Gokul Swamy

CONTACT

gswamy@cmu.edu
https://gokul.dev

EDUCATION

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

University of California, Berkeley

M.S. in Computer Science, Thesis: [Learning with Humans in the Loop] **Aug. 2019 - May 2020**
B.S. in Electrical Engineering and Computer Science, High Honors **Aug. 2016 - May 2019**

RESEARCH EXPERIENCE

Robotics Institute @ CMU, Pittsburgh, PA
Graduate Student Researcher **Sep. 2020 - Present**
Collaborating with Profs. Drew Bagnell and Steven Wu on *efficient interactive learning with unobserved confounders*. Currently researching imitation learning, game theory, and causal inference with applications to self-driving and recommender systems. Published at ICML 2021, 2022, 2023 and NeurIPS 2022 (2x).

Berkeley Artificial Intelligence Research Lab, Berkeley, CA
Graduate Student Researcher **Jan. 2018 - May 2020**
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.

PROFESSIONAL EXPERIENCE

Google Research, Seattle, WA
Student Researcher **May 2023 - Aug. 2023**
Collaborating with Dr. Alekh Agarwal on game-theoretic algorithms for reinforcement learning from human feedback for better fine-tuning of large language models.

Microsoft Research, Montreal, CA
Graduate Research Intern **June 2022 - Aug. 2022**
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.

Aurora, Pittsburgh, PA
Motion Planning ML Intern **May 2020 - Aug. 2020**
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.

NVIDIA, Santa Clara, CA
Autonomous Vehicles Perception Intern **May 2019 - Aug. 2019**
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.

SpaceX, Hawthorne, CA
Data Engineering Intern **May 2018 - Aug. 2018**
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.

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.

Conference Publications

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%)**, Internat. 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]

Workshop Papers

Konwoo Kim*, Gokul Swamy*, Zuxin Liu, Ding Zhao, Sanjiban Choudhury, Zhiwei Steven Wu, *Learning Shared Safety Constraints from Multi-task Demonstrations*, Interactive Learning with Implicit Human Feedback, Adversarial ML Workshops @ ICML 2023. [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]

Activities

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 T-RO, International Journal of Robotics Research (IJRR)
Conference Reviewer: NeurIPS 2021/2022/2023, ICRA 2022, ICML 2023
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

Mentorship

Konwoo Kim, Learning safety constraints from multi-task demonstrations (next: Jump Trading) 2022–2023
▷ Won 2023 CMU SCS Allen Newell Award for Excellence in Undergraduate Research.
Matt Peng, Practical, minimax-optimal online imitation learning (next: Applied Intuition) 2021–2022

Honors

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

Press Coverage

Inverse Reinforcement Learning Without Reinforcement Learning
- ▷ TWiML AI Podcast, 2023. [Link]

Talks

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 University, 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

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

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