at Georgia Tech, now at Blue River Technology
- Sumedh Naik, M.Sc. Student at Georgia Tech, now at Intel
- Van Duong, M.Sc. Student at Georgia Tech, now at Jet Propulsion Laboratory

Teaching Assistantship

- *Interactive Robot Learning* (CS 7648, Graduate-level), Georgia Tech (Spring 2021)
- *Machine Learning* (OMSCS 7641, Graduate-level), Georgia Tech (Fall 2020, Spring 2019)
- *Introduction to Computation* (Undergraduate-level), Peking University (Fall 2016)

Guest Lecture

- *Interactive Robot Learning* (CS 7648, Graduate-level), Georgia Tech (Spring 2024, Fall 2022, Spring 2021)

Leadership

Academic President of Student Association, School of Psychological and Cognitive Sciences, Peking University (2015 - 2016)

ACADEMIC SERVICE

Reviewer for

- International Conference on Human Robot Interaction (HRI)
- Conference on Robot Learning (CoRL)
- International Symposium of Robotic Research (ISRR)
- IEEE Robotics and Automation Letters (RA-L)
- International Conference on Intelligent Robots and Systems (IROS)
- International Conference on Robotics and Automation (ICRA)
- International Conference on Autonomous Agents and Multiagent Systems (AAMAS)
- Robotics: Science and Systems (RSS)
- International Conference on Machine Learning (ICML)
- Annual Conference on Neural Information Processing Systems (NeurIPS)
- Transactions on Human-Robot Interaction (THRI)
- International Conference on Learning Representations (ICLR)
- International Conference on Artificial Intelligence and Statistics (AISTATS)
- International Conference on Robotics and Automation (ICRA)
- Conference on Computer Vision and Pattern Recognition (CVPR)
- International Conference on Machine Learning (ICML)