
| | |
|-------------------------|--|
| CONTACT | <p>gswamy@cmu.edu https://gokul.dev</p> |
| EDUCATION | <p>Carnegie Mellon University Ph.D. in Robotics (GPA: 4.04/4.00) Sept. 2020 - Present ▷ Thesis Committee: J. Andrew Bagnell, Zhiwei Steven Wu, Geoffrey J. Gordon, Arthur L. Gretton ▷ Publications: [Google Scholar]</p> <p>University of California, Berkeley M.S. in Computer Science Aug. 2019 - May 2020 ▷ Thesis: [Learning with Humans in the Loop] B.S. in Electrical Engineering and Computer Science, High Honors Aug. 2016 - May 2019</p> |
| RESEARCH EXPERIENCE | <p>Robotics Institute @ CMU, Pittsburgh, PA Sept. 2020 - Present <i>Graduate Student Researcher</i> Collaborating with Profs. Drew Bagnell and Steven Wu on <i>interactive learning from implicit human feedback</i>. Currently researching imitation learning, game theory, and causal inference with applications to robotics and large language models. Published at ICML 2021, 2022, 2023, 2024 (4x), NeurIPS 2022 (2x), 2023.</p> <p>Berkeley Artificial Intelligence Research Lab, Berkeley, CA Jan. 2018 - May 2020 <i>Graduate Student Researcher</i> Collaborated with Prof. Anca Dragan on comparing increasingly structured models of human drivers as modeling assumptions are broken and with Anca Dragan and Sergey Levine on allowing one person to supervise and provide corrections to a fleet of learning robots. Published at HRI 2019, ICRA 2020.</p> |
| PROFESSIONAL EXPERIENCE | <p>Google Research, Seattle, WA May 2023 - Nov. 2023 <i>Student Researcher</i> Collaborated with Drs. Alekh Agarwal, Chris Dann, and Rahul Kidambi on game-theoretic algorithms for reinforcement learning from conflicting human feedback for preference fine-tuning of large language models.</p> <p>Microsoft Research, Montreal, CA June 2022 - Aug. 2022 <i>Graduate Research Intern</i> Collaborated with Prof. Geoff Gordon on learning factorized dynamics models from visual observations. Investigating game-theoretic methods for learning structured latent spaces and sequence model architectures.</p> <p>Aurora, Pittsburgh, PA May 2020 - Aug. 2020 <i>Motion Planning ML Intern</i> Collaborated with Prof. Sanjiban Choudhury on learning deep driving policies that respected safety constraints (e.g. avoiding cyclists). Built C++ data pipelines / simulation tools and implemented constrained training of deep networks.</p> <p>NVIDIA, Santa Clara, CA May 2019 - Aug. 2019 <i>Autonomous Vehicles Perception Intern</i> Collaborated with Dr. Trung Pham on single-image weakly-supervised 3D structure estimation of intersection entry/exit lines. Designed CNN to recover lines in 3D that learned from 2D key points and 3D geometric constraints. Produced significant improvement over existing predict-then-project method.</p> <p>SpaceX, Hawthorne, CA May 2018 - Aug. 2018 <i>Data Engineering Intern</i> Collaborated with Dr. Anthony Rose on estimating shop-floor operation durations. Used hierarchical navigable small world graphs on top of word2vec to build approximate nearest neighbors engine that significantly outperformed domain experts. Created classical computer vision algorithm to detect flight risks.</p> <p>Intuit, San Diego, CA May 2017 - Aug. 2017 <i>Software Engineering Intern</i> Worked on using CNNs to classify the sentiments of product reviews. Created themable iOS UI framework and Java/PHP/React-based internal tools.</p> |

TEACHING
EXPERIENCE

Carnegie Mellon University, Pittsburgh, PA

Teaching Assistant

Sept. 2022 - Dec. 2022

TA for 10-732: Robustness and Adaptation in Shifting Environments w/ Prof. Zachary Lipton. Covered causally-structured and feedback-driven distribution shifts, online adaptation, and adv. robustness. [\[Site\]](#)

University of California, Berkeley, Berkeley, CA

Teaching Assistant

Jan. 2019 - Dec. 2019

Helped teach CS 188: Introduction to AI and CS 189: Introduction to Machine Learning. Gave guest lecture and created worksheet walkthrough videos.

Course Instructor

Aug. 2018 - May 2019

Created material for and taught course twice on societal impacts and ethical considerations of AI, covering topics like automation, bias in AI, data privacy, artificially generated data, and human-compatible AI.

ACTIVITIES

Communications Chair, Reinforcement Learning Conference (RLC) 2024
Workshop Organizer, Reinforcement Learning Beyond Rewards (RLC 2024) 2024
Future Faculty Program Participant, Eberly Center @ CMU 2022–
Teaching & Learning Summit Advisory Board, Eberly Center @ CMU 2022
Undergraduate Research Engagement Working Group, CMU SCS Dean’s Advisory Committee 2021–
 ▷ Put together interactive workshop for graduate students on best practices for mentoring undergraduates on research projects. Conducted IRB-approved study to measure training effects. Results selected for presentation at Eberly Teaching Summit. [\[Workshop Slides and Results Poster\]](#)
Graduate Application Support Program, CMU Robotics Institute, Mentor 2021–2022
Graduate Application Support Program, CMU Robotics Institute, Organizer 2021
Undergraduate AI Mentoring Program, CMU 2021
Journal Reviewer: IEEE Transactions on Robotics (IEEE T-RO), IEEE Transactions on Information Theory (IEEE T-IT), International Journal of Robotics Research (IJRR), Transactions on ML Research (TMLR)
Conference Reviewer: NeurIPS 2021-2023, ICRA 2022/2024, ICML 2023/2024, ICLR 2024, IROS 2024
Workshop Reviewer: Strategic ML @ NeurIPS 2021, Real World Reinforcement Learning @ NeurIPS 2022, Interactive Learning with Implicit Human Feedback @ ICML 2023, Frontiers for Learning, Dynamics, and Control @ ICML 2023, Robot Learning Workshop @ NeurIPS 2023

MENTORSHIP

Nicolas Espinosa Dice (Cornell), Speeding up inverse RL with suboptimal data. 2024–
Kensuke Nakamura (CMU), Interactive imitation learning on real robots. 2024–
Abigail DeFranco (CMU), Interactive imitation learning on real robots. 2024–
Jinwu Tang (CMU), Multi-agent imitation learning. 2023–
Juntao Ren (Cornell), Hybrid algorithms for inverse reinforcement learning. 2023–
Vibhakar Mohta (CMU), Interactive imitation learning on real robots. (next: Plus AI) 2023–
Silvia Sapora (Oxford), Faster inverse reinforcement learning via evolutionary algorithms. 2023–2024
Konwoo Kim (CMU), Learning safety constraints from multi-task demos. (next: Stanford PhD) 2022–2023
 ▷ Won 2023 CMU SCS Allen Newell Award for Excellence in Undergraduate Research.
Matthew Peng (Berkeley), Minimax-optimal online imitation learning. (next: Applied Intuition) 2021–2022

HONORS

Best Reviewer Award, ICML 2024
Invited to Jane Street Graduate Research Fellowship Research Workshop 2024
Finalist for Two Sigma PhD Fellowship, 1/6 students nominated by CMU Robotics Institute 2023
Finalist for JP Morgan PhD Fellowship, 1/3 students nominated by CMU School of Computer Science 2023
Top Reviewer, NeurIPS 2022
Finalist for IBM PhD Fellowship, 1/3 students nominated by CMU Robotics Institute 2022
Finalist for Apple PhD Fellowship, 1/2 students nominated by CMU Robotics Institute 2022
Finalist for Microsoft Research PhD Fellowship, 1/4 students nominated by CMU Robotics Institute 2022
Finalist for NVIDIA Graduate Research Fellowship, recieved GPU award 2021
NSF GRFP, Honorable Mention 2020

GRADUATE
COURSEWORK

Carnegie Mellon University: Convex Optimization, Computer Vision, Kinematics/Dynamics/Controls, Statistical Methods in ML, Computational Game Solving, Intermediate Statistics, Philosophical Foundations of ML, Optimal Control and Reinforcement Learning, Advanced Statistical Theory I, Game-Theoretic Statistics

University of California, Berkeley: Computer Vision, AI Safety, Information Theory, Linear Systems Theory, Advanced Robotics, Natural Language Processing

Jingwu Tang, **Gokul Swamy**, Fei Fang, Zhiwei Steven Wu, *Multi-Agent Imitation Learning: Value is Easy, Regret is Hard*. TTIC 2024 AdaptLearn Workshop, RL Safety (**Oral Presentation**) and RL Beyond Rewards Workshops @ RLC 2024, Models of Human Feedback for AI Alignment and Humans, Algorithms, and Society, and Aligning RL Experimentalists and Theorists Workshops @ ICML 2024. [\[PDF\]](#)

Zhaolin Gao, Jonathan D. Chang, Wenhao Zhan, Owen Oertell, **Gokul Swamy**, Kianté Brantley, Thorsten Joachims, J. Andrew Bagnell, Jason D. Lee, Wen Sun, *REBEL: RL via Regressing Relative Rewards*. TTIC 2024 AdaptLearn Workshop, RL Beyond Rewards Workshop @ RLC 2024, Models of Human Feedback for AI Alignment, Aligning RL Exps. and Theorists Workshops (**Oral Presentation**) @ ICML 2024. [\[PDF\]](#)

Nicolas Espinosa Dice, **Gokul Swamy**, Sanjiban Choudhury, Wen Sun *Efficient Inverse Reinforcement Learning without Compounding Errors*. TTIC 2024 AdaptLearn Workshop, RL Safety (**Oral Presentation**) + RL Beyond Rewards Workshops @ RLC 2024, Models of Human Feedback for AI Alignment Workshop @ ICML 2024. [\[PDF\]](#)

Yuda Song, **Gokul Swamy**, Aarti Singh, J. Andrew Bagnell, Wen Sun, *The Importance of Online Data: Understanding Preference Fine-tuning via Coverage*. TTIC 2024 AdaptLearn Workshop, RL Beyond Rewards Workshop @ RLC, 2024, Aligning RL Experimentalists and Theorists Workshop @ ICML 2024. [\[PDF\]](#)

Gokul Swamy, Christoph Dann, Rahul Kidambi, Zhiwei Steven Wu, Alekh Agarwal, *A Minimaximalist Approach to Reinforcement Learning from Human Feedback*. Internat. Conf. on ML (ICML), 2024. [\[PDF\]](#)

Juntao Ren*, **Gokul Swamy***, Zhiwei Steven Wu, J. Andrew Bagnell, Sanjiban Choudhury, *Hybrid Inverse Reinforcement Learning*. International Conference on ML (ICML), 2024. [\[Site\]](#)

Silvia Sapora, **Gokul Swamy**, Chris Lu, Yee Whye Teh, Jakob Nicolaus Foerster, *EvIL: Evolution Strategies for Generalisable Imitation Learning*, Internat. Conf. on ML (ICML), 2024. [\[PDF\]](#)

My Phan, Kianté Brantley*, Stephanie Milani*, Soroush Mehri*, **Gokul Swamy***, Geoffrey J. Gordon, *When is Transfer Learning Possible?*, Internat. Conf. on ML (ICML), 2024. [\[PDF\]](#)

Konwoo Kim*, **Gokul Swamy***, Zuxin Liu, Ding Zhao, Sanjiban Choudhury, Zhiwei Steven Wu, *Learning Shared Safety Constraints from Multi-task Demonstrations*, Neural Information Processing Symposium (NeurIPS), 2023. [\[Site\]](#)

Gokul Swamy, Sanjiban Choudhury, J. Andrew Bagnell, Zhiwei Steven Wu, *Inverse Reinforcement Learning without Reinforcement Learning*, Internat. Conf. on ML (ICML), 2023. [\[Site\]](#)

Gokul Swamy, Sanjiban Choudhury, J. Andrew Bagnell, Zhiwei Steven Wu, *Sequence Model Imitation Learning with Unobserved Contexts*, Neural Information Processing Symposium (NeurIPS), 2022. [\[Site\]](#)

Gokul Swamy*, Nived Rajaraman*, Matt Peng, Sanjiban Choudhury, J. Andrew Bagnell, Zhiwei Steven Wu, Jiantao Jiao, Kannan Ramchandran, *Mimimax Optimal Imitation Learning via Replay Estimation*, Neural Information Processing Symposium (NeurIPS), 2022. [\[Site\]](#)

Gokul Swamy, Sanjiban Choudhury, J. Andrew Bagnell, Zhiwei Steven Wu, *Causal Imitation Learning under Temporally Correlated Noise*, **Oral Presentation (2.1%)**, Int. Conf. on ML (ICML), 2022. [\[Site\]](#)

Gokul Swamy, Sanjiban Choudhury, J. Andrew Bagnell, Zhiwei Steven Wu, *Of Moments and Matching: A Game-Theoretic Framework for Closing the Imitation Gap*, Internat. Conf. on ML (ICML), 2021. [\[Site\]](#)

Gokul Swamy, Siddharth Reddy, Sergey Levine, Anca D. Dragan, *Scaled Autonomy: Enabling Human Operators to Control Robot Fleets*, International Conf. on Robotics and Automation (ICRA), 2020. [\[PDF\]](#)

Gokul Swamy, Jens Schulz, Rohan Choudhury, Dylan Hadfield-Menell, Anca D. Dragan, *On the Utility of Model Learning in HRI*, International Conf. on Human-Robot Interaction (HRI), 2019. [\[PDF\]](#)

Gokul Swamy, Sanjiban Choudhury, J. Andrew Bagnell, Zhiwei Steven Wu, *Complementing a Policy with a Different Observation Space*, Interactive Learning with Implicit Human Feedback, Spurious Correlations, Invariances, and Stability Workshops @ ICML 2023. [\[PDF\]](#)

Gokul Swamy, Sanjiban Choudhury, J. Andrew Bagnell, Zhiwei Steven Wu, *Game Theoretic Algorithms for Conditional Moment Matching*, Neglected Assumptions in Causal Inference @ ICML 2021. [\[PDF\]](#)

TALKS

Efficient Reductions for Reinforcement Learning from Human Feedback

- ▷ INFORMS Annual Meeting, 2024
- ▷ Foerster Lab for AI Research (FLAIR) @ Oxford, 2024
- ▷ Foundations Team @ Google DeepMind, 2024
- ▷ Adaptive Learning in Complex Environments @ TTIC, 2024
- ▷ Emma Brunskill's Group @ Stanford, 2024
- ▷ Learning Theory Team @ Google Research, 2024
- ▷ NLP Lunch @ Google Research, 2024

Efficient Reductions for Interactive Learning from Implicit Feedback

- ▷ Control and Learning Seminar @ CMU, 2024
- ▷ CSAIL Seminar @ MIT, 2024
- ▷ Northeast Systems and Control Workshop, 2024

Efficient Reductions for Inverse Reinforcement Learning

- ▷ Safe Robotics Lab @ Princeton, 2024
- ▷ General Purpose Robotics and AI Lab @ NYU, 2024
- ▷ Coordinated Science Laboratory Student Conference @ UIUC, 2024
- ▷ Robotics Seminar @ Cornell, 2024

REBEL: Reinforcement Learning by Regressing Relative Rewards

- ▷ Aligning RL Experimentalists and Theorists Workshop @ ICML, 2024

Causal Confounds in Imitation Learning

- ▷ Causality for Robotics Workshop @ IROS, 2023

Efficient Algorithms for Interactive Imitation Learning

- ▷ Shock Lab @ University of Cape Town, 2023

Learning Shared Safety Constraints from Multi-task Demonstrations

- ▷ Oral Presentation at the Adversarial ML Workshop @ ICML 2023

An Interactive Workshop on Undergraduate Research Mentorship for Graduate Students

- ▷ Eberly Center Teaching & Learning Summit, 2022

Learning Modular World Models

- ▷ MSR Montreal, 2022

On Interaction, Imitation and Causation

- ▷ Approximately Correct Machine Intelligence Lab @ CMU, 2023
- ▷ Guest Lecture, Learning for Robot Decision Making @ Cornell, 2022
- ▷ Personal Autonomous Robotics Lab @ UT Austin, 2022
- ▷ Reinforcement Learning Discussion Group @ MSR-NYC, 2022
- ▷ Robots Perceiving and Doing Lab @ CMU, 2022

Causal Imitation Learning under Temporally Correlated Noise

- ▷ Long Talk at ICML 2022
- ▷ Oral Presentations at Causal Sequential Decision Making, Offline RL, and Safe and Robust Control of Uncertain Systems Workshops @ NeurIPS 2021

Of Moments and Matching: A Game Theoretic Framework for Closing the Imitation Gap

- ▷ Robots Perceiving and Doing Lab @ CMU, 2021

Leveraging Human Input for Training Self-Driving Cars

- ▷ Guest Lecture, Human-AI Interaction @ CMU, 2022
- ▷ Guest Lecture, Human-AI Interaction @ CMU, 2020

PRESS COVERAGE

Inverse Reinforcement Learning Without Reinforcement Learning

- ▷ TWiML AI Podcast, 2023.

[\[Link\]](#)