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 2020B.S. in Electrical Engineering and Computer Science, High HonorsAug. 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 interactive learning from implicit human feedback.       Currently researching imitation learning, game theory, and causal inference with applications to autonomous         vehicles, large language models, and recommender systems. Published at ICML 2021, 2022, 2023 and       NeurIPS 2022 (2x), 2023.
	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 - Nov. 2023         Collaborated with Drs. Alekh Agarwal, Chris Dann, and Rahul Kidambi on game-theoretic algorithms for reinforcement learning from conflicting 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 InternMay 2019 - Aug. 2019Collaborated with Dr. Trung Pham on single-image weakly-supervised 3D structure estimation of intersectionentry/exit lines. Designed CNN to recover lines in 3D that learned from 2D key points and 3D geometricconstraints. 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 navi- gable 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.
	Intuit, San Diego, CA Software Engineering Intern May 2017 - Aug. 2017 Worked on using CNNs to classify the sentiments of product reviews. Created themable iOS UI framework and Java/PHP/React-based internal tools.

Teaching Experience	Carnegie Mellon University, Pittsburgh, PA         Teaching Assistant         TA for 10-732: Robustness and Adaptation in Shifting Environments w/ Prof. Zachary Lipton.	ec. 2022 Covered
	causally-structured and feedback-driven distribution shifts, online adaptation, and adv. robustnes	s. [Site]
	University of California, Berkeley, Berkeley, CA	oo 2010
	Helped teach CS 188: Introduction to AI and CS 189: Introduction to Machine Learning. Gave gue and created worksheet walkthrough videos.	ec. 2019 est lecture
	Course Instructor Aug. 2018 - M	fay 2019
	Created material for and taught course twice on societal impacts and ethical considerations of AI topics like automation, bias in AI, data privacy, artificially generated data, and human-compatible	, covering e AI.
ACTIVITIES	Communications Chair, Reinforcement Learning Conference (RLC)	2024
	Workshop Organizer, Reinforcement Learning Beyond Rewards (RLC)	2024
	Teaching & Learning Summit Advisory Board, Eberly Center @ CMU	2022- 2022
	Undergraduate Research Engagement Working Group, CMU SCS Dean's Advisory Committee	2022
	> Put together interactive workshop for graduate students on best practices for mentoring under	graduates
	on research projects. Conducted IRB-approved study to measure training effects. Results se	lected for
	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 Informatic	on Theory
	Conference Reviewer: NeurIPS 2021-2023, ICRA 2022/2024, ICML 2023/2024, ICLB 2024, IROS	2024
	Workshop Reviewer: Strategic ML @ NeurIPS 2021, Real World Reinforcement Learning @ Neur	IPS 2022,
	Interactive Learning with Implicit Human Feedback @ ICML 2023, Frontiers for Learning, Dyna Control @ ICML 2023, Robot Learning Workshop @ NeurIPS 2023	mics, and
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-
	Jinwgu Tang (CMU), Multi-agent imitation learning.	2023-
	Silvia Sanora (Oxford). Faster inverse reinforcement learning via evolutionary algorithms.	2023-
	Juntao Ren (Cornell), Hybrid algorithms for inverse reinforcement learning.	2023-
	Konwoo Kim (CMU), Learning safety constraints from multi-task demos. (next: Stanford PhD) 2	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) 2	2021-2022
Honors	Invited to Jane Street Graduate Research Fellowship Research Workshop	2024
110110110	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 Sci	ence 2023
	10p Reviewer, Neurir's Finalist for IBM PhD Fellowshin 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 Instit	ute 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/Cont tistical Methods in ML, Computational Game Solving, Intermediate Statistics, Philosophical Found	trols, Sta- dations 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 System Advanced Robotics, Natural Language Processing	is Theory,

Pre-prints	Juntao Ren <sup>*</sup> , Gokul Swamy <sup>*</sup> , Zhiwei Steven Wu, J. Andrew Bagnell, Sanjiban Choudhury, <i>Hybrid Inverse Reinforcement Learning</i> . Presented at NeurIPS 2023 Workshop on Robot Learning and ISAIM Special Session on Deep Reinforcement Learning. [PDF]
	Gokul Swamy, Christoph Dann, Rahul Kidambi, Zhiwei Steven Wu, Alekh Agarwal, A Minimaximalist Approach to Reinforcement Learning from Human Feedback. Presented at DIMACS Workshop on Foundation Models, Large Language Models, and Game Theory. [PDF]
Conference Publications	Konwoo Kim <sup>*</sup> , Gokul Swamy <sup>*</sup> , Zuxin Liu, Ding Zhao, Sanjiban Choudhury, Zhiwei Steven Wu, <i>Learn-</i> <i>ing Shared Safety Constraints from Multi-task Demonstrations</i> , 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 Learn- ing with Unobserved Contexts, Neural Information Processing Symposium (NeurIPS), 2022. [Site]
	Gokul Swamy <sup>*</sup> , Nived Rajaraman <sup>*</sup> , Matt Peng, Sanjiban Choudhury, J. Andrew Bagnell, Zhiwei Steven Wu, Jiantao Jiao, Kannan Ramchandran, <i>Mimimax Optimal Imitation Learning via Replay Estimation</i> , Neural Information Processing Symposium (NeurIPS), 2022. [Site]
	Gokul Swamy, Sanjiban Choudhury, J. Andrew Bagnell, Zhiwei Steven Wu, <i>Causal Imitation Learning under Temporally Correlated Noise</i> , <b>Oral Presentation (2.1%)</b> , 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, <i>Scaled Autonomy: Enabling Human Operators to Control Robot Fleets</i> , 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 and Posters	Silvia Sapora, Chris Lu <sup>*</sup> , Gokul Swamy <sup>*</sup> , Yee Whye Teh, Jakob Nicolaus Foerster <i>EvIL: Evolution Strategies</i> for Generalisable Imitation Learning, NeurIPS 2023 Workshop on Robot Learning [PDF]
	Gokul Swamy, Sanjiban Choudhury, J. Andrew Bagnell, Zhiwei Steven Wu, <i>Complementing a Policy with a Different Observation Space</i> , 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, <i>Game Theoretic Algorithms for Conditional Moment Matching</i> , Neglected Assumptions in Causal Inference @ ICML 2021. [PDF]

Efficient Reductions for Reinforcement Learning from Human Feedback

▷ Learning Theory Team @ Google Research, 2024

 $\triangleright$  NLP Lunch @ Google Research, 2024

Efficient Reductions for Interactive Learning from Implicit Feedback

▷ Control and Learning Seminar @ CMU, 2024

- ▷ CSAIL Seminar @ MIT, 2024
- ▷ Emma Brunskill's Group @ Stanford, 2024
- $\triangleright$  Northeast Systems and Control Workshop, 2024

Efficient Reductions for Inverse Reinforcement Learning

- ▷ General Purpose Robotics and AI Lab @ NYU, 2024
- ▷ Coordinated Science Laboratory Student Conference @ UIUC, 2024
- $\triangleright$  Robotics Seminar @ Cornell, 2024
- Causal Confounds in Imitation Learning

 $\triangleright$  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

 $\triangleright$  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

- $\triangleright$  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

Inverse Reinforcement Learning Without Reinforcement Learning ▷ TWiML AI Podcast, 2023.

[Link]