Josiah P. Hanna

Computer Sciences Department
The University of Wisconsin–Madison
Madison, WI 53706

jphanna@cs.wisc.edu
(608) 262-0719
https://www.cs.wisc.edu/~jphanna

EDUCATION

The University of Texas at Austin — Austin, TX
Ph.D. in Computer Science — 2014 – 2019

- ◇ Advisor: Prof. Peter Stone
- ◇ Dissertation: Data Efficient Reinforcement Learning with Off-policy and Simulated Data
- ◇ Research: Artificial Intelligence, Reinforcement Learning, Robotics

The University of Kentucky — Lexington, KY
B.S. in Computer Science and Mathematics — 2010 – 2014

- ◇ GPA: 4.0
- ◇ Summa Cum Laude

WORK EXPERIENCE

University of Wisconsin – Madison, Computer Sciences Department — Madison, WI
Assistant Professor — August 2021 – Present

- ◇ Lead research in reinforcement learning and robotics.
- ◇ Supervise Ph.D., MS, and undergraduate student researchers.
- ◇ Teach courses in artificial intelligence and reinforcement learning.

University of Edinburgh, School of Informatics — Edinburgh, U.K.
Postdoctoral Research Associate — January 2020 – July 2021

- ◇ Advised by Prof. Stefano Albrecht.
- ◇ Conducted research in reinforcement learning and multi-agent systems.
- ◇ Informally advised three Ph.D. students and two M.Sc. students.

FiveAI, Ltd. — Edinburgh, U.K.
Consultant — March 2020 – June 2021

- ◇ Developed planning and prediction algorithms for autonomous vehicles.
- ◇ Supervised two research interns.

Google, Inc. — Mountain View, CA
Software Engineering Intern — May 2017 – Sept. 2017

- ◇ Advised by Craig Boutillier.
- ◇ Developed reinforcement learning algorithms with application to Google products.

University of Texas at Austin — Austin, TX
IBM PhD Research Fellow — September 2018 – December 2019

- ◇ Developed algorithms for correcting inaccuracy from random sampling in reinforcement learning.
- ◇ Mentored three undergraduate students on projects relating to reinforcement learning and optimization.

Teaching Assistant — September 2017 – December 2018

- ◇ CS 343H – Honors Artificial Intelligence
- ◇ CS 393R – Autonomous Robotics

NSF Graduate Research Fellow — August 2014 – August 2017

- ◇ Developed an algorithm allowing robot learning in simulation to transfer to the real world.

- ◇ Developed algorithms for evaluating the performance of untested robot behaviors.
- ◇ Developed a novel tolling scheme for autonomous vehicles that reduced traffic congestion in road networks.

Computer Science Department, University of Kentucky Lexington, KY
Undergraduate Research Assistant May 2013 – May 2014
- ◇ Investigated leveraging structure in artificial intelligence planning under uncertainty problems.

Laboratoire d'Informatique de Paris 6 Paris, France
Research Intern May 2012 – Aug. 2012
- ◇ Developed algorithms for solving multi-objective planning problems.

College of Arts and Sciences, University of Kentucky Lexington, KY
Software Developer June 2011 – May 2012
- ◇ Developed a clustering algorithm for student academic data.

HONORS AND AWARDS

- ◇ RoboCup Standard Platform League Challenge Shield Division 3rd Place 2023
- ◇ Madison Teaching and Learning Excellence Fellow 2022
- ◇ IBM Ph.D. Fellowship 2018
- ◇ Robocup Standard Platform League Runner-Up 2016
- ◇ RoboCup 3D Simulation League Champions 2015
- ◇ National Science Foundation Graduate Research Fellowship 2014
- ◇ Barry M. Goldwater Scholarship 2013
- ◇ Phi Kappa Phi 2013
- ◇ Astronaut Scholarship 2013
- ◇ Duncan E. Clarke Memorial Scholarship 2012
- ◇ Barry M. Goldwater Scholarship, Honorable Mention 2012
- ◇ Tau Beta Pi 2012

TEACHING

- ◇ Assistant Professor at the University of Wisconsin–Madison: August 2021 – Present
  - CS 540: Introduction to Artificial Intelligence Fall 2021
  - CS 839: Advanced Topics in Reinforcement Learning Fall 2022
  - CS 540: Introduction to Artificial Intelligence Spring 2023
  - CS 760: Machine Learning Fall 2023

SERVICE ACTIVITIES

University and Department Service

- ◇ Computer Sciences Graduate Admissions Committee 2021 – 2024
- ◇ Computer Sciences Professional Masters Admissions Committee 2023

Reviewing

- ◇ Action Editor, Machine Learning Journal (MLJ) 2023
- ◇ Editorial Board, Machine Learning Journal (MLJ) 2021 – present
- ◇ Meta-Reviewer, AAAI 2023
- ◇ Area Chair, NeurIPS 2023
- ◇ Program Committee, ICML 2023
- ◇ Reviewer, Robotics and Automation Letters (RA-L) 2023
- ◇ Reviewer, The Artificial Intelligence Journal (AIJ) 2022
- ◇ Meta-Reviewer, AAAI 2022
- ◇ Program Committee, NeurIPS 2022
- ◇ Senior Program Committee, CoLLAS 2022
- ◇ AISTATS Mentorship Program Mentor 2022
- ◇ Program Committee, ICML 2022
- ◇ Reviewer, RSS 2021
- ◇ Reviewer, IROS 2021
- ◇ Program Committee, ICML 2021
- ◇ Program Committee, AAMAS Workshop on Adaptive Learning Agents (ALA) 2021
- ◇ Reviewer, Journal of Artificial Intelligence Research (JAIR) 2020
- ◇ Program Committee, ICML 2020
- ◇ Program Committee, ICML 2019
- ◇ Program Committee, AAMAS 2019
- ◇ Program Committee, AAAI Conference on Artificial Intelligence 2019
- ◇ Reviewer, Neural Information Processing Systems (NeurIPS) 2018
- ◇ Reviewer, International Conference on Machine Learning (ICML) 2018
- ◇ Program Committee, AAAI Spring Symposium on Data Efficent Reinforcement Learning 2018
- ◇ Reviewer, Neural Information Processing Systems (NeurIPS) 2017
- ◇ Program Committee, Workshop on Scaling Up Reinforcement Learning 2017
- ◇ Review Assistant, International Joint Conference on Artificial Intelligence (IJCAI) 2017
- ◇ Reviewer, Neural Information Processing Systems (NeurIPS) 2016

Conference, Workshop, and Competition Organization

- ◇ Reinforcement Learning Conference Workshop Co-Chair 2024
- ◇ RoboCup Symposium Co-Chair 2024
- ◇ RoboCup Standard Platform League, Technical Committee 2023 – 2024
- ◇ RoboCup Standard Platform League, Organizing Committee 2018

Other Service

- ◇ U.S. Robotics Research Roadmapping Workshop Contributor 2023
- ◇ WISCERS Faculty Mentor 2022, 2024
- ◇ Mercile Lee Scholars Mentor 2021 – 2022

THESIS COMMITTEES

Doctoral Committee Member: (University of Wisconsin – Madison)

- ◇ Yuzhe Ma, Computer Sciences. Supervisor: Jerry Zhu.
- ◇ Young Wu, Computer Sciences. Supervisor: Jerry Zhu.
- ◇ Matt Dutson, Computer Sciences. Supervisor: Mohit Gupta.
- ◇ Toygun Basaklar, Electrical and Computer Engineering. Supervisor: Umit Ogras.
- ◇ Jeremey McMahan, Computer Sciences. Supervisor: Jerry Zhu
- ◇ Yeping Wang, Computer Sciences. Supervisor: Michael Gleicher

OTHER ADVISING

- ◇ Current Wisconsin PhD Students: Brahma Pavse, Subhojyotee Mukherjee, Nicholas Corrado, Adam Labiosa.
- ◇ Current Wisconsin PhD Independent Study: Yunfu Deng, Abhinav Harish, Andrew Wang, Will Cong
- ◇ Current Wisconsin Undergraduate Students: Ben Hong, Chen Li, Kwasi Debrah-Pinamang, Edbert Wang, Lucas Poon
- ◇ Past Wisconsin Undergraduate Students: Yuxiao Qu (December 2022), Adhit Sankaran (May 2022), Will Cong (December 2022), Paul Pak (May 2023).
- ◇ Past Wisconsin MS Students: Yoon Chae Na (December 2022), Arun Ravi (December 2022), Shreyansh Sharma (December 2022), Duohan Zhang (May 2023), John Balis (December 2023).
- ◇ University of Alberta M.Sc. Thesis: Hager Radi (2022)
- ◇ Five AI Interns: Elliott Fosong (2020), Arrasy Rahman (2021)
- ◇ University of Edinburgh M.Sc. Thesis: Rujie (Jerry) Zhong (2021), Panagiotis Kyriakou (2021)
- ◇ UT Austin MS Thesis: Brahma Pavse (2019-2020)
- ◇ UT Austin Undergraduate Research: Xiang Gu (2018), John Fang (2018-2019), Harsh Goyal (2018-2019)

PUBLICATIONS

Works Under Review

- ◇ Corrado, N, Qu, Y, Balis, J, Labiosa, A, **Hanna, J.P.**. "Guided Data Augmentation for Offline Reinforcement Learning and Imitation Learning." *To be Submitted to the IEEE Conference on Intelligent Robots and Systems (IROS)*, 2024.
- ◇ Corrado, N, **Hanna, J.P.**. "On-Policy Policy Gradient Reinforcement Learning Without On-Policy Sampling." *Under Review at the International Conference on Machine Learning (ICML)*, 2024.
- ◇ Pavse, B, Zurek, M, Chen, Y, Xie, Q, **Hanna, J.P.**. "Tackling Unbounded State Spaces in Continuing Task Reinforcement Learning." In *Under Review at the International Conference on Machine Learning (ICML)*, 2024.
- ◇ Mukherjee, S, Xie, Q, **Hanna, J.P.**, Nowak, R. "Pretraining Decision Transformers with Reward Prediction for In-Context Structured Bandit Learning." *Under Review at the International Conference on Machine Learning*, 2024.
- ◇ Kwon, J, Yang, L, **Hanna, J.P.**, Nowak, R. "Future Prediction Can Be a Strong Evidence of Good History Representation in Partially Observable Environments." *Under Review at the International Conference on Machine Learning (ICML)*, 2024.
- ◇ **Hanna, J.P.**, Chandak, Y., Thomas, P., White, M., Stone, P., Niekum, S. "Data-Efficient Policy Evaluation Through Behavior Policy Search." *Under Review at the Journal of Machine Learning Research (JMLR)*, 2023.

Journal Articles

- ◇ **Hanna, J.P.**, Niekum, S., Stone, P. "Importance Sampling in Reinforcement Learning with an Estimated Behavior Policy." In *Machine Learning (MLJ)*, 2021.
- ◇ **Hanna, J.P.**, Desai, S., Karnan, H., Warnell, G., Stone, P. "Grounded Action Transformation for Sim-to-Real Reinforcement Learning." In *Machine Learning (MLJ): Special Issue on Reinforcement Learning for Real Life*, 2021.
- ◇ Pavse, B.S., Torabi, F., **Hanna, J.P.**, Warnell, G., Stone, P. "RIDM: Reinforced Inverse Dynamics Modeling for Learning From a Single Observed Demonstration." In *IEEE Robotics and Automation Letters*, 2020.
- ◇ Sharon, G., Levin, M.W., **Hanna, J.P.**, Rambha, T., Boyles, S.D., Stone, P. "Network-wide Adaptive Tolling for Connected and Automated Vehicles." In *Transportation Research Part C*, 2017.
- ◇ Chen, T, Kockelman, K, **Hanna, J.P.**. "Operations of a Shared, Autonomous, Electric Vehicle Fleet: Implications of Vehicle & Charging Infrastructure Decisions." In *Transportation Research Part A: Policy and Practice*, 2016.

Refereed Conference Proceedings

- ◇ Mukherjee, S, Xie, Q, **Hanna, J.P.**, Nowak, R. "SPEED: Experimental Design for Policy Evaluation in Linear Heteroscedastic Bandits." In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2024.
- ◇ Corrado, N, **Hanna, J.P.**. "Understanding When Dynamics-Invariant Data Augmentations Benefit Model-Free Reinforcement Learning Updates." In *Proceedings of the International Conference on Representation Learning (ICLR)*, 2024.
- ◇ Dunion, M, McInroe, T, Luck, K, **Hanna, J.P.**, Albrecht, S. "Conditional Mutual Information for Disentangled Representations in ReinforcementLearning." In *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, 2023.
- ◇ Mukherjee, S, Xie, Q, **Hanna, J.P.**, Nowak, R. "Multi-task Representation Learning for Pure Exploration in Bilinear Bandits." In *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, 2023.
- ◇ Pavse, B, **Hanna, J.P.**. "State-Action Similarity-Based Representations for Off-Policy Evaluation." In *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, 2023.
- ◇ Dunion, M., McInroe, T., Luck, K.S., **Hanna, J.P.**, Albrecht, S.V. "Temporal Disentanglement of Representations for Improved Generalisation in Reinforcement Learning." In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2023.
- ◇ Pavse, B., **Hanna, J.P.**. "Scaling Marginalized Importance Sampling To High-Dimensional State-Spaces Via State Abstraction." In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, 2023.
- ◇ Zhong, R., Zhang, D., Schäfer, L., Albrecht, S.V., **Hanna, J.P.**. "Robust On-Policy Sampling for Data-Efficient Policy Evaluation." In *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- ◇ Mukherjee, S., **Hanna, J.P.**, Nowak, R. "ReVar: Strengthening Policy Evaluation Via Reduced Variance Sampling." In *Proceedings of the 38th International Conference on Uncertainty in Artificial Intelligence (UAI)*, 2022.
- ◇ Corrado, N., Qu, Y., **Hanna, J.P.**. "Simulation-Acquired Latent Action Spaces for Dynamics Generalization." In *Proceedings of the 1st Conference on Lifelong Learning Agents (CoLLAs)*, 2022.
- ◇ Schäfer, L., Christianos, F., **Hanna, J.P.**, Albrecht, S.V. "Decoupled Reinforcement Learning To Stabilise Intrinsically-Motivated Exploration." In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2022.

- ◇ Ahmed, I.H., **Hanna, J.P.**, Fosong, E., Albrecht, S.V. "Towards Quantum-Secure Authentication and Key Agreement Via Abstract Multi-Agent Interaction." In *International Conference on Practical Applications of Agents and Multi-Agent Systems (PAAMS)*, 2021.
- ◇ **Hanna, J.P.**, Rahman, A., Fosong, E., Eiras, F., Dobre, M., Redford, J., Ramamoorthy, S., Albrecht, S.V. "Interpretable Goal Recognition in the Presence of Occluded Factors for Autonomous Vehicles." In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021.
- ◇ Dey, S., Pendurkar, S., Sharon, G., **Hanna, J.P.**. "A Joint Imitation-Reinforcement Learning Framework for Reduced Baseline Regret." In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021.
- ◇ Desai, S, Durugkar, I, Karnan, H, Warnell, G, **Hanna, J.P.**, Stone, P. "An Imitation From Observation Approach To Transfer Learning with Dynamics Mismatch." In *Proceedings of the 33rd Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- ◇ Karnan, H, Desai, S, **Hanna, J.P.**, Warnell, G, Stone, P. "Reinforced Grounded Action Transformation for Sim-to-Real Transfer." In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020.
- ◇ Desai, S, Karnan, H, **Hanna, J.P.**, Warnell, G, Stone, P. "Stochastic Grounded Action Transformation for Robot Learning in Simulation." In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020.
- ◇ Pavse, B, Durugkar, I, **Hanna, J.P.**, Stone, P. "Reducing Sampling Error in Batch Temporal Difference Learning." In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, 2020.
- ◇ Ault, J., **Hanna, J.P.**, Sharon, G. "Learning an Interpretable Traffic Signal Control Policy." In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2020.
- ◇ **Hanna, J.P.**, Niekum, S, Stone, P. "Importance Sampling Policy Evaluation with an Estimated Behavior Policy." In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, 2019.
- ◇ **Hanna, J.P.**, Stone, P. "Reducing Sampling Error in the Monte Carlo Policy Gradient Estimator." In *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2019.
- ◇ **Hanna, J.P.**, Sharon, G., Boyles, S.D., Stone, P. "Selecting Compliant Agents for Opt-in Micro-tolling." In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, 2019.
- ◇ Chen, H., An, B., Sharon, G., **Hanna, J.P.**, Stone, P., Miao, C., Soh, Y.C. "DyETC: Dynamic Electronic Toll Collection for Traffic Congestion Alleviation." In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, 2018.
- ◇ **Hanna, J.P.**, Thomas, P., Stone, P., Niekum, S. "Data-Efficient Policy Evaluation Through Behavior Policy Search." In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, 2017.
- ◇ Sharon, G., **Hanna, J.P.**, Rambha, T., Levin, M.W., Albert, M., Boyles, S.D., Stone, P. "Real-time Adaptive Tolling Scheme for Optimized Social Welfare in Traffic Networks." In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2017)*, 2017.
- ◇ **Hanna, J.P.**, Stone, P., Niekum, S. "Bootstrapping with Models: Confidence Intervals for Off-Policy Evaluation." In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2017.
- ◇ **Hanna, J.P.**, Stone, P. "Grounded Action Transformation for Robot Learning in Simulation." In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, 2017.
- ◇ Perny, P., Weng, P., Goldsmith, J., **Hanna, J.P.**. "Approximation of Lorenz-Optimal Solutions in Multiobjective Markov Decision Processes." In *Proceedings of the International Conference on Uncertainty in Artificial Intelligence (UAI)*, 2013.

Book Chapters / Refereed Workshops / Symposium Proceedings

- ◇ Mukherjee, S., Xie, Q., **Hanna, J.P.**, Nowak, R. "SPEED: Experimental Design for Policy Evaluation in Linear Heteroscedastic Bandits." In *ICML Workshop on the Many Facets of Preference-Based Learning*, 2023.
- ◇ Pavse, B., **Hanna, J.P.**. "Scaling Marginalized Importance Sampling To High-Dimensional State-Spaces Via State Abstraction." In *NeurIPS Workshop on Offline Reinforcement Learning (OfflineRL)*, 2022.
- ◇ Dunion, M., McInroe, T., Luck, K.S., **Hanna, J.P.**, Albrecht, S.V. "Temporal Disentanglement of Representations for Improved Generalisation in Reinforcement Learning." In *NeurIPS Workshop on Deep Reinforcement Learning*, 2022.
- ◇ Zhang, C, Papaemmanouil, O, **Hanna, J.P.**, Akella, A. "Multi-agent Databases Via Independent Learning." In *Proceedings of the 4th International Workshop on Applied AI for Database Systems and Applications*, 2022.
- ◇ Lobo, E., Chandak, Y., Subramanian, D., **Hanna, J.P.**, Petrik, M. "Behavior Policy Search for Risk Estimators in RL." In *NeurIPS Workshop on Safe and Robust Control of Uncertain Systems*, 2021.
- ◇ Zhong, R., **Hanna, J.P.**, Schäfer, L., Albrecht, S.V. "Robust On-Policy Data Collection for Data-Efficient Policy Evaluation." In *NeurIPS Workshop on Offline Reinforcement Learning (OfflineRL)*, 2021.
- ◇ Radi, H, **Hanna, J.P.**, Stone, P, Taylor, M. "Safe Evaluation for Offline Learning: Are We Ready To Deploy?" In *NeurIPS Workshop on Deployable Decision Making in Embodied Systems (DDM)*, 2021.
- ◇ Schäfer, L., Christianos, F., **Hanna, J.P.**, Albrecht, S.V. "Decoupling Exploration and Exploitation in Reinforcement Learning." In *ICML Workshop on Unsupervised Reinforcement Learning (URL)*, 2021.
- ◇ Pavse, B.S., **Hanna, J.P.**, Durugkar, I., Stone, P. "On Sampling Error in Batch Action-Value Prediction Algorithms." In *In the Offline Reinforcement Learning Workshop at Neural Information Processing Systems (NeurIPS)*, 2020.
- ◇ Pavse, B.S., Torabi, F., **Hanna, J.P.**, Warnell, G., Stone, P. "RIDM: Reinforced Inverse Dynamics Modeling for Learning From a Single Observed Demonstration." In *Imitation, Intent, and Interaction (I3) Workshop at ICML 2019*, 2019.
- ◇ **Hanna, J.P.**, Stone, P. "Towards a Data Efficient Off-Policy Policy Gradient." In *AAAI Spring Symposium on Data Efficient Reinforcement Learning*, 2018.
- ◇ Menashe, J., Kelle, J., Genter, K., **Hanna, J.P.**, Liebman, E., Narvekar, S., Zhang, R., Stone, P. "Fast and Precise Black and White Ball Detection for RoboCup Soccer." In *RoboCup-2017: Robot Soccer World Cup XXI*, 2017.
- ◇ MacAlpine, P., **Hanna, J.P.**, Liang, J., Stone, P. "UT Austin Villa: RoboCup 2015 3D Simulation League Competition and Technical Challenges Champions." In *RoboCup-2015: Robot Soccer World Cup XIX*, 2016.
- ◇ **Hanna, J.P.**, Albert, M., Chen, D., Stone, P. "Minimum Cost Matching for Autonomous Carsharing." In *Proceedings of the 9th IFAC Symposium on Intelligent Autonomous Vehicles (IAV 2016)*, 2016.
- ◇ Guerin, J.T., **Hanna, J.P.**, Ferland, L., Mattei, N., Goldsmith, J. "The Academic Advising Planning Domain." In *Proceedings of the 3rd Workshop on the International Planning Competition at ICAPS*, 2012.

FUNDING

Current Support

- ◇ PI: American Family Data Science Funding Initiative
  - – *Counterfactual Evaluation of Sequential Decision Policies*
  - – Award amount: $96,000
  - – Dates: September 1, 2022 – August 31, 2023
- ◇ PI: American Family Data Science Funding Initiative
  - – *Learning What is Relevant for Counterfactual Policy Evaluation*
  - – Award amount: $99,999
  - – Dates: September 1, 2023 – August 31, 2024
- ◇ PI: Sandia University Partnership Network
  - – *Discovery of Conductive Inks and Electronic Devices co-Designed with Closed-Loop, Autonomous, Reinforcement Ecosystems*
  - – Award amount: $49,450
  - – Dates: November 13, 2023 – September 30, 2024

TALKS

- ◇ AAAI Conference New Faculty Highlight. *Scaling Offline Evaluation of Reinforcement Learning Agents through Abstraction.* February 2024.
- ◇ Tulane University Computer Science Department Colloquium. *Towards Reinforcement Learning for Real-time and Dynamic Robotic Tasks.* December 2023.
- ◇ University of Kentucky Keeping Current Seminar. *Towards Reinforcement Learning for Real-time and Dynamic Robotic Tasks.* November 2023.
- ◇ Sony AI. *Towards Data Efficient Monte Carlo Estimates in Reinforcement Learning.* November 2021.
- ◇ University of Wisconsin – Madison SILO Seminar Series. *Towards Data Efficient Monte Carlo Estimates in Reinforcement Learning.* September 2021.
- ◇ University of Wisconsin – Madison Robotics Seminar Series. *Better Prediction for Reinforcement Learning in Robotics and Autonomous Driving.* October 2021.
- ◇ University of Edinburgh AIAI Institute Seminar. *Data Efficient Reinforcement Learning from Reweighted and Simulated Data.* November 2020.
- ◇ University of Wisconsin – Madison SILO Seminar Series. *Data Re-weighting for Data Efficient Reinforcement Learning.* 2020.
- ◇ Microsoft Research Seminar. *Date Efficient Reinforcement Learning for Autonomous Robots* June 2019.
- ◇ AAAI Spring Symposium on Data Efficient Reinforcement Learning, Invited Talk. *Data Efficient Reinforcement Learning with Off-policy and Simulated Data.* April 2018.