Peter Herald Stone Dept. of Computer Science

7208 Lamplight Lane **Austin, TX 78731**

(512) 346–3809

pstone@cs.utexas.edu www.cs.utexas.edu/~pstone

University of Texas at Austin **Austin, TX 78712** (512) 471–9796

RESEARCH STATEMENT

I am a computer scientist with research specialization in artificial intelligence. My long-term research goal is to create complete, robust, autonomous agents that can learn to interact with other intelligent agents in a wide range of complex, dynamic environments. These agents must sense their environment; engage in high-level cognitive decision-making; and then execute their actions in the environment. Moreover, to be effective, they should improve their performance automatically over time and reason explicitly about the behaviors of other agents, including teammates and adversaries. Thus, my research contributions are mainly in the areas of machine learning; autonomous agents and multiagent systems; and robotics. Application domains have included robot soccer, autonomous bidding agents, intelligent traffic management, generalpurpose service robots, and autonomous vehicles.

PROFESSIONAL PREPARATION

• Carnegie Mellon University, Pittsburgh, PA

Ph.D., Computer Science, December 1998.

Dissertation: Layered Learning in Multi-Agent Systems.

Thesis committee: Manuela Veloso (chair), Andrew Moore, Herbert Simon, Victor Lesser.

M.S., Computer Science, December 1995.

• The University of Chicago, Chicago, IL

B.S., Mathematics with honors and a concentration in Computer Science, June 1993.

APPOINTMENTS

- The University of Texas at Austin:
 - Chair of the Department of Computer Science, September 2025 Present.
 - Truchard Foundation Chair in Computer Science, September 2022 Present.
 - Director of Texas Robotics, September 2017 2025.
 - Associate Chair of the Department of Computer Science, September 2015 2025.
 - David Bruton, Jr. Centennial Professor of Computer Science, September 2014 August 2022.
 - *Professor* of Computer Science, September 2012 August 2014.
 - Associate Professor of Computer Science, September 2007 August 2012.
 - Assistant Professor of Computer Sciences, June 2002 August 2007.
- Sony:
 - Chief Scientist, Sony AI, July 2024 Present.
 - Executive Director, Sony Research, July 2024 Present.
 - President, Sony AI America, July 2024 Present.
 - Senior Vice President, Sony Corporation of America, November 2020 Present.
 - Executive Director, Sony AI America, November 2020 June 2024.
- Cogitai, Inc., September 2015 November 2020.

President, COO, and co-Founder. Acquired by Sony.

- The Hebrew University of Jerusalem and Bar Ilan University, September 2008 June 2009. Visiting Professor in the Computer Science Departments.
- **RobotsLab**, September 2012 December 2015.

Consultant. Helped design educational robotics curriculum.

- SAIC, then Leidos April 2013 May 2014.
 - Consultant. Worked on DARPA project pertaining to reinforcement learning.
- Sidley Austin LLP, May July 2009.

Consultant. Patent infringement case expert.

- Corporation for National Research Initiatives (CNRI), May September 2002. Consultant. Developed possibilities for future DARPA programs pertaining to multiagent systems.
- New York University, September 2001 January 2002. *Adjunct Professor* in the Computer Science Department.
- AT&T Labs Research, September 1999 March 2002. Senior Researcher in the Artificial Intelligence Department.
- Carnegie Mellon University, January 1999 August 1999. Postdoctoral Fellow in the Computer Science Department.
- **Perspectives, Inc.**, April 1998 March 1999.
 - Consultant. Created a comprehensive report on the state of the art in multiagent systems.
- Carnegie Mellon University, August 1993 December 1998.

 Graduate Research Assistant. Created a framework by which multiple intelligent agents can learn to act both individually and in coordination with one another in real-time, noisy, collaborative, and adversarial environments. Developed a flexible commitment strategy for interleaving planning and execution in the PRODIGY planner.
- NASA Jet Propulsion Laboratory, May August 1995.

 Summer intern. Worked on automatic planning and scheduling for the New Millennium Project.

AWARDS AND RESEARCH DISTINCTIONS

- Best RoboCup-Themed Paper, ICRA workshop on Roboletics 2.0, 2025.
- ACM/AAAI Allen Newell Award, 2024.
- Outstanding Paper Award, Applications of RL, Reinforcement Learning Conference (RLC), 2024.
- Advisor to Winner of UT Austin's CS Best Undergraduate Thesis Award:
 - Stephane Hatgis-Kessell, 2024
 - Dustin Carlino, 2013.
- Sony Technology Exchange Fair (STEF) Technology Award for GT Sophy in GT 7, 2024.
- Sonv Research CEO Award. 2023.
- ACM/SIGAI Industry Award for Excellence in Artificial Intelligence, July 2022.
- Sony AI CEO Award, 2022.
- Sony Technology Award, 2022.
- Best Paper Award, AAMAS workshop on Adaptive Learning Agents (ALA), May 2022.
- Elected **ACM Fellow** by the Association for Computing Machinery, 2021.
- Named **AAIA Fellow** by the Asia-Pacific Artificial Intelligence Association, 2021.
- IFAAMAS Influential Paper Award, 2020.
- Elected **AAAS Fellow** by the American Association for the Advancement of Science, 2019.
- Minnie Stevens Piper Professorship, in recognition of superior teaching at the college level, 2019.
- Elected **IEEE Fellow** by the Institute of Electrical and Electronics Engineers, 2018.
- **Best Robotics Track Paper**, International Conference on Autonomous Agents and Multiagent Systems (AAMAS), July 2018.
- ACM/SIGAI Autonomous Agents Research Award, 2016.
- College of Natural Sciences **Teaching Excellence Award**, 2015.
- Outstanding Paper Award, Computational Sustainability Track, International Joint Conference on Artificial Intelligence (IJCAI), July 2015.
- Elected to UT Austin Academy of Distinguished Teachers, 2014.
- World Champion team member in 17 RoboCup events: 3D simulator competition, June 2021; July 2019; July 2018; July 2017, July 2016, July 2015, July 2014, June 2012, July 2011; standard platform competition, July 2024, June 2012; simulator coach competition, July 2005, July 2003; simulator competition, August 1999, July 1998; small-size robot competitions, July 1998, August 1997.
- RoboCup **US Open Champion** team leader, standard platform league, April 2018; April 2017, April 2016, April 2012, May 2010, May 2009.

- **Best Contribution Award** at NIPS workshop on Machine Learning for Sustainability (MLSUST), December 2013.
- Best Paper Award, International Conference on Social Robotics (ICSR), October 2013.
- The University of Texas System **Regents' Outstanding Teaching Award**, 2013.
- Leader of **1st-place** teams in the **Trading Agent Competition** (TAC), July 2013; July 2011; June 2010; July 2009; July 2008; May 2006; August 2005; August 2003; October 2001; July 2000.
- Elected **AAAI Fellow** by the Association for the Advancement of Artificial Intelligence, 2012.
- Google Faculty Research Award, 2012.
- Yahoo! Faculty Research and Engagement (FREP) award, 2011.
- Elected **Board Member**, International Machine Learning Society (IMLS), March 2011–2015.
- Advisor to Winner of UT Austin's **Best Dissertation Award** (Kurt Dresner), May 2010.
- **Best Student Paper Award**, International Conference on Autonomous Agents and Multiagent Systems (AAMAS), May 2010.
- Best Student Paper Award, RoboCup Symposium, July 2009.
- William David Blunk Memorial Professorship, in recognition of undergraduate teaching, 2008–09.
- Fulbright Award, 2008–09.
- Guggenheim Fellow, 2008–09.
- Elected **Board Member**, International Foundation of Autonomous Agents and Multi-Agent Systems (IFAAMAS), 2008–14.
- Austin Business Journal **Tech Innovation Award**, November 2007.
- Best Paper Award, RoboCup Symposium, July 2007.
- IJCAI Computers and Thought Award, January 2007.
- Best Paper Award, Genetic and Evolutionary Computation Conference, GA Track, July 2006.
- Best Student Paper Award, RoboCup Symposium, June 2006.
- Elected Councilor, Association for the Advancement of AI (AAAI), July 2005–2008.
- Alfred P. Sloan Research Fellow, September 2004 2006.
- Office of Naval Research (ONR) Young Investigator, June 2004 May 2007.
- **IBM Faculty Award**, 2005, 2004, 2003.
- National Science Foundation **CAREER Award**, February 2003 January 2008.
- Best Paper Award, Autonomous Agents Conference, May 2001.
- AT&T Labs **Research Innovator**, 2000.
- NASA Graduate Student Research Program Fellowship, 1997 1999.
- Allen Newell Medal for Research Excellence, August 1997.
- NASA Certificate of Recognition for the creative development of a technical innovation entitled "DCAPS Iterative Repair Planning and Scheduling System," June 1997.
- Pennsylvania Space Grant Consortium fellowship, 1996.
- Hertz Foundation Research Fellowship Grant, 1995. (note: not the Hertz Graduate Fellowship)
- National Science Foundation honorable mention, 1993, 1994.
- Undergraduate Research Stipend Florida State University, June August 1992.
- State Farm Exceptional Student Fellowship, June 1992.
- The University of Chicago: Phi Beta Kappa, Sigma Xi, Dean's List every year, College Honor Scholarship: merit-based 4-yr, full-tuition scholarship, National Merit Scholarship, Maroon Key Society, Student Marshall, Scholar-Athlete Award: 4-yr varsity letterman with highest GPA.

TEACHING

• **Professor** at **The University of Texas at Austin**: September 2012 – present.

CS 394R Reinforcement Learning: Theory and Practice — online. Summer 2025. rating: 4.3/5.0 CS 309 AI Literacy: The Essentials of AI for Life and Society. Autumn 2024. rating: 4.7/5.0 CS 394R Reinforcement Learning: Theory and Practice — online. Summer 2024. rating: 4.2/5.0 CS 394R Reinforcement Learning: Theory and Practice. Spring 2024. Instructor rating: 4.7/5.0 CS 394R Reinforcement Learning: Theory and Practice — online. Autumn 2023. rating: 3.9/5.0

```
CS 109 The Essentials of AI for Life and Society. Autumn 2023.
                                                                           Instructor rating: 4.3/5.0
  CS 394R Reinforcement Learning: Theory and Practice — online. Autumn 2022.
                                                                                     rating: 3.7/5.0
  CS 395T Introduction to Ethical Artificial Intelligence and Robotics. Autumn 2022.
                                                                                    rating: 4.9/5.0
  CS 394R Reinforcement Learning: Theory and Practice. Spring 2022.
                                                                           Instructor rating: 4.4/5.0
  CS 394R Reinforcement Learning: Theory and Practice — online. Spring 2021.
                                                                                     rating: 4.3/5.0
  CS 343 Artificial Intelligence Spring 2021.
                                                                           Instructor rating: 4.2/5.0
  CS 394R Reinforcement Learning: Theory and Practice — online. Summer 2020.
                                                                                     rating: 4.1/5.0
  CS 394R Reinforcement Learning: Theory and Practice — online. Spring 2020.
                                                                                     rating: 4.8/5.0
  CS 394R Reinforcement Learning: Theory and Practice. Autumn 2019.
                                                                           Instructor rating: 4.5/5.0
  CS 393R Autonomous Robots. Autumn 2018.
                                                                           Instructor rating: 4.8/5.0
  CS 343H Artificial Intelligence: Honors. Autumn 2017.
                                                                           Instructor rating: 4.5/5.0
  CS 394R Reinforcement Learning: Theory and Practice. Autumn 2016.
                                                                           Instructor rating: 4.8/5.0
  CS 393R Autonomous Robots. Autumn 2015.
                                                                           Instructor rating: 4.7/5.0
  CS 343H Artificial Intelligence: Honors. Spring 2015.
                                                                           Instructor rating: 4.4/5.0
  CS 393R Autonomous Robots. Autumn 2013.
                                                                           Instructor rating: 4.9/5.0
  CS 311H Discrete Math for Computer Science: Honors. Autumn 2013.
                                                                           Instructor rating: 4.4/5.0
  CS 394R Reinforcement Learning: Theory and Practice. Spring 2013.
                                                                           Instructor rating: 5.0/5.0
  CS 313H Logic, Sets, and Functions: Honors. Autumn 2012.
                                                                           Instructor rating: 4.6/5.0
• Associate Professor at The University of Texas at Austin: September 2007 – August 2012.
  CS 343 Artificial Intelligence. Spring 2012.
                                                                           Instructor rating: 5.0/5.0
  CS 393R Autonomous Robots. Autumn 2011.
                                                                           Instructor rating: 4.9/5.0
  CS 394R Reinforcement Learning: Theory and Practice. Spring 2011.
                                                                           Instructor rating: 4.8/5.0
  CS 344M Autonomous Multiagent Systems. Autumn 2010.
                                                                           Instructor rating: 4.8/5.0
  CS 343 Artificial Intelligence. Spring 2010.
                                                                           Instructor rating: 4.7/5.0
  CS 393R Autonomous Robots. Autumn 2009.
                                                                           Instructor rating: 4.9/5.0
  CS 344M Autonomous Multiagent Systems. Spring 2008.
                                                                           Instructor rating: 4.9/5.0
  CS 394R Reinforcement Learning: Theory and Practice. Autumn 2007.
                                                                           Instructor rating: 4.9/5.0
• Assistant Professor at The University of Texas at Austin: June 2002 – August 2007.
  CS 378 Autonomous Vehicles — Driving in Traffic. Spring 2007.
                                                                           Instructor rating: 4.7/5.0
  CS 395T Agent-Based Electronic Commerce. Autumn 2006.
                                                                           Instructor rating: 4.9/5.0
  CS 378 Autonomous Multiagent Systems. Spring 2006.
                                                                           Instructor rating: 4.8/5.0
  CS 395T Autonomous Robots. Autumn 2005.
                                                                           Instructor rating: 4.7/5.0
                                                                           Instructor rating: 4.9/5.0
  CS 378 Autonomous Multiagent Systems. Spring 2005.
  CS 395T Reinforcement Learning: Theory and Practice. Autumn 2004.
                                                                           Instructor rating: 4.7/5.0
  CS 378 Autonomous Multiagent Systems. Spring 2004.
                                                                           Instructor rating: 4.8/5.0
  CS 395T Agent-Based Electronic Commerce. Autumn 2003.
                                                                           Instructor rating: 4.6/5.0
  CS 395T Multi-Robot Systems. Spring 2003.
                                                                           Instructor rating: 4.3/5.0
  CS 378 Autonomous Multiagent Systems. Autumn 2002.
                                                                           Instructor rating: 4.9/5.0
• Adjunct Professor at New York University: September 2001 – January 2002.
  Graduate class Autonomous Multiagent Systems. Autumn 2001.
                                                                           Instructor rating: 4.6/5.0
• Tutorial on learning motion control for mobile robot navigation at ICRA-22, May 2022.
```

- **Tutorial** on *multiagent learning: foundations and recent trends* at IJCAI-17, August 2017. • Tutorials on autonomous bidding agents at AAMAS-07 and AAAI-07, May – July 2007.
- Tutorials on robot soccer at AAAI-99, Agents-99, and IJCAI-99, May August 1999.
- Teaching Assistant, How to Think Like a Computer Scientist with Prof. Steven Rudich. Spring 1996.
- Teaching Assistant, Introduction to Artificial Intelligence with Prof. Jaime Carbonell. Spring 1995.
- College Mathematics Tutor at the University of Chicago. 1992-93.
- Private Violin Teacher in Buffalo, NY. Taught 40 students individually. August 1989-August 1991.

THESIS COMMITTEES

- **Doctoral Committee Supervisor:** (The University of Texas at Austin)
 - Jiaxun Cui, defended August 2025. Communication and Generalization in Multi-Agent Learning.
 - Yuqian Jiang, defended July 2025.

Towards Long-Term Deployment of General-Purpose Service Robots.

Yifeng Zhu, defended March 2025.

Efficient Sensorimotor Learning for Open-world Robot Manipulation.

Bo Liu, defended October 2024.

From Diversity to Adaptivity: Effective Multiobjective Learning and Continual Learning Neural Architectures.

- Haresh Karnan, defended February 2024.

Aligning Robot Navigation Behaviors with Human Intentions and Preferences.

- Ishan Durugkar, defended March 2023.

Estimation and Control of Visitation Distributions for Reinforcement Learning.

Faraz Torabi, defended April 2021.

Imitation Learning from Observation.

- Sanmit Narvekar, defended December 2020.

Curriculum Learning in Reinforcement Learning.

- Josiah Hanna, defended August 2019.

Data Efficient Reinforcement Learning with Off-policy and Simulated Data.

- Elad Liebman, defended February 2019.

Sequential Decision Making in Musical Intelligence.

- Patrick MacAlpine, defended July 2017.

Multilayered Skill Learning and Movement Coordination for Autonomous Robotic Agents.

UT Austin Computer Science Bert Kay Outstanding Dissertation Award.

- Katie Genter, defended June 2017.

Fly with Me: Algorithms and Methods for Influencing a Flock.

- Piyush Khandelwal, defended May 2017.

On-Demand Coordination of Multiple Service Robots.

Matthew Hausknecht, defended November 2016.

Cooperation and Communication in Multiagent Deep Reinforcement Learning.

- Daniel Urieli, defended November 2015.

Autonomous Trading in Modern Electricity Markets.

Samuel Barrett, defended October 2014.

Making Friends on the Fly: Advances in Ad Hoc Teamwork.

- Todd Hester, defended December 2012.

TEXPLORE: Temporal Difference RL for Robots and Time-Constrained Domains.

- W. Bradley Knox, defended August 2012.

Learning from Human-Generated Reward.

UT Austin Computer Science Bert Kay Outstanding Dissertation Award.

- Doran Chakraborty, defended August 2012.

Sample Efficient Multiagent Learning in the Presence of Markovian Agents.

- Juhyun Lee, defended November 2011.

Robust Color-based Vision for Mobile Robots.

- Shivaram Kalyanakrishnan, defended November 2011.

Learning Methods for Sequential Decision Making with Imperfect Representations.

David Pardoe, defended April 2010.

Adaptive Trading Agent Strategies Using Market Experience.

Nicholas K. Jong, defended December 2010.

Structured Exploration for Reinforcement Learning.

- Gregory Kuhlmann, defended August 2010.

Automated Domain Analysis for General Game Playing.

Kurt Dresner, defended October 2009.

Autonomous Intersection Management.

UT Austin Outstanding Dissertation Award.

- Matthew E. Taylor, defended June 2008.

Autonomous Inter-Task Transfer in Reinforcement Learning Domains.

- Daniel Stronger, defended June 2008.

Autonomous Sensor and Action Model Learning for Mobile Robots.

- Shimon Whiteson, defended April 2007.
 Adaptive Representations for Reinforcement Learning.
- Mohan Sridharan, defended April 2007.
 Robust Structure-Based Autonomous Color Learning on a Mobile Robot.
- **Doctoral Committee Member:** (The University of Texas at Austin)
 - Oguzhan Akcin, Electrical and Computer Eng. Supervisor: Sandeep Chinchali.
 Data Collection for Robust Machine Learning in Multi-Agent Systems,
 defended June 2025.
 - Harshit Sikchi, Computer Science. Supervisors: Amy Zhang and Scott Niekum.
 Reinforcement Learning Beyond Rewards: Decision-Making in the Language of Visitation Distributions,
 defended March 2025.
 - Caleb Chuck, Computer Science. Supervisors: Scott Niekum, Yuke Zhu, and Amy Zhang.
 Control-based Factorization through Causal Interactions and Hierarchical Reinforcement Learning,
 - defended October 2024.
 - Dian Chen, Computer Science. Supervisor: Philipp Kräehenbühl.
 Towards Richly Supervised Autonomous Driving Policies,
 defended July 2023.
 - Keya Rajesh Ghonasgi, Mechanical Engineering. Supervisor: Ashish Deshpande.
 Practice Makes Perfect: Leveraging Exoskeleton Interactions To Elucidate The Motor Learning Process,
 defended July 2023.
 - Santhosh Ramakrishnan, Computer Science. Supervisor: Kristen Grauman. Predictive Scene Representations for Embodied Visual Search, defended May 2023.
 - Ziyang Tang, Computer Science. Supervisor: Qiang Liu.
 Efficient And Safe Off-Policy Evaluation: From Point Estimation to Interval Estimation, defended March 2023.
 - Wonjoon Goo, Computer Science. Supervisor: Scott Niekum.
 Imitation Learning with Auxiliary, Suboptimal, and Task-Agnostic Data, defended November 2022.
 - Sadegh Rabiee, Computer Science. Supervisor: Joydeep Biswas.
 Introspective Perception for Mobile Robots,
 defended September 2022.
 - Prasoon Goyal, Computer Science. Supervisors: Raymond Mooney and Scott Niekum.
 Using Natural Language for Task Specification in Sequential Decision Making Problems, defended July 2022.
 - Mai Lee Chang, Electrical and Computer Eng. Supervisor: Andrea Thomaz.
 Optimizing for Task Performance and Fairness for Human-Robot Teamwork,
 defended July 2022.
 - Mauricio Tec, Statistics. Supervisors: James Scott and Corwin Zigler.
 Spatial Applications of Markov Random Fields and Neural Networks for Spatio-temporal Denoising, Causal Inference and Reinforcement Learning, defended June 2022.
 - Taylor Kessler, Electrical and Computer Eng. Supervisor: Andrea Thomaz.
 Learning Robot Policies from Imperfect Human Teacher,
 defended May 2022.
 - Carlin Liao, Civil Engineering. Supervisor: Steve Boyles.
 Modular autonomous intersection management simulation for stochastic and priority auction paradigms,
 defended November 2021.
 - Akanksha Saran, Computer Science. Supervisor: Scott Niekum.
 Leveraging Multimodal Human Cues to Enhance Robot Learning from Demonstration,

- defended November 2021.
- Yuchen Cui, Computer Science. Supervisor: Scott Niekum.
 Efficient Algorithms for Low-Effort Human Teaching of Robots, defended November 2021.
- Yan (Francis) Pei, Computer Science. Supervisor: Keshav Pingali.
 Introducing Principled Approximation and Online Control into Streaming Applications, defended July 2021.
- Abhinav Verma, Computer Science. Supervisor: Swarat Chaudhuri. Programmatic Reinforcement Learning, defended July 2021.
- Suda Bharadwaj, Aerospace Engineering. Supervisor: Ufuk Topcu. *Assured Decision-Making for Autonomous Systems*, defended July 2021.
- Ruohan Zhang, Computer Science. Supervisor: Dana Ballard.
 A Modular Attention Hypothesis for Modeling Visuomotor Behaviors, defended April 2021.
- Shih-Yun Lo, Mechanical Engineering. Supervisors: Andrea Thomaz and James Sulzer.
 Communicative Behavior Generation for Navigational Robots,
 defended March 2021.
- Lijia Liu, Computer Science. Supervisor: Dana Ballard.
 Cognitive Control of Motor Synergies,

defended March 2021.

- Venktesh Pandey, Civil Engineering. Supervisor: Steve Boyles
 Dynamic Pricing and Long-term Planning Models for Managed Lanes with Multiple Entrances and Exits,
 defended February 2021.
- Aishwarya Padmakumar, Computer Science. Supervisor: Raymond Mooney.
 Dialog as a Vehicle for Lifelong Learning of Grounded Language Understanding Systems, defended August 2020.
- Daniel Brown, Computer Science. Supervisor: Scott Niekum.
 Safe and Efficient Inverse Reinforcement Learning, defended July 2020.
- Venktesh Pandey, Civil Engineering. Supervisor: Stephen Boyles.
 - Dynamic Pricing and Long-term Planning Models for Managed Lanes with Multiple Entrances and Exits,

defended February 2020.

Jason Liang, Computer Science. Supervisor: Risto Miikkulainen.
 Evolutionary Neural Architecture Search for Deep Learning,

defended November 2018.

- Jesse Thomason, Computer Science. Supervisor: Raymond Mooney. *Continually Improving Grounded Natural Language Understanding*, defended April 2018.
- Donghyun Kim, Mechanical Engineering. Supervisor: Luis Sentis.
 Sensor-Based Robust Whole-Body Control of Highly Dynamic Legged Robots, defended November 2017.
- Kwan Suk Kim, Mechanical Engineering. Supervisor: Luis Sentis.
 Intelligent Collision Management in Dynamic Environments for Human-Centered Robots, defended August 2017.
- Karl Pichotta, Computer Science. Supervisor: Raymond Mooney.
 Advances in Statistical Script Learning,

defended July 2017.

- Subhashini Venugopalan, Computer Science. Supervisor: Raymond Mooney. Natural-Language Video Description with Deep Recurrent Neural Networks, defended June 2017.
- Wesley Tansey, Computer Science. Supervisor: James Scott.

Scalable Smoothing Algorithms for Massive Graph-Structured Data, defended May 2017.

Michael Levin, Civil Engineering. Supervisor: Stephen Boyles.
 Modeling and Optimizing Network Infrastructure for Autonomous Vehicles,

defended March 2017.

- Pei-Chi Huang, Computer Science. Supervisor: Al Mok.

Real-Time Robotic Tasks for Cyber-Physical Avatars,

defended February 2017.

- Ye Zhao, Mechanical Engineering. Supervisor: Luis Sentis.

A Planning And Control Framework Of Humanoid Systems: Robust, Optimal And Real-Time Performance,

defended August 2016.

- Christian Miller, Computer Science. Supervisor: Donald Fussell.

Derivative-Free Motion Optimization for Animated Characters, defended August 2016.

- Tarun Rambha, Civil Engineering. Supervisor: Stephen Boyles.

Dynamic Congestion Pricing in Within-Day and Day-to-Day Network Equilibrium Models, defended July 2016

- Chao-Yeh Chen, Computer Science. Supervisor: Kristen Grauman.

Learning Human Activities and Poses with Interconnected Data Sources, defended January 2016.

Ki Jung Yoon, Electrical and Computer Eng. Supervisors: Sriram Vishwanath and Ila Fiete.
 Unraveling the Dynamics and Structure of Grid Cells as a Spatial Map in the Brain, defended November 2015.

- Leif Johnson, Computer Science. Supervisor: Dana Ballard.

Redundancy Reduction in Motor Control,

defended September 2015.

- Donna Chen, Civil Engineering. Supervisor: Kara Kockleman.

Management of a Shared, Autonomous, Electric Vehicle Fleet: Vehicle Choice, Charging Infrastructure Planning, & Pricing Strategy,

defended July 2015.

- Gabriel Lopez-Mobilia, Psychology. Supervisor: Jacqueline Woolley.

Children's Psychological and Moral Attributions to a Humanoid Robot,

defended June 2015.

Jeremy Stober, Computer Science. Supervisors: Benjamin Kuipers and Risto Miikkulainen.
 Sensorimotor Embedding: A developmental approach to learning geometry,
 defended May 2015.

- Wenke Li, Neuroscience. Supervisor: Mike Mauk.

Timing in the Cerebellum: A Matter Of Network Inhibition,

defended January 2015.

- Nicholas Paine, Mechanical Engineering. Supervisor: Luis Sentis.

High-Performance Series Elastic Actuation,

defended August 2014.

Dan Fagnant, Civil Engineering. Supervisor: Kara Kockelman.

The Future of Fully Automated Vehicles: Opportunities for Vehicle- and Ride-Sharing, with Cost and Emission Savings,

defended June 2014.

Fangkai Yang, Computer Science. Supervisor: Vladimir Lifschitz.

Representing Actions in Logic-Based Languages,

defended March 2014.

- Jacob Schrum, Computer Science. Supervisor: Risto Miikkulainen.

Evolving Multimodal Behavior Through Modular Multiobjective Neuroevolution, defended March 2014.

Joohyun Kim, Computer Science, Supervisor: Raymond Mooney.
 Grounded Language Learning Models for Ambiguous Supervision,

- defended August 2013.
- Joseph Cooper, Computer Science, Supervisor: Dana Ballard.

Analysis and Synthesis of Bipedal Humanoid Movement: A Physical Simulation Approach, defended August 2013.

- Jaechul Kim, Computer Science, Supervisor: Kristen Grauman.

Region Detection and Matching for Object Recognition, defended July 2013.

Yaroslav Rosokha, Economics, Supervisor: Dale Stahl.

Capacity Of Multi-Antenna Ad Hoc Networks Via Stochastic Geometry, defended April 2013.

- Andrew Hunter, Electrical and Computer Engineering, Supervisor: Jeff Andrews.

Capacity Of Multi-Antenna Ad Hoc Networks Via Stochastic Geometry, defended November 2012.

- Rahul Iyer, Computer Science, Supervisor: Dana Ballard.

Efficient Muscle Representation for Human Walking,

defended September 2012.

- Bryan Silverthorn, Computer Science, Supervisor: Risto Miikkulainen.

A Probabilistic Architecture for Algorithm Portfolios,

defended April 2012.

- Dmitry Kit, Computer Science, Supervisor: Dana Ballard.

Change Detection Models for Mobile Cameras,

defended April 2012.

- David Chen, Computer Science, Supervisor: Raymond Mooney.

Learning Language from Ambiguous Perceptual Context,

defended January 2012.

Chinmayi Krishnappa, Computer Science, Supervisor: Greg Plaxton.

Unit-Demand Auctions: Bridging Theory and Practice,

defended December 2011.

- Changhai Xu, Computer Science, Supervisors: Benjamin Kuipers and Kristen Grauman.

Steps Towards the Object Semantic Hierarchy,

defended August 2011.

- Shilpa Gulati, Mechanical Engineering, Supervisors: Raul Longoria and Benjamin Kuipers.

A Framework for Characterization and Planning of Safe, Comfortable, and Customizable Motion of Assistive Mobile Robots,

defended June 2011.

- David Han, Electrical and Computer Engineering, Supervisor: K. Suzanne Barber.

Action Selection and Coordination of Autonomous Agents for UAV Surveillance, defended December 2010.

Vinod Valsalam, Computer Science, Supervisor: Risto Miikkulainen.

Utilizing Symmetry In Evolutionary Design,

defended August 2010.

- Jonathan Mugan, Computer Science, Supervisor: Benjamin Kuipers.

Autonomous Qualitative Learning of Distinctions and Actions in a Developing Agent, defended August 2010.

- Yuliya Lierler, Computer Science, Supervisor: Vladimir Lifschitz.

SAT-Based Answer Set Programming,

defended April 2010.

Yiu Fai Sit, Computer Sciences, Supervisor: Risto Miikkulainen.

A Population Gain Control Model of Spatiotemporal Responses in the Visual Cortex, defended August 2009.

- Nate Kohl, Computer Science, Supervisor: Risto Miikkulainen.

Learning in Fractured Problems with Constructive Neural Network Algorithms, defended August 2009.

- Lilyana Mihalkova, Computer Science, Supervisor: Ray Mooney.

Learning with Markov Logic Networks: Transfer Learning, Structure Learning, and an Appli-

- cation to Web Query Disambiguation, defended July 2009.
- Aniket Murarka, Computer Sciences, Supervisor: Ben Kuipers.
 Building Safety Maps using Vision for Safe Local Mobile Robot Navigation, defended August 2009.
- Michael Bond, Computer Sciences, Supervisor: Kathryn McKinley.
 Diagnosing And Tolerating Bugs In Deploymed Systems,
 defended September 2008.
- Patrick Beeson, Computer Sciences. Supervisor: Benjamin Kuipers.
 Creating And Utilizing Hybrid Representations Of Spatial Knowledge Using Mobile Robots, defended August 2008.
- Selim Erdogan, Computer Sciences. Supervisor: Vladimir Lifschitz.
 A Library of General-Purpose Action Descriptions, defended July 2008.
- Tal Tversky, Computer Sciences. Supervisor: Risto Miikkulainen, Bill Geisler.
 Motion Perception and Scene Statistics of Motion,
 defended April 2008.
- Nedialko Dimitrov, Computer Sciences. Supervisor: Greg Plaxton.
 Coping with Dynamic Membership, Selfishness, and Incomplete Information: Applications of Probabilistic Analysis and Game Theory, defended April 2008.
- Youngin Shin, Computer Sciences, Supervisor: Don Fussell.
 Parametric Kernels for Structured Data Analysis,
 defended December 2007.
- Karen Fullam, Electrical and Computer Engineering. Supervisor: K. Suzanne Barber.
 Adaptive Trust Modeling in Multi-Agent Systems: Utilizing Experience and Reputation, defended November 2007.
- Rohit Kate, Computer Sciences. Supervisor: Raymond Mooney.
 Learning For Semantic Parsing With Kernels Under Various Forms Of Supervision, defended August 2007.
- Jefferson Provost, Computer Sciences. Supervisors: Benjamin Kuipers, Risto Miikkulainen. Reinforcement Learning in High-Diameter Continuous Environments, defended August 2007.
- Joseph Modayil, Computer Sciences. Supervisor: Ben Kuipers.
 Robot Developmental Learning of an Object Ontology Grounded in Sensorimotor Experience, defended June 2007.
- Wallace Earl Depue, Jr. (Music), Supervisor: Andrew Dell'Antonio.
 Central Park Reel for Violin and Piano,
 defended November 2006.
- Bobby Bryant, Computer Sciences, Supervisor: Risto Miikkulainen.
 Evolving Visibly Intelligent Behavior For Embedded Game Agents,
 defended July 2006.
- Mikhail Bilenko, Computer Sciences, Supervisor: Raymond Mooney. Learnable Similarity Functions and Their Applications, defended July 2006.
- Prem Melville, Computer Sciences, Supervisor: Raymond Mooney *Creating Diverse Ensemble Classifiers to Reduce Supervision*, defended November 2005.
- Joohyung Lee, Computer Sciences, Supervisor: Vladimir Lifschitz.
 Automated Reasoning about Actions,
 defended May 2005.
- Brett Mitchell, Music, Supervisors: Byron Almén, Kevin Noe.
 Mahler and the Art of Self-borrowing,
 defended May 2005.
- Joon Woo Kim, Electrical and Computer Engineering, Supervisor: K. Suzanne Barber.

Trusting Information and Sources in Open Multi-Agent Systems, defended November 2003.

• **Doctoral Committee Member:** (External)

Andrew Levy, Computer Science, Brown University.

Supervisor: George Konidaris.

Unsupervised Skill Discovery with Empowerment.

August 2025.

- Tonghan Wang, Computer Science, Harvard University.

Supervisors: David Parkes and Milind Tambe.

Advancing Deep Learning for Multiagent AI: Mechanisms, Dynamics, and Organizations. June 2025.

- Rex Chen, Software & Societal Systems, Carnegie Mellon University.

Supervisors: Fei Fang and Norman Sadeh.

Rethinking the Design of Human-Centric AI Systems for Deployment in Transportation. June 2025.

- Ryan Sullivan, Computer Science, University of Maryland.

Supervisor: John Dickerson.

Applying Policy Gradient Methods to Open-Ended Domains.

April 2025.

- Lin Guan, Computer Science, Arizona State University.

Supervisor: Subbarao Kambhampati.

Taming the Sample Complexity in Agentifying AI Systems by the Exploitation of Explicit Human Knowledge.

October 2024.

- Serena Booth, Computer Science, Massachusetts Institute of Technology.

Supervisor: Julie Shah.

Building Blocks for Human-AI Alignment: Specify, Inspect, Model, and Revise.

October 2023.

Pascal Klink, Computer Science, TU Darmstadt.

Supervisor: Jan Peters.

Reinforcement Learning Curricula as Interpolations between Task Distributions.

November 2023.

- Rohan Paleja, Computer Science, Georgia Institute of Technology.

Supervisor: Matthew Gombolay.

Interpretable Artificial Intelligence for Personalized Human-Robot Collaboration.

August 2023.

Rémy Portelas, Computer Science, University of Bordeaux.

Supervisors: Pierre-Yves Oudeyer and Katja Hofmann.

Automatic Curriculum Learning for Developmental Machine Learners.

February 2022.

Sina Ghiassian, Computer Science, University of Alberta.

Supervisors: Richard Sutton and Adam White.

Online Off-policy Prediction.

January 2022.

- Max Korein, Computer Science, Carnegie Mellon University.

Supervisor: Manuela Veloso.

Planning to Optimize and Learn Reward in Navigation Tasks in Structured Environments with Time Constraints.

July 2021.

Mohammad Rostami, Electrical and Systems Engineering, University of Pennsylvania.

Supervisors: Eric R. Eaton and Daniel D. Lee.

Learning Transferable Knowledge through Embedding Spaces.

July 2019.

- Jakob Foerster, Computer Science, University of Oxford.

Supervisor: Shimon Whiteson.

Deep Multi-Agent Reinforcement Learning.

January 2019.

Dhanvin Mehta, Computer Science and Engineering, University of Michigan.

Supervisor: Edwin Olson.

Multi-Policy Decision Making for Reliable Navigation in Dynamic Uncertain Environments.

November 2018.

Kaushik Subramanian, Computer Science, Georgia Institute of Technology.

Supervisors: Charles Isbell and Andrea Thomaz.

Policy-Based Exploration for Efficient Reinforcement Learning.

May 2017.

- Timothy Wiley, Computer Science and Engineering, University of New South Wales.

Supervisors: Claude Sammut and Bernhard Hengst.

A Planning and Learning Hierarchy for the Online Acquisition of Robot Behaviours.

June 2017.

- Junqing Wei, Robotics Institute, Carnegie Mellon University.

Supervisor: John Dolan.

Autonomous Vehicle Social Behavior.

May 2017.

- Nolan Bard, Computing Science, University of Alberta.

Supervisor: Michael Bowling.

Online Agent Modelling in Human-Scale Problems.

March 2016.

- Jason Pazis, Computer Science, Duke University.

Supervisor: Ron Parr.

PAC-optimal, Non-parametric Algorithms and Bounds for Exploration in Concurrent MDPs with Delayed Updates.

October 2015.

- Guni Sharon, Computer Science, Ben Gurion University.

Supervisor: Ariel Felner.

Multi-Agent Path-Finding and Agent Centered Search.

July 2015.

- Mariano Schein, Computer Science, Tel Aviv University.

Supervisor: Yishay Mansour.

Machine Learning Algorithms and Robustness.

March 2015.

Sayan Sen, Computer Science, Vanderbilt University.

Supervisor: Julie Adams.

An Intelligent and Unified Framework for Multiple Robot and Human Coalition Formation.

January 2015.

- Reshef Meir, Computer Science, Hebrew University.

Supervisor: Jeffrey Rosenschein.

Mechanisms for Stability and Welfare: Increasing Cooperation among Self-interested Agents. August 2013.

- Xiang Li, Computer Science, Texas Tech.

Supervisor: Mohan Sridharan.

Autonomous Learning of Object Models on Mobile Robots Using Visual Cues.

July 2013.

- Somchaya Liemhetcharat, Computer Science, Carnegie Mellon University.

Supervisor: Manuela Veloso.

Representation, Planning, and Learning of Dynamic Ad Hoc Robot Teams.

July 2013.

- Haitham Bou Ammar, Artificial Intelligence, Maastricht University.

Supervisor: Karl Tuyls.

Automated Transfer for Reinforcement Learning.

June 2013.

- Timothy Mann, Computer Science, Texas A&M.

Supervisor: Yoonsuck Choe.

Scaling up RL without Sacrificing Optimality by Constraining Exploration.

October 2012.

Jason Kulk, Electrical Engineering and Computer Science, University of Newcastle, Australia.
 Supervisor: James Welsh.

Improved Humanoid Robot Movement through Impact Perception and Walk Optimisation. August 2012.

- Matthew Robards, Computer Science, Australian National University, Australia.

Supervisor: Peter Sunehag.

Online Learning for Reinforcement Learning with Function Approximation.

January 2012.

- Matteo Leonetti, Ingegneria Informatica, University of Rome.

Supervisor: Luca Iocchi.

Robot Teams for Multi-Objective Tasks.

November 2010.

- Alessandro Lazaric, Elettronica e Informazione, Politecnico Di Milano.

Supervisor: Andrea Bonarini.

Knowledge Transfer in Reinforcement Learning.

January 2008.

Min-Sub Kim, Computer Science and Engineering, University of New South Wales, Australia.

Supervisor: Will Uther.

Reinforcement Learning by Incremental Patching.

January 2008.

- Vittorio Ziparo, Ingegneria Informatica, University of Rome.

Supervisor: Daniele Nardi.

Robot Teams for Multi-Objective Tasks.

November 2007.

Christian Quintero, Electronics, Computer Science and Automatic Control, University of Girona.
 Supervisor: Josep Ll. de la Rosa.

Introspection on Control-grounded Capabilities. An Agent-inspired Approach for Control . October 2007.

- Robert Abbott, Computer Science, U. of New Mexico.

Supervisor: Stephanie Forrest.

Automated Tactics Modeling: Techniques and Applications.

April 2007.

- Jelle Kok, Computer Science, University of Amsterdam, Netherlands.

Supervisor: Nikos Vlassis.

Coordination and Learning in Cooperative Multiagent Systems.

November 2006.

Michael Quinlan, Computer Science and Software Engineering, U. of Newcastle, Australia.

Supervisor: Stephan Chalup.

Machine Learning on AIBO Robots.

June 2006.

- Jeff Riley, RMIT University, Australia.

Supervisor: Victor Ciesielski.

Evolving Fuzzy Rules for Goal-Scoring Behaviour in a Robot Soccer Environment.

February 2006.

• Masters Thesis Supervisor: (The University of Texas at Austin)

- Ajith Kemisetti, Computer Science, Spring 2025.

Stochastic SIPP: A Novel Approach for Motion Planning in Stochastic Dynamic Environments.

- Faraz Heravi, Computer Science, Spring 2025.

Curriculum Learning in Reinforcement Learning for Multi-Step Manipulation Tasks.

- Abayomi Adekanmbi, Computer Science, Spring 2024.

Deep Reinforcement Learning in RoboCup Keepaway.

- Sai Kiran Narayanaswami, Computer Science, Summer 2023.

Decision-Making Problems in Computationally Constrained Robot Perception.

William Macke, Computer Science, Spring 2023.

Optimizing and Planning with Queries in Communication in Ad Hoc Teamwork.

- Bharath Masetty, Mechanical Engineering, Summery 2021.

Modeling Human Motor Learning Traits using Reinforcement Learning.

- Qiping Zhang, Computer Science, Summer 2021.

Iteractive Learning from Ipmlicit Human Feedback: The EMPATHIC Framework.

- Brahma Pavse, Computer Science, Spring 2020.

Reducing Sampling Error in Batch Temporal Difference Learning.

- Prabhat Nagarajan, Computer Science, Summer 2018.

Nondeterminism as a Reproducibility Challenge for Deep Reinforcement Learning.

- Rolando Fernandez Jr., Spring 2017.

Light-Based Nonverbal Signaling with Passive Demonstrations for Mobile Service Robots.

- Priyanka Kante, Computer Science, Spring 2017.

Learning Attributes of Real-world Objects by Clustering Multimodal Sensory Data.

- Yuchen He, Computer Science, Autumn 2013.

Localization using Natural Landmarks Off-Field for Robot Soccer.

- Alon Farchy, Computer Science, Spring 2012.

Learning in Simulation for Real Robots.

- Neda Shahidi, ECE, Summer 2010.

A Response Delayed Policy for Autonomous Intersection Management.

- Gurushyam Hariharan, ECE, Spring 2004.

News Mining Agent for Automated Stock Trading.

- Harish Subramanian, ECE, Summer 2004.

Evolutionary Algorithms in Optimization of Technical Rules for Automated Stock Trading.

• Masters Thesis Reader: (The University of Texas at Austin)

- Xinyi Liu, Operations Research and Industrial Engineering, Spring 2024.

Supervisor: Qianwen Wang (University of Minnesota).

Adaptive Explainable AI: Designing User-Centric Explanation Systems for Enhanced Interaction.

- Josh Kelle, Computer Science, Spring 2017.

Supervisor: Kristen Grauman.

Frugal Forests: Learning a Dynamic and Cost Sensitive Feature Extraction Policy for Anytime Activity Classification.

- Shun Zhang, Computer Science, Summer 2015.

Supervisor: Dana Ballard.

Parameterized Modular Inverse Reinforcement Learning.

- Jason Liang, Computer Science, Spring 2015.

Supervisor: Risto Miikkulainen.

Evolutionary Bilevel Optimization for Complex Control Problems and Blackbox Optimization.

Anand Subramoney, Computer Science, Summer 2012.

Supervisor: Risto Miikkulainen.

Evaluating ESP in the robot soccer keepaway domain.

- Aravind Gowrisankar, Computer Sciences, Autumn 2008.

Supervisor: Risto Miikkulainen.

Evoloving Controllers for Simulated Car Racing Using Neuroevolution.

- Travis Mercker, Aerospace Engineering, Spring 2008.

Supervisor: Maruthi Akella.

Self-Organization and Navigation Algorithms for Deployable Decentralized Sensor Networks.

- Karen Fullam, ECE, Autumn 2003.

Supervisor: K. Suzanne Barber.

An Expressive Belief Revision Framework Based on Information Valuation.

• Masters Thesis Reader: (External)

- Leonid Trainer, Computer Science, Hebrew University, December 2014.

Supervisor: Jeff Rosenschein.

Collaboration in Ad Hoc Settings: Novel Approaches to Implementation.

• Undergraduate Honors Thesis Supervisor (The University of Texas at Austin)

- William Yue, Computer Science, Autumn 2024.

Towards General Purpose Robots at Scale: Lifelong Learning and Learning to Use Memory.

- Stephane Hatgis-Kessell, Computer Science, Autumn 2023.

A New Model of Human Preferences for Learning Reward Functions.

UT Austin CS Outstanding Undergraduate Thesis Award.

- Brahma Pavse, Computer Science, Spring 2019.

RIDM: Reinforced Inverse Dynamics Modeling for Learning from a Single Observed Demonstration.

- Harsh Goyal, Computer Science, Spring 2019.

Holistic Action Transform.

- Sean Geiger, Computer Science, Spring 2019.

Sample-efficient Imitation from Observation on a Robotic Arm.

- Avilash Rath, Computer Science, Spring 2019.

Learning Social Behavior from Human Feedback in Ad Hoc Teamwork.

- John Fang, Computer Science, Spring 2019.

Black-Box Optimization of Parameterized Link-Dependent Road Tolling.

- Virin Tamprateep, Computer Science, Spring 2017.

Of Mice and Mazes: Simulating Mice Behavior with Reinforcement Learning.

- Yuqian Jiang, Computer Science, Autumn 2016.

Efficient Symbolic Task Planning for Multiple Mobile Robots.

- Patricio Lankenau, Computer Science, Summer 2016.

Virtour: Telepresence System for Remotely Operated Building Tours.

- Mike Depinet, Computer Science, Spring 2014.

Keyframe Sampling, Optimization, and Behavior Integration: A New Longest Kick in the RoboCup 3D Simulation League.

- Christopher Donahue, Computer Science, Autumn 2013.

Applications of genetic programming to digital audio synthesis.

- Dustin Carlino, Computer Science, Autumn 2013.

Approximately Orchestrated Routing and Transportation Analyzer: City-scale autonomous traffic simulation.

UT Austin CS Outstanding Undergraduate Thesis Award.

- Adrian Lopez-Mobilia, Computer Science, Spring 2012.

Inverse Kinematics Kicking in the Humanoid RoboCup Simulation League.

- Jason Weng, Computer Science, Spring 2012.

Identifying the Content and Location of Objects in a Roadside Image through Computer Vision.

- Nick Collins, Computer Science, Spring 2012.

Transformation of robot model to facilitate optimization of locomotion.

- Chau Nguyen, Computer Science, Autumn 2009.

Constructing Drivability Maps From 3D Laser Range Data for Autonomous Vehicles.

- Adam Setapen, Computer Science, Spring 2009.

Exploiting Human Motor Skills for Training Bipedal Robots.

- Tarun Nimmagadda, Computer Sciences, Spring 2008.

Building an Autonomous Ground Traffic System.

Ryan Madigan, Computer Sciences, Spring 2007.

Control Module for an Autonomous Mobile Robot Operating in an Urban Environment.

- Jan Ulrich, Computer Sciences, Spring 2006.

An Analysis of the 2005 TAC SCM Finals.

- Irvin Hwang, Computer Sciences, Spring 2005.

Discovering Conditions for Intermediate Reinforcement with Causal Models.

- Ellie Lin, Computer Sciences, Autumn 2003.

Creation of a Fine Controlled Action for a Robot.

• Undergraduate Thesis Reader: (The University of Texas at Austin)

- Anant Bhatia, Computer Science, Spring 2025.

Supervisor: Amy Zhang.

Is Bootstrapping Necessary for Offline Goal Conditioned Quasimetric Learning?

- Emily Sturman, Computer Science, Spring 2025.

Supervisor: Luis Sentis.

Sharp and Agile Locomotion Turns With TOVI-MPC.

Anthony Wang, Computer Science, Spring 2025.

Supervisor: W. Bradley Knox.

Talk, Watch, Steer: Leveraging Thing-Aloud Human Interventions and Vision-Language Models to Refine Offline Driving Policies.

- Sonesh Patel, Computer Science, Spring 2024.

Supervisor: Amy Zhang.

Crossing Paths: Incentivising Exploration using Distributional Reinforcement Learning.

- Rahul Menon, Computer Science, Spring 2023.

Supervisor: Joydeep Biswas.

Terrain-Adaptive Global Planning from Local Demonstrations.

- Thomas Nathaniel Plaxton, Computer Science, Spring 2023.

Supervisor: Joydeep Biswas.

Estimating Kinodynamic Uncertainty Using Learned Gaussian Noise Models.

- Elvin Yang, Computer Science, Spring 2023.

Supervisor: Joydeep Biswas.

Wait, That Feels Familiar: Learning to Extrapolate Human Preferences for Preference-Aligned Path Planning.

- Evonne Ng, Computer Science, Spring 2019.

Supervisor: Kristen Grauman.

You2Me: Inferring Body Pose in Egocentric Video via First and Second Person Interactions.

- Victoria Zhou, Computer Science, Spring 2019.

Supervisor: Justin Hart.

Exploration of Neural Networks for Stereo Vision.

- Michael Langford, Computational Engineering, Spring 2019.

Supervisor: Bruce Pennycock.

A Comparison of Recurrent Neural Network Effectiveness in Generating Music in the Style of 18th-Century Counterpoint.

- Jackson Haenchen, Plan II, Spring 2017.

Supervisor: David Prindle.

Artificial Intelligence: Predictions for the Future through a Political Lens.

- Rodolfo Corona, Computer Science, Autumn 2016.

Supervisor: Raymond Mooney.

An Analysis of Using Semantic Parsing for Speech Recognition.

- Michael Levin, Computer Science, Autumn 2012.

Supervisor: Stephen Boyles.

A Comparative Analysis of Heuristics for the Improved Convergence of Dynamic Traffic Assignment Models.

- David Robson, Computer Science, Spring 2010.

Supervisor: Risto Miikkulainen.

Hierarchical Neural Networks for Behavior-Based Decision Making.

Laurel Issen, Computer Sciences, Spring 2006.

Supervisor: Bill Geisler.

Using Edge Statistics for Object Recognition.

- Clare Richardson, Computer Sciences, Autumn 2005.

Supervisor: Benjamin Kuipers.

Rapid, High Precision Control in Tightly Constrained Environments.

OTHER ADVISING

- Postdoctoral Fellows: Nathan Tsoi (2005–), Junhong Xu (2005–), Shuijing Liu (2005–), Claas Voelcker (2005–), Arrasy Rahman (2022–), Chen Tang (2023–25), Alexander Levine (2023–25), Yoonchang Song (2021-25), Rohan Chandra (2022–24), Shahaf Shperberg (2021–22), Yulin Zhang (2021–22), Xuesu Xiao (2019–22), Reuth Mirsky (2019–22), Harel Yedidsion (2017–21), Shani Alkoby (2017–19), Patrick MacAlpine (2017–18), Guni Sharon (2015–18), Justin Hart (2016–17), Stefano Albrecht (2016–2017), Jivko Sinapov (2014–2017), Michael Albert (2015–2016), Shiqi Zhang (2014–2016), Matteo Leonetti (2013–2015), Todd Hester (2013), Noa Agmon (2010–2012), Tsz-Chiu Au (2008–2012), Michael Quinlan (2007–2011), Tobias Jung (2008–2010), Patrick Beeson (2008–2009), Ian Fasel (2007–2008), Yaxin Liu (2004–2007), Bikramjit Banerjee (2006).
- Other Current UT Austin Ph.D. students: Siddhant Aggarwal, Jiaheng Hu, Yu-Sian Jiang, Michael Munje, Jin Soo Park, Caroline Wang, Zhihan Wang, Zizhao Wang, Zifan Xu.
- AT&T Labs Research summer intern: Paul Reitsma (2001).
- CMU undergraduate research (informal): Patrick Riley (1998–1999), Michael Bowling (1996).

PROFESSIONAL ACTIVITIES

- Major event coordination:
 - **General co-chair**, Conference on Robot Learning (CoRL), 2026.
 - General co-chair, International Conference on Distributed Artificial Intelligence (DAI), 2023.
 - Conference chair, International Joint Conference on Artificial Intelligence (IJCAI), 2023.
 - Program co-chair, Multi-disciplinary Conference on Reinforcement Learning and Decision Making (RLDM), 2022.
 - General chair, RoboCup, 2021.
 - **Journal Track chair**, International Joint Conference on Artificial Intelligence (IJCAI), 2017.
 - **Journal Track chair**, Autonomous Agents and Multi-Agent Systems (AAMAS), 2016.
 - Co-chair, Machine Learning Summer School (MLSS), 2015.
 - **Program co-chair**, AAAI Conference on Artificial Intelligence (AAAI), 2014.
 - Tutorial chair, International Conference on Machine Learning (ICML), 2013.
 - Video Track Chair, International Joint Conference on Artificial Intelligence (IJCAI), 2011.
 - General co-chair, Autonomous Agents and Multi-Agent Systems (AAMAS), 2011.
 - **Tutorial co-chair**, Conference on Artificial Intelligence (AAAI), 2008.
 - **Program co-chair**, Autonomous Agents and Multi-Agent Systems (AAMAS), 2006.
 - Workshop co-chair, Conference on Artificial Intelligence (AAAI), 2005.
 - Chair, RoboCup US Open simulation league committee, 2005, 2004.
 - Tutorial chair, International Joint Conference on Artificial Intelligence (IJCAI), 2003.
 - Entry coordinator, Trading Agent Competition, 2001.
 - **Associate chair** in charge of simulation events for RoboCup, 2001.
 - Co-chair, RoboCup simulator competition organizing committee, 1997–1999.

• Workshop/Symposium coordination:

- Co-chair, RoboCup Symposium, 2012.
- **Program Co-chair**, Humanoids Workshop on *Humanoid Soccer Robots*, 2010, 2011.
- **Co-chair**, AAAI workshop on *Multiagent Learning*, 2005.
- **Co-chair**, ICML Workshop on *Physiological Data Mining Contest*, 2004.
- Chair, Information Science and Technology (ISAT) study on
 - Distributed Cognitive Systems Focused on Team/Multiagent Learning, 2004, 2005.
- **Co-chair**, IJCAI workshop on *Trading Agent Design and Analysis*, 2003.
- **Co-chair**, AAAI Fall Symposium on *Personalized Agents*, 2002.
- Co-chair, AAAI Spring Symposium on Collaborative Learning Agents, 2002.
- **Co-chair**, Agents Workshop on *Learning Agents*, 2001.
- Co-chair, RoboCup Workshop, 2000.
- **Co-chair**, Agents Workshop on *Learning Agents*, 2000.
- Editor-in-chief, Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS), 2010–2016.
- Editor, Synthesis Lectures on Artificial Intelligence and Machine Learning, 2012–present.
- Associate editor:

- International Conference on Robotics and Automation (ICRA), 2011.
- Artificial Intelligence Journal (AIJ), 2007–2014.
- J. of Autonomous Agents and Multi-Agent Systems (JAAMAS), 2003–2009, 2016–present.
- ACM Transactions on Internet Technology (TOIT), 2003–2005.
- International Journal of Image and Graphics (IJIG), 2002–2006.

• Assistant editor:

- ACM SIGecom Exchanges, 2004–2005.

• Senior Editorial board:

- IEEE Journal on *Robotics and Automation Practice* (RA-P), 2024–present.

• Editorial board:

- ACM Journal on *Responsible Computing*, 2022–present.
- Machine Learning Journal (MLJ), 2003-present.
- Springer Verlag's Encyclopedia of Machine Learning, 2005–2010.
- Journal of Artificial Intelligence Research (JAIR), 2002–2005.

• Guest editor:

- AIJ special issue on Autonomous Agents Modelling Other Agents, 2020.
- IEEE Intelligent Systems special issue on *Multi-Robot Systems*, 2017.
- JAAMAS special issue on Multiagent Interaction without Prior Coordination, 2016.
- ACM SIGecom Exchanges special issue on Trading Agent Design and Analysis, 2004.
- IEEE Intelligent Systems special issue on Agents and Markets, 2003.

• Organizing committee member:

- AAMAS workshop on Rebellion and disobedience in AI (RaD-AI), 2022, 2023.
- ICRA workshop on *Machine Learning for Motion Planning*, 2021.
- AAAI Spring Symposium on Machine Learning for Mobile Robot Navigation in the Wild, 2021.
- AAAI Fall Symposium on Reasoning and Learning in Real-World Systems for Long-Term Autonomy (LTA), 2018.
- IJCAI workshop on Explainable Artificial Intelligence (XAI), 2017.
- AAAI Spring Syposium on Challenges and Opportunities in Multiagent Learning for the Real World, 2016.
- AAAI Spring Syposium on Intelligent Systems for Supporting Distributed Human Teamwork, 2016.
- NIPS workshop on Learning, Inference and Control of Multi-Agent Systems, 2015.
- AAAI Spring Symposium on Applied Computational Game Theory, 2015.
- AAMAS workshop on Agent Technologies for Energy Systems (ATES), 2013.
- IJCAI Workshop on General Game Playing, 2009.
- 2nd Reinforcement Learning Competition, 2008.
- ICAPS Workshop on AI Planning and Learning, 2007.
- NIPS Workshop on The Inaugural Reinforcement Learning Competition, 2006.
- AAAI Fall Symposium on Real Life Reinforcement Learning, 2004.
- AAMAS Workshop on Learning and Evolution in Agent Based Systems, 2004.

• Book reviewer:

- Elsevier, 2015, 2016.
- Synthesis Lecture Series, 2011.
- Cambridge University Press, 2010.
- John Wiley & Sons, 2006, 2007.
- Morgan Kaufmann, 2001.

• Journal article reviewer:

- Adaptive Behavior, 2006.
- Advanced Robotics Journal, 1999.
- ACM Transactions on Intelligent Systems and Technology (TIST), 2010
- AI Communication (AICOM), 2005.
- − AI Magazine, 2010.
- Artificial Intelligence (AIJ), 2002, 2005, 2006, 2013, 2014, 2016, 2018, 2019.
- Autonomous Agents and Multi-Agent Systems (JAAMAS), 2000, 2002–2007.
- Autonomous Robots, 1999.

- *Communications of the ACM* (CACM), 2009, 2010, 2022.
- Computational Intelligence, 2003.
- Data Mining and Knowledge Discovery (DMKD) 2007.
- Decision Support Systems (DSS), 2003, 2004, 2006, 2007.
- Electronic Commerce (EC)
- Electronic Communication of the EASST (ECEASST)
- Electronic Markets (EM), 2002.
- Engineering Applications of Artificial Intelligence (EAAI), 2011
- IEEE Internet Computing, 2006.
- IEEE Robotics and Automation Letters (RA-L), 2020.
- IEEE Transactions on Games, 2020.
- IEEE Transactions on Intelligent Transportation Systems, 2010.
- IEEE Transactions on Knowledge and Data Engineering (IEEE TKDE), 1999, 2002.
- IEEE Transactions on Robotics (IEEE TRO), 2004–2007.
- IEEE Transactions on Robotics and Automation (IEEE TRA), 2001, 2002.
- International Journal of Robotics Research (IJRR), 2011–2013.
- International Journal of Social Robotics (IJSR), 2011.
- INFORMS Journal on Computing, 2006.
- International Journal of Advanced Robotic Systems (IJARS), 2012.
- Journal of Artificial Intelligence Research (JAIR), 2000–2005, 2012, 2013.
- Journal of Behavioral Robotics, 2009.
- Journal of Intelligent and Robotic Systems (JINT) 2013.
- Journal of Intelligent Traffic Systems (JITS), 2010.
- Journal of Machine Learning Research (JMLR), 2003, 2005, 2009, 2011.
- Knowledge and Information Systems (KAIS), 2000, 2002.
- Knowledge Engineering Review, 2003.
- *Machine Learning Journal* (MLJ), 2003, 2005–2013, 2015.
- Nature, 2018, 2024.
- Neural Networks (NN), 2007, 2008.
- Proceedings of the National Academy of Sciences (PNAS), 2024.
- Robotics and Autonomous Systems (RAS), 2003, 2007.
- Science, 2018, 2022.
- Systems, Man and Cybernetics (SMC), 2005.
- Texas Undergraduate Research Journal, 2013.
- Transportation Research Board (TRB), 2014.
- Transportation Research Part C (TRC), 2010.
- Wiley Interdisciplinary Reviews: Cognitive Science, 2009.

• Conference advisory/steering committee member:

- Edge of Now (EON), 2025–present.
- IEEE International Conference on Simulation, Modeling, and Programming for Autonomous Robots (SIMPAR), 2024–present.
- Reinforcement Learning Conference (RLC), 2023–present.
- Conference on Lifelong Learning Agents (CoLLA), 2022–present.
- Conference on Artificial Intelligence (AAAI), 2015.
- International Joint Conference on Artificial Intelligence (IJCAI), 2009.
- IPTO Cognitive Systems Conference, 2005–2006.

• Conference senior area chair:

- International Conference on Machine Learning (ICML), 2021–2025.
- Neural Information Processing Systems (NeurIPS), 2002, 2003, 2017, 2018, 2020–2025.
- International Joint Conference on Artificial Intelligence (IJCAI), 2021.

• Conference area chair:

- AAAI Conference on Artificial Intelligence (AAAI), 2018, 2021, 2023, 2026.
- International Joint Conference on Artificial Intelligence (IJCAI), 2009, 2013, 2016, 2018, 2019, 2022, 2025.
- Reinforcement Learning Conference (RLC), 2024, 2025.

- Conference on Robot Learning (CoRL), 2024, 2025.
- Autonomous Agents and Multiagent Systems (AAMAS), 2019.
- International Conference on Machine Learning (ICML), 2003, 2012, 2015, 2016.
- European Conference on Machine Learning (ECML), 2005.

• Conference senior program committee member:

- International Conference on Intelligent Robots and Systems (IROS), 2020.
- Autonomous Agents and Multiagent Systems (AAMAS), 2003, 2004, 2007, 2016.
- International Joint Conference on Artificial Intelligence (IJCAI), 2007, 2015.
- International Conference on Machine Learning (ICML), 2006.
- AAAI Conference on Artificial Intelligence, 2002, 2004.

• Conference program committee member:

- Multi-disciplinary Conference on Reinforcement Learning and Decision Making (RLDM), 2015, 2017.
- Conference on Artificial Intelligence (AAAI), 2012 (Computational Sustainability Track), 2007, 2010 (Integrated Intelligence Track), 2000.
- International Conference on Machine Learning (ICML), 2000, 2008, 2010.
- Autonomous Agents and Multiagent Systems (AAMAS), 2008.
- International Conf. on Automated Planning and Scheduling. (ICAPS), 2003, 2007.
- Robotics: Science and Systems (RSS), 2006.
- International Joint Conference on Artificial Intelligence (IJCAI), 2001, 2003, 2005.
- International Conference on Autonomic Computing (ICAC), 2004, 2005.
- ACM Conference on Electionic Commerce (EC), 2005.
- Neural Information Processing Systems (NIPS), 2002, 2003.
- European Conference on Machine Learning (ECML), 2001–2003.
- Autonomous Intelligent Networks and Systems Conference (AINS), 2003.
- Distributed Autonomous Robotic Systems (DARS), 2000, 2002.
- Intelligent Autonomous Systems (IAS), 2002.
- International Conference on Artificial Intelligence (IC-AI), 2001.
- Autonomous Agents (AA), 2000, 2001.
- International Conference on Multi-Agent Systems (ICMAS), 2000.
- International Conference on Enterprise Information Systems (ICEIS), 2000.

• Conference paper reviewer:

- International Conference on Learning Representations (ICLR), 2025.
- Multi-disciplinary Conference on Reinforcement Learning and Decision Making (RLDM), 2019
- Autonomous Agents and Multiagent Systems (AAMAS), Blue Sky track, 2015.
- Conference on Artificial Intelligence (AAAI), Senior Member track, 2015.
- International Conference on Intelligent Robots and Systems (IROS), 2013, 2015.
- ASME International Design Engineering Technical Conference (IDETC), 2012.
- Neural Information Processing Systems (NIPS), 2011.
- Autonomous Agents and Multiagent Systems (AAMAS), 2010.
- IEEE International Conference on Robotics and Automation (ICRA), 2005, 2008–2010.
- IEEE Conference on Intelligent Transportation Systems (ITSC), 2009.

• Workshop/symposium advisory committee member:

- Adaptive and Learning Agents Workshop (ALA), 2008–present.
- Workshop Series on Generalization in Planning (GenPlan), 2021-present.
- ICRA Workshop on Human-Centered Robot Learning in the Era of Big Data, 2025.
- AAAI Workshop on Ad Hoc Teamwork (WAHT), 2024.
- LPNMR workshop on Knowledge Representation and Planning in Robotics and Autonomous Systems (KRPRAS), 2017.
- AAAI workshop on Multiagent Interaction without Prior Coordination, 2014–2016.
- AAAI workshop on *Transfer Learning for Complex Tasks*, 2008.

• Workshop/symposium program committee member:

- IJCAI Doctoral Consortium, 2024.
- NeurIPS workshop on *Cooperative AI*, 2021.

- IJCAI workshop on AI for Sports Analytics, 2021.
- AAAI Spring Symposium on Challenges and Opportunities in Multi-Agent Reinforcement Learning (COMARL), 2020.
- ECAI workshop on Safe Machine Learning, 2020.
- IJCAI workshop on *Education in AI K-12*, 2019.
- AAAI Fall Symposium on Natural Communication for Human-Robot Collaboration, 2017.
- RoboCup Symposium, 2001–2005, 2010, 2013, 2015, 2016, 2021, 2023.
- RSS workshop on Combining AI Reasoning and Cognitive Science with Robotics, 2015.
- IJCAI AI video competition, 2009, 2015.
- AAMAS workshop on Multiagent Sequential Decision Making Under Uncertainty, 2013.
- Technical track on Cooperative Multi-Agent Systems and Applications (CMASA),
 ACM Symposium on Applied Computing (SAC), 2013.
- Humanoids workshop on *Humanoid Soccer*, 2012.
- AAMAS workshop on Autonomous Robots and Multirobot Systems (ARMS), 2012.
- AAAI Doctoral Consortium, 2011.
- IEEE International Symposium on Robot and Human Interaction Communication (Ro-Man) associate editor, 2011.
- IJCAI workshop on Trading Agent Design and Analysis (TADA), 2011.
- ECAI workshop on Benchmarking Intelligent (Multi-)Robot Systems (BIMRS), 2010.
- AAAI AI video competition, 2010.
- AAMAS 2010 workshop on Agents in Real-Time and Dynamic Environments, 2010.
- AAAI workshop on *Trading Agent Design and Analysis* (TADA), 2008.
- ECAI workshop on Cognitive Robotics, 2008.
- AAMAS workshop on Formal Models and Methods for Multi-Robot Systems, 2008.
- ICAPS workshop on AI Planning and Learning (AIPL), 2007.
- ICAC workshop on Adaptive Methods in Autonomic Computing (AMACS), 2007.
- Second workshop on Tackling Computer Systems Problems with Machine Learning Techniques (SysML), 2007.
- ICML workshop on Structural Knowledge Transfer for Machine Learning, 2006.
- 3rd International IEEE Latin American Robotic Symposium (LARS), 2006.
- AAMAS workshop on Agent Mediated Electronic Commerce (AMEC/TADA), 2006.
- IJCAI workshop on Planning and Learning in A Priori Unknown or Dynamic Domains, 2005.
- IJCAI workshop on Trading Agent Design and Analysis (TADA), 2005.
- IJCAI workshop on Agents in Real-Time and Dynamic Environments, 2005.
- AAMAS workshop on Learning and Adaptation in MAS (LAMAS), 2005.
- AAMAS workshop on *Teaching Multiagent Systems* (TeachMAS), 2005.
- AAMAS workshop on *Trading Agent Design and Analysis* (TADA), 2004.
- ICML workshop on *Predictive Representations of World State*, 2004.
- IJCAI Workshop on Agents in Dynamic Real-Time Environments, 2003.
- AAMAS Workshop on Resource, Role, and Task Allocation in MAS, 2003.
- Pacific Rim MultiAgent Workshop (PRIMA), 2000–2002.
- AAMAS Workshop on *Coalitions and Team Formation*, 2002.
- AAMAS Workshop on MAS Problem Spaces and Their Implications to Achieving Globally Coherent Behavior Coalitions and Team Formation, 2002.
- Agents, Theories, Architectures, and Languages (ATAL), 2001.
- Agents Workshop on Infrastructure for Agents, Multi-Agent Systems, and Scalable Multi-Agent Systems, 2001.
- ECAI Workshop on Balancing Reactivity and Social Deliberation in Multi-Agent Systems, 2000.
- ICMAS *Collective Robotics* Workshop, 1998.
- IROS Workshop on *RoboCup*, 1996.

• **RoboCup Federation**:

- Past-President 2022–2024.
- President, 2019–2022.
- Vice President, 2013–2019.

- Trustee, 2003-present.
- Executive Committee 1999–present.
- Director, RoboCup US, 2017-present.
- Secretary, RoboCup US, 2016–2017.

• One Hundred Year Study on Artificial Intelligence (AI100):

- Past-Chair, Standing Commitee, 2023–2025.
- Chair, Standing Committee, 2018–2023.
- Chair, First Study Panel, 2015–2016.

• International Joint Conferences on Artificial Intelligence (IJCAI):

- Trustee, 2020–2025.
- Awards Committee, 2010–2015.

• Association for the Advancement of Artificial Intelligence (AAAI):

- Mentor, Undergraduate Mentoring Program, 2021, 2022.
- Conference Committee, 2014–2016, 2021–Present.
- Conference Committee Chair, 2018–2020.
- Fellows Committee, 2016–2018.
- Councilor, 2005–2008.
- Executive Committee, The Multi-disciplinary Conference on Reinforcement Learning and Decision Making (RLDM), 2022–present.
- Selection Committee, ACM SIGAI Autonomous Agents Research Award, 2023.
- Feigenbaum Prize Committee, AAAI 2014–2017.
- Advisory/Steering committee:
 - Center for Artificial Intelligence at University of Sao Paulo, Brazil, 2020–present.
 - MDPI Robotics Journal, 2021–2024.
 - RoboCup Asia-Pacific Tianjin Invitational Tournament, 2021.
 - AAAS committee on AI & Judiciary, 2021.
 - Pacific Rim Trading Agent Competition, 2007.
 - Springer-Verlag book on Balancing Reactivity and Social Deliberation in MAS, 2000–2001.
 - National Academy of Engineering DARPA Prize Authority Workshop, 2000.
- **Board of directors**, Association for Trading Agent Research, 2003–2009.
- Consultant, Information Science and Technology (ISAT) Summer Study on *Automated Intent Recognition on Distributed Organizations* (AIRDO), 2003.

• Project reviewer:

- European Commission (EC) "Composing Learning systems for Artificial Cognitive Systems," (CompLACS) 2012.
- European Commission (EC) "Ubiquitous Networking Robotics in Urban Settings," (URUS) 2008–2010.
- Science Foundation of Ireland (SFI) "Integrated Analysis of System of Systems," 2009.
- UK EPSRC "Market-Based Control of Complex Computational Systems," 2008.

• Proposal reviewer:

- Computing Innovation Fellows, 2020.
- Army Research Office, 2020, 2019.
- United Arab Emirates University, 2019, 2018.
- Czech Science Foundation, 2015.
- European Commission (EC) 2013, 2012, 2010.
- Army Research Office (ARO) 2011.
- Instituto de Telecomunicações (IT), Portugal, 2011.
- National Science Foundation (NSF), 2016, 2014, 2013, 2008–11, 2002–04, 1999.
- US-Israel Binational Science Foundation (BSF), 2007, 2003, 2000.
- Microsoft Scholarship, 2007.
- French Agence Nationale de la Recherche (ANR), 2006.
- City University of New York internal research award program, 2006.
- Research Council of Norway, 2003.
- Israel Science Foundation (ISF), 2010, 2004, 2002.
- Alberta Circle of Research Excellence (iCORE), 2000.

• Departmental and university service:

- Chair, Computer Science Department, 2025–present.
- Chair, Graduate Portfolio Program in Robotics, 2015–present.
- Scientific Board, Machine Learning Laboratory, 2020–present.
- Executive Team Member, Good Systems Bridging Barriers Initiative 2017—present.
- Graduate Studies Committee (GSC) of UT Austin Operations Research and Industrial Engineering (ORIE), 2017–present.
- Director, Texas Robotics, 2017–2025.
- Associate Chair, Computer Science Department, 2015–2025.
- Department Endowment Committee, 2025.
- Interdisciplinary Studies Working Group, 2024–25.
- Search committee, Dir. of TX Institute for Therapeutic Neurotechnology (TITAN), 2024–25.
- Bridging Disciplines Program Committee for Smart Cities, 2019–2025.
- Chair, Graduate Studies Committee (GSC) of Online Masters Program in AI, 2022–2024.
- Faculty Evaluation committee, 2022–23, 2003–04.
- Faculty Recruiting committee, 2020–21(co-chair), 2015–16, 2014–15, 2012–13, 2006–07.
- Online Masters Admissions Committee, 2019–2021.
- President's Council on TEXAS Impact, 2019–20.
- Mechanical Engineering Faculty Search Committee, 2019–20.
- SURA Distinguished Scientist Awards, Internal Review Committee, 2019–20.
- Undergraduate Studies Committee, 2019–20, 2018–19.
- Chair, Provost's Future of Computing Task Force, 2017.
- Chair, Faculty Awards and Honors Committee, 2016–17.
- GDC Advisory Board, 2015–16.
- Search Committee for Neuroscience Department Chair, 2015–16.
- Colloquia committee, 2013–14.
- College of Natural Sciences Strategic Planning Task Force, 2013.
- Doctoral Admissions committee, 2019–20 (ex officio), 2011–12(chair), 2009–10, 2007–08, 2005–06.
- Mechanical Engineering recruiting committee for "Intelligent Physical Systems," 2011.
- Selection committee for the University Best Dissertation award, 2011.
- Chair, Special Events committee, 2010–11, 2009–10, 2004–05.
- Strategic Initiatives committee, 2011–12
- Blunk Memorial Professorship award committee, 2010, 2009.
- AI lab steering committee, 2004–present.
- Graduate Studies Committee (GSC) of UT Austin ECE department, 2008–present.
- Turing Scholars committee, 2008.
- Ad hoc committee on the new WLC policy, 2006.
- Ad hoc Japan Prize committee, 2006.
- Ad hoc GSC committee on diversity proposals, 2005.
- Evaluation of Graduate Programs committee, 2004–05.
- Departmental Best Dissertation committee, 2003.

INVITED DISTINGUISHED LECTURES

- "Outracing Champion Gran Turismo Drivers with Deep Reinforcement Learning" Hans J. Berliner Lecture in Artificial Intelligence at Carnegie Mellon University. Pittsburgh, PA. April 2025.
- "Practical Reinforcement Learning: Lessons from 30 Years of Research"
 - **Keynote Talk at Inaugural Reinforcement Learning Conference.**
 - Amherst, MA. August 2024.
- "Practical Reinforcement Learning: Lessons from 30 Years of Research"
 - **Keynote talk at the Upper Bound Conference.**
 - Edmonton, AB. May 2023.
- "Topics in Multiagent Learning Motivated by Ad Hoc Teamwork"

Keynote talk at The Fourth International Conference on Distributed Artificial Intelligence (DAI).

Tianjin, China (Virtual). December 2022.

• "Outracing Champion Gran Turismo Drivers with Deep Reinforcement Learning"

Computer Science Department Distinguished Lecture at Tulane University.

New Orleans, LA. December 2022.

• "Outracing Champion Gran Turismo Drivers with Deep Reinforcement Learning"

University of Memphis CS Distinguished Colloquium Speaker.

Memphis, TN (Virtual). April 2022.

• "Ad Hoc Autonomous Agent Teams: Collaboration without Pre-Coordination"

Plenary talk at International Joint Conferences on Artificial Intelligence (IJCAI).

Yokohama, Japan (virtual). January 2021.

- "Efficient Robot Skill Learning: Grounded Sim. Learning and Imitation Learning from Observation"
 Instituto Superior Técnico (IST) Distinguished Lecture, co-sponsored by American Corner.
 Lisbon, Portugal (virtual). November 2020.
- "Advances in Ad Hoc Teamwork and Adaptive Tolling for Multiagent Traffic Optimization" **Keynote talk at IEEE International Symposium on Multi-Robot and Multi-Agent Systems (MRS)**. New Brunswick, NJ. August 2020.
- "Efficient Robot Skill Learning: Grounded Sim. Learning and Imitation Learning from Observation" **Distinguished Lecture at Microsoft Research**.

Seattle, WA. August 2019.

• "Efficient Robot Skill Learning: Grounded Sim. Learning and Imitation Learning from Observation" Keynote talk at Amazon Machine Learning Conference.

Seattle, WA. July 2019.

• "Robot Skill Learning: From the Real World to Simulation and Back"

Keynote talk at NVIDIA NTECH Conference.

Santa Clara, CA. September 2017.

• "Artificial Intelligence and Life in 2030"

Keynote talk at Finish Prime Minister's symposium on Artificial Intelligence. Helsinki, Finland. February 2017.

• "Learning and Multiagent Reasoning for Autonomous Robots"

Keynote talk at International Conference on Computer Aided Design (ICCAD).

Austin, TX. November 2016.

• "Autonomous Learning Agents: Layered Learning and Ad Hoc Teamwork"

Keynote talk at Autonomous Agents and Multiagent Systems Conference.

Singapore. May 2016.

• "Learning and Multiagent Reasoning for Autonomous Robots"

Yahoo! Big Thinkers Distinguished Lecture Series.

Sunnyvale, CA. May 2015.

• "Practical RL: Representation, Interaction, Synthesis, and Mortality (PRISM)"

Keynote talk at Florida Artificial Intelligence Research Society (FLAIRS) Conference. Hollywood, Florida. May 2015.

• "Learning and Multiagent Reasoning for Autonomous Robots"

Department of EECS Distinguished Lecture, Vanderbilt University.

Nashville, TN. December 2014.

• "Learning and Multiagent Reasoning for Autonomous Robots"

Department of Computer Science Distinguished Lecture, George Mason University.

Fairfax, VA. November 2014.

• "Practical RL: Representation, Interaction, Synthesis, and Mortality (PRISM)"

ML/Google Distinguished Lecture, Carnegie Mellon University.

Pittsburgh, PA. November 2014.

• "Ad Hoc Autonomous Agent Teams: Collaboration without Pre-Coordination"

UCLA Electrical Engineering Distinguished Seminar Series.

Los Angeles, CA. October 2014.

• "Learning and Multiagent Reasoning for Autonomous Robots"

Drexel University College of Computing & Informatics 6th Annual **Jay Modi Memorial Lecture**. Philadelphia, PA. March 2014.

• "PRISM – Practical RL: Representation, Interactions, Synthesis, and Mortality"

 $\label{lem:conditional} \textbf{Keynote talk at International Symposium on Artificial Intelligence and Math (ISAIM)}.$

Ft. Lauderdale, FL. January 2014.

• "PRISM – Practical RL: Representation, Interactions, Synthesis, and Mortality"

Keynote talk at 9th European Workshop on Reinforcement Learning (EWRL).

Athens, Greece. September 2011.

• "Learning and Multiagent Reasoning for Autonomous Agents"

UC Irvine Computer Science Department Distinguished Lecturer.

Irvine, California. October 2009.

• "Teaching Teammates in Ad Hoc Teams"

Keynote talk at 10th Bar-Ilan Symposium on the Foundations of Artificial Intelligence (BIS-FAI).

Ramat-Gan, Israel. June 2009.

• "Learning and Multiagent Reasoning for Autonomous Agents"

Keynote talk at 4th IEEE Latin American Robotic Symposium (LARS).

Monterrey Mexico. November 2007.

• "Learning and Multiagent Reasoning for Autonomous Agents"

UT Austin Visions of Computing Lecture.

Austin, Texas. November 2007.

• "Learning and Multiagent Reasoning for Autonomous Agents"

IJCAI Computers and Thought Award Lecture.

Hyderabad, India. January 2007.

• "Robot Learning"

National Academy of Sciences spring symposium.

Washington, DC. April 2006.

• "Machine Learning on Physical Robots"

Keynote talk at International Conference on Computing (CIC).

Mexico City. October 2004.

• "The Trading Agent Competition: Two Champion Adaptive Bidding Agents"

Keynote talk at Computer Games Conference.

Edmonton, Alberta. July, 2002.

INVITED TALKS

• "The Role of AI in Robotics"

Special Competitive Studies Project *Exchange on Innovation and National Security.*

Austin, TX. September 2025.

• "Human-in-the-Loop Machine Learning for Robot Navigation and Manipulation"

Keynote talk at the **Thompson Reuters Institute** *Emerging Technology and Generative AI Forum*. Austin, TX. September 2025.

• "Getting a Job after Grad School"

AAAI Doctoral Consortium.

Montreal, Canada. August 2025.

• "Building Minimal and Reusable Causal State Abstractions for Reinforcement Learning"

ONR Science of Autonomy Program Review.

Virtual. August 2025.

• "Human-in-the-Loop Machine Learning for Robot Navigation and Manipulation"

Austin Forum on Technology & Society.

Austin, TX. August 2025.

• "Human-in-the-Loop Machine Learning for Robot Navigation and Manipulation"

Keynote talk at **Texas Department of Information Resources** AI Day.

Austin, TX. August 2025.

• "Human-in-the-Loop Machine Learning for Robot Navigation and Manipulation"

AIM 2025 Workshop on Data-Enabled Learning Control for Intelligent Mechatronics and Robotics. Hangzhou, China (Virtual). July 2025.

• "Leveraging Human Social Navigation Prowess for Robot Social Navigation"

ICRA 2025 Workshop on Advances in Social Navigation.

Atlanta, GA. May 2025.

• "Leveraging Human Social Navigation Prowess for Robot Social Navigation"

ICRA 2025 Workshop on Building a Common Humanoid Platform Infrastructure.

Atlanta, GA. May 2025.

• "Human-in-the-Loop Machine Learning for Robot Navigation and Manipulation"

Keynote talk at Information Systems and Computing Technology Network Symposium.

Tuscson, AZ (Virtual). April 2025.

• "Human-in-the-Loop Machine Learning for Robot Navigation and Manipulation"

(Georgia) Tech AI Fest.

Atlanta, GA. March 2025.

- SxSW, panel on *Pushing Creativity to New Bounds: Future Robot Applications* Austin, TX. March 2025.
- "Advances in Ad Hoc Teamwork: Multiagent Collaboration without Pre-Coordination"

AAAI 2025 Workshop on Cooperative Multi-Agent Systems Decision-Making and Learning. Philadelphia, PA. March 2025.

• "Learning for Integrated Task and Motion Planning"

AAAI 2025 Bridge Program on Learning for Integrated Task and Motion Planning.

Philadelphia, PA. February 2025.

• "Outracing Champion Gran Turismo Drivers with Deep Reinforcement Learning"

Texas Tech University CS Seminar.

Lubbock, TX (online). February 2025.

• "Multiagent Reinforcement Learning: Cooperation and Competition"

MIT Physics Colloquium.

Boston, MA. February 2025.

• "Human-in-the-Loop Machine Learning for Robot Navigation and Manipulation"

University of Washington Robotics Seminar.

Seattle, WA. February 2024.

• "Multiagent Reinforcement Learning: Cooperation and Competition"

Keynote Talk at NSF Workshop on Reinforcement Learning.

Boston, MA. January 2025.

• "Human-in-the-Loop Machine Learning for Robot Navigation and Manipulation"

Keynote Talk at Southwest Robotics Symposium.

Tempe, AZ. November 2024.

• "Human-in-the-Loop Machine Learning for Robot Navigation and Manipulation"

Duke University Robotics Seminar.

Durham, NC. October 2024.

• "Human in the Loop Learning for Robot Navigation and Task Learning from Implicit Human Feedback"

University of Toronto Robotics Institute Seminar.

Toronto, Canada (Virtual). September 2024.

• "The Future"

Keynote at UT Austin Portugal Colloquium.

Austin, TX. September 2024.

• "The Importance of Public Engagement"

Panel Moderator at **the 33rd International Joint Conference on Artificial Intelligence (IJCAI).** Jeju, South Korea. August 2024.

• "Human in the Loop Learning for Robot Navigation and Task Learning from Implicit Human Feedback"

RSS 2024 Workshop on Task Specification for General-Purpose Intelligent Robots.

Delft, The Netherlands (Virtual). July 2024.

• "Ethics, AI, and the Future"

Congregation Agudas Achim.

Austin, TX. July 2024.

• "Future of Science and Technology with AI"

Keynote at Inaugural UT System AI Symposium in Health Care.

Dallas, TX. May 2024.

• "Human-in-the-Loop Machine Learning"

Interagency Working Group (IWG) on Intelligent Robotics & Autonomous Systems (IRAS) Seminar Series.

Washington, DC (Virtual). May 2024.

• "Humans vs. AI: Robot Soccer and Gran Turismo"

Hot Science - Cool Talks Outreach Lecture Series.

Austin, TX. April 2024.

• "Human-in-the-Loop Machine Learning

Interagency Working Group (IWG) on AI Research & Development Seminar Series.

Washington, DC (Virtual). April 2024.

• "Getting a Job after Grad School"

AAAI Doctoral Consortium.

Vancouver, BC, Canada. February 2024.

• "Practical Reinforcement Learning: Lessons from 30 Years of Research"

Texas A&M University-Kingsville, College of Engineering and EECS Department Virtual Distinguished Lecture Series.

Kingsville, TX (virtual). February 2024.

• "Practical Reinforcement Learning: Lessons from 30 Years of Research"

8th International Conference on Robotics, Control and Automation (ICRCA).

Shanghai, China (virtual). January 2024.

• "Causal Dynamics Learning for Task-Independent State Abstraction"

NeurIPS 2023 workshop on Generalization in Planning (GenPlan).

New Orleans, LA. December 2023.

• "Broader Societal Concerns with LPTMs"

Army Strategy Planning and Science Meeting on Large Pre-Trained Models for DoD AI/ML: Opportunities, Challenges, and Applications.

Arlington, VA. November 2023.

• "Overview of work at Texas Robotics"

Zebra Technologies Innovation Expo2028.

Virtual. November 2023.

• "Causal Dynamics Learning for Task-Independent State Abstraction"

CoRL 2023 workshop on Learning Effective Abstractions for Planning (LEAP).

Atlanta, GA (Virtual). November 2023.

• "Human in the Loop Learning for Robot Navigation and Task Learning from Implicit Human Feedback"

Keynote at IEEE Latin American Robotics Symposium.

Salvador, Brazil (Virtual). October 2023.

• "AI and its Effects on Education and Society"

Texas A&M — Commerce.

Commerce, TX (Virtual). October 2023.

• "Repurposing Facilities for Robotics Education, Research, and Partnerships"

Tradeline Conference on University Science and Engineering Facilities.

Austin, TX. October 2023.

• "DM²: Decentralized Multiagent Reinforcement Learning via Distribution Matching"

IROS 2023 workshop on Advances in Multi-Agent Learning - Coordination, Perception, and Control.

Detroit, MI (Virtual). October 2023.

• "DM2: Decentralized Multiagent Reinforcement Learning via Distribution Matching"

IROS 2023 workshop on MAD-Games: Multi-Agent Dynamic Games.

Detroit, MI (Virtual). October 2023.

• "Human in the Loop Learning for Robot Navigation and Task Learning from Implicit Human Feedback"

Cornell Robotics Seminar.

New York, NY (Virtual). August 2023.

• "Ad Hoc Teamwork with Communication"

ONR Science of Autonomy Meeting.

Virtual. August 2023.

• "Autonomous Agents and Multiagent Systems for Social Good"

Keynote at RTX ISaCTN Symposium.

Virtual. June 2023.

• "Coopernaut: End-to-End Driving with Cooperative Perception for Networked Vehicles"

ICRA 2023 workshop on CoPerception: Collaborative Perception and Learning.

London, UK. June 2023.

• "Practical Reinforcement Learning: Lessons from 30 Years of Research"

AAMAS 2023 workshop on Adaptive Learning Agents (ALA).

London, UK. May 2023.

• "Human in the Loop Learning for Robot Navigation and Task Learning from Implicit Human Feedback"

ICRA 2023 workshop on Life-Long Learning with Human Help (L3H2).

London, UK. May 2023.

• "AI100: The One Hundred Year Study on AI"

Headliner's Club.

Austin, TX. May 2023.

• "AI100: The One Hundred Year Study on AI"

Keynote at the Florida State Machine Learning Expo (MLX23).

Tallahassee, FL (Virtual). April 2023.

• "Outracing Champion Gran Turismo Drivers with Deep Reinforcement Learning"

Missouri University of Science and Technology Department of Computer Science Seminar Series.

Rolla, MO (Virtual). April 2023.

• "Outracing Champion Gran Turismo Drivers with Deep Reinforcement Learning"

University of California San Diego Contextual Robotics Institue Seminar.

San Diego, CA (Virtual). March 2023.

• "(Mis)design for Autonomous Driving and Accumulating Safety Rules from Catastrophic Action Effects"

AAAI 2023 workshop on Deployable AI).

Washington, DC. February 2023.

• "Outracing Champion Gran Turismo Drivers with Deep Reinforcement Learning"

NeurIPS 2022 workshop on Reinforcement Learning for Real Life.

New Orleans, LA. December 2022.

• "Human in the Loop Learning for Robot Navigation and Task Learning from Implicit Human Feedback"

NeurIPS 2022 workshop on Human in the Loop Learning.

New Orleans, LA. December 2022.

• "Outracing Champion Gran Turismo Drivers with Deep Reinforcement Learning"

Keynote at University of Georgia AI Research Day.

Athens, GA (Virtual). November 2022.

• "The Value of Communication in Ad Hoc Teamwork"

IROS 2022 workshop on Decision Making in Multi-Agent Systems.

Kyoto, Japan (Virtual). October 2022.

• "(Mis)design for Autonomous Driving and Accumulating Safety Rules from Catastrophic Action Effects".

IROS 2022 workshop on Behavior-driven Autonomous Driving in Unstructured Environments.

Kyoto, Japan (Virtual). October 2022.

• "Outracing Champion Gran Turismo Drivers with Deep Reinforcement Learning"

Edge of Now (EON).

Palm Springs, CA. October 2022.

• "Outracing Champion Gran Turismo Drivers with Deep Reinforcement Learning" University of Wiconsin seminar on Systems, Information, Learning, Optimization (SILO). Madison, WI (Virtual). September 2022.

• "Outracing Champion Gran Turismo Drivers with Deep Reinforcement Learning"

9th International Conference on Signal Processing and Integrated Networks (SPIN).

Amity University, Uttar Pradesh, India (virtual). August 2022.

• "Reward (Mis)design for Autonomous Driving and Accumulating Safety Rules from Catastrophic Action Effects"

ICML 2022 workshop on Safe Learning for Autonomous Driving (SL4AD).

Baltimore, MD (Virtual). July 2022.

• "Topics in Multiagent Learning Motivated by Ad Hoc Teamwork"

Berkeley MARL Seminar.

Virtual. July 2022.

• "The EMPATHIC Framework for Task Learning from Implicit Human Feedback"

RLDM 2022 workshop on Reinforcement Learning with Humans in (and around) the Loop. Providence, RI. June 2022.

• "Grounded Simulation Learning for Sim2Real"

ICLR 2022 workshop on Generalizable Policy Learning in the Physical World.

Virtual. April 2022.

• "Outracing Champion Gran Turismo Drivers with Deep Reinforcement Learning"

UT Joint Forum for AI (FAI) and Machine Learning Lab (ML+X) seminar.

Austin, TX. March 2022.

• "Outracing Champion Gran Turismo Drivers with Deep Reinforcement Learning" **SysML Workshop**.

Austin, TX. March 2022.

• "Outracing Champion Gran Turismo Drivers with Deep Reinforcement Learning"

NVIDIA GPU Technology Conference (GTC).

Virtual. March 2022.

• "Autonomous Agents and Multiagent Systems for Social Good"

ISI AI Seminar Series.

Marina Del Rey, CA (Virtual). March 2022.

• "Outracing Champion Gran Turismo Drivers with Deep Reinforcement Learning"

AAAI Workshop on Reinforcement Learning in Games.

Virtual. February 2022.

• "Estimation and Control of Visitation Distributions for Reinforcement Learning"

Highligted talk at NVIDIA Reinforcement Learning Workshop.

Virtual. January 2022.

• "Machine Learning for Robot Locomotion: Grounded Simulation Learning and Adaptive Planner Parameter Learning"

Keynote at IEEE International Conference on Big Data.

Virtual. December 2021.

• "Machine Learning for Robot Locomotion: Grounded Simulation Learning and Adaptive Planner Parameter Learning"

Keynote at **6th International Conference on Robotics and Automation Engineering (ICRAE)**. Guangzhou, China (virtual). November 2021.

• "Machine Learning for Robot Locomotion: Grounded Simulation Learning and Adaptive Planner Parameter Learning"

University of Sheffield Automatic Control and Systems Engineering department Seminar. Sheffield, UK (virtual). November 2021.

• "Efficient Robot Skill Learning: Grounded Sim. Learning and Imitation Learning from Observation"

Oden Institute Babuška Forum.

Austin, TX. November 2021.

• "Machine Learning for Robot Locomotion: Grounded Simulation Learning and Adaptive Planner

Parameter Learning"

Oxford Robotics Institute Seminar.

Oxford, UK (virtual). October 2021.

• "Grounded Simulation Learning for Sim2Real"

Réseau Régional de Recherche en Robotique (R4) Seminar.

Bordeaux, France. October 2021.

• "Machine Learning for Robot Locomotion: Grounded Simulation Learning and Adaptive Planner Parameter Learning"

Robert Bosch Centre for Data Science and Artificial Intelligence LatentView Colloquium at IIT, Madras.

Madras, India (virtual). September 2021.

• "Artificial Intelligence and Machine Learning Research"

Jefferies Asia Forum.

Hong Kong (virtual). September 2021.

• "Task Planning and Learning for General Purpose Service Robots"

Keynote at 9th ICAPS Workshop on Planning and Robotics (PlanRob).

Guangzhou, China (virtual). August 2021.

• "The RoboCup Grand Challenge for AI and Robotics"

World Artificial Intelligence Conference.

Shanghai, China (virtual). July 2021.

• "Coach-Player Multi-Agent Reinforcement Learning for Dynamic Team Composition"

Lockheed Martin AI Summit.

Virtual. July 2021.

• "Learning and Multiagent Reasoning for Autonomous Robots"

Texas McCombs Real Estate Center.

Austin, TX (virtual). May 2021.

• "Machine Learning for Robot Locomotion: Grounded Simulation Learning and Adaptive Planner Parameter Learning"

Keynote at The IV Brazilian Humanoid Robot Workshop and the V Brazilian Workshop on Service Robotics.

Virtual. May 2021.

• "Efficient Robot Skill Learning: Grounded Sim. Learning and Imitation Learning from Observation" Secondmind Labs Research Seminar.

London, UK (virtual). May 2021.

• "Machine Learning for Robot Locomotion: Grounded Simulation Learning and Adaptive Planner Parameter Learning"

Technion Robotics Seminar.

Haifa, Israel (virtual). May 2021.

• "Autonomous Learning Agents and Multiagent Systems for Social Good"

Keynote at AAMAS workshop on Autonomous Agents for Social Good (AASG)

London, UK (virtual). May 2021.

• "Efficient Robot Skill Learning: Grounded Sim. Learning and Imitation Learning from Observation". Keynote at IEEE International Conference on Autonomous Robot Systems and Competitions.

Santa Maria de Feira, Portugal (virtual). April 2021."Topics in Multiagent Learning Motivated by Ad Hoc Teamwork"

Seminar on Challenges and Opportunities for Multi-Agent Reinforcement Learning (COMARL). Virtual. February 2021.

• "Ad Hoc Autonomous Agent Teams: Collaboration without Pre-Coordination"

AAAI Workshop on Plan, Activity, and Intent Recognition (PAIR).

Virtual. February 2021.

• "Ad Hoc Autonomous Agent Teams: Collaboration without Pre-Coordination"

NeurIPS Workshop on Cooperative AI.

Virtual. December 2020.

• "Grounded Sim. Learning for Sim2Real with Connections to Off-Policy Reinforcement Learning"
NeurIPS Workshop on Deep Reinforcement Learning.

Virtual. December 2020.

• "Task-Motion Navigation Planning with Learning for Adaptable Mobile Service Robots"

ICAPS Workshop on Bridging the Gap between Planning and Reinforcement Learning (PRL). Nancy, France (virtual). October 2020.

• "Efficient Robot Skill Learning: Grounded Sim. Learning and Imitation Learning from Observation" West Virginia University Robotics Seminar Series.

Morgantown, WV (virtual). September 2020.

• "Artificial Inteligence Research"

Youth AI Lab and OpenCode Foundation Seminar.

Virtual. September 2020.

• "Off-Policy Evaluation for Grounded Simulation Learning"

ONR Science of Autonomy Meeting.

Arlington, VA (virtual). August 2020.

• "Efficient Robot Skill Learning via Grounded Sim. Learning, Imitation Learning from Observation, and Off-Policy Reinforcement Learning"

Institute of Advanced Studies Theoretical Machine Learning Seminar.

Princeton, NJ (virtual). July 2020.

• "Grounded Simulation Learning"

Army Reserach Lab Workshop on Synthetic Data in AI/ML.

Virtual. July 2020.

• "Efficient Robot Skill Learning: Grounded Sim. Learning and Imitation Learning from Observation" Apptronik.

Austin, TX (virtual). June 2020.

• "Learning Robot Behaviors from Other Agents" (with Garrett Warnell)

ICRA 2020 Workshop on Interactive Robot Learning.

Virtual. June 2020.

• "Imitation Learning from Observation"

On-device Intelligence Workshop at Machine Learning and Systems (MLSys).

Austin, TX. March 2020.

• "Task-Motion Planning with Learning for Adaptable Mobile Service Robots"

AAAI 2020 Workshop on Generalization in Planning.

New York, NY. February 2020.

• "Efficient Robot Skill Learning: Grounded Sim. Learning and Imitation Learning from Observation" Rice Computer Science Department Colloquium.

Houston, TX. January 2020.

• "Machine Learning and Multiagent Reasoning for Autonomous Robots"

JASON fall meeting.

McLean, VA. November 2019.

• "Efficient Robot Skill Learning: Grounded Sim. Learning and Imitation Learning from Observation" Cornell Computer Science Department Colloquium.

Ithaca, NY. November 2019.

• "Adaptive Tolling for Multiagent Traffic Optimization and Imitation Learning from Observation"

MIT Civil and Environmental Engineering Henry L. Pierce Laboratory Seminar Series.

Boston, MA. November 2019.

• "Machine Learning and Artificial Intelligence for Autonomous Robots"

Sage Perspectives on the Future Conference.

Austin, TX. November 2019.

• "Efficient Robot Skill Learning: Grounded Sim. Learning and Imitation Learning from Observation" Northwestern University CS Colloquium.

Chicago, IL. October 2019.

• "Efficient Robot Skill Learning: Grounded Sim. Learning and Imitation Learning from Observation" CTRL-labs.

New York, NY. August 2019.

• "Efficient Robot Skill Learning: Grounded Sim. Learning and Imitation Learning from Observation" IJCAI workshop on Scaling Up Reinforcement Learning (SURL).

Macau, China. August 2019.

• "Learning Curricula for Transfer Learning in RL"

ICML Workshop on Multi-Task and Lifelong Reinforcement Learning.

Long Beach, CA. June 2019.

• "Adaptive Tolling for Multiagent Traffic Optimization"

ICML Workshop on AI in Finance: Applications and Infrastructure for Multi-Agent Learning. Long Beach, CA. June 2019.

• "Learning and Multiagent Reasoning for Autonomous Robots"

Silicon Labs Technical Symposium

Austin, TX. May 2019.

• "Efficient Robot Skill Learning: Grounded Sim. Learning and Imitation Learning from Observation" SUNY Binghamthon Computer Science Invited Speaker Series.

Binghamton, NY. April 2019.

• "The Robocup Challenge (2050) and Human in the Loop Machine Learning"

U.S. Army Mad Scientist Conference.

Austin, TX. April 2019.

• "Artificial Intelligence and Life in 2030"

Washington University in St. Louis Division of Computational and Data Sciences Colloquium. St. Louis, MO. April 2019.

• "Efficient Robot Skill Learning: Grounded Sim. Learning and Imitation Learning from Observation" Washington University in St. Louis Computer Science & Engineering Colloquia Series. St. Louis, MO. April 2019.

• "Efficient Robot Skill Learning: Grounded Sim. Learning and Imitation Learning from Observation" University of Maryland Robotics Center Seminar.

College Park, MD. March 2019.

• "Efficient Robot Skill Learning: Grounded Sim. Learning and Imitation Learning from Observation" University of Washington Robotics Colloquium.

Seattle, WA. February 2019.

• "Control Algorithms for Imitation Learning from Observation"

RE•WORK Deep Learning Summit. San Francisco, CA. January 2019.

• "Knowledge and Planning for Autonomous Service Robots"

Army Science Planning and Strategy Meeting Intelligent Systems with Real-Time Learning, Knowledge Bases, and Information Retrieval.

Austin, TX. January 2019.

• "Control Algorithms for Imitation Learning from Observation"

NeurIPS workshop on Learning by Instruction.

Montreal, Canada. December 2018.

• "Control Algorithms for Imitation Learning from Observation"

NeurIPS workshop on RL Under Partial Observability.

Montreal, Canada. December 2018.

• "Control Algorithms for Imitation Learning from Observation"

NeurIPS workshop on Imitation Learning and its Challenges in Robotics.

Montreal, Canada. December 2018.

• "Artificial Intelligence: Applications and Developments in Industry"

Huawei Technology Symposium.

Dublin, Ireland. September 2018.

• "Robot Skill Learning: From the Real World to Simulation and Back"

Toyota Technical Institute — Chicago.

Chicago, IL. August 2018.

• "Off-Policy Evaluation for Grounded Simulation Learning"

ONR Science of Autonomy Meeting.

Arlington, VA. August 2018.

• "Robot Skill Learning: From the Real World to Simulation and Back"

Federated AI for Robotics Workshop.

Stockholm, Sweden. July 2018.

• "Robot Skill Learning: From the Real World to Simulation and Back"

McGill Computer Science Department Seminar.

Montreal, CA. June 2018.

• "Robot Skill Learning: From the Real World to Simulation and Back"

ARM technical talk.

Austin, TX. June 2018.

• "AI and Life in 2030"

Huawei Strategic Technologies Workshop.

Shenzhen, China. May 2018.

• "Machine Learning and Artificial Intelligence for Autonomous Robots"

University of Chicago Alumni Association of Austin. Austin, TX. May 2018.

• "Lifelong Learning of Perception and Action in Autonomous Systems"

Kickoff meeting of DARPA Lifelong Learning Machines (L2M) Program.

Arlington, VA. March 2018.

• "AI and Robotics"

Futures Panel at MBN conference.

Seoul, South Korea. February 2018.

• "Machine Learning and Artificial Intelligence for Autonomous Robots"

AAAS 2018 Flash Talk on Developing Robotics to Assist Human.

Austin, TX. February 2018.

• "AI and Travel in 2030"

HEDNA 2018 Austin Global Distribution Conference.

Austin, TX. January 2018.

• "Artificial Intelligence and Life in 2030"

UT Austin OLLI Seminar for Adult Growth and Enrichment (SAGE).

Austin, TX. January 2018.

• "Robotics and AI"

Keynote at SparkCognition Time Machine Conference.

Austin, TX. December 2017.

• "Machine Learning and AI for Autonomous Robots"

Filene Research Institute's big. bright. minds. conference.

Austin, TX. December 2017.

• "Machine Learning and AI for Autonomous Robots"

Michael and Susan Dell Foundation meeting.

Austin, TX. November 2017.

• "Machine Learning and AI for Autonomous Robots"

International Lawyers Network meeting.

Austin, TX. November 2017.

• "Machine Learning and AI for Autonomous Robots"

Keynote at NXP Foundation Women's STEM Leadership Forum for Women.

Austin, TX. October 2017.

• "Robot Skill Learning: From the Real World to Simulation and Back"

Google Brain Tech Talk.

Mountain View, CA. August 2017.

• "Machine Learning and AI for Autonomous Robots"

Keynote at **Deutsche Bank Internet Conference**.

Palo Alto, CA. August 2017.

• "Reasoning about Hypothetical Agent Behaviours and their Parameters"

ONR Science of Autonomy Meeting.

Arlington, VA. August 2017.

• "Learning and Multiagent Reasoning for Autonomous Robots"

Sony Deep Learning Seminar.

Nagoya, Japan. July 2017.

• "Intersection of the Future: Possibilities for Autonomous Vehicles"

World Robotics x AI Seminar.

Nagoya, Japan. July 2017.

• "Robot Skill Learning: From the Real World to Simulation and Back"

AAMAS 2017 workshop on Optimisation in Multi-Agent Systems (OptMAS).

Sao Paulo, Brazil. May 2017.

• "Robot Skill Learning: From the Real World to Simulation and Back"

Intel Tech Talk.

Austin, TX. April 2017.

• "Robot Skill Learning: From the Real World to Simulation and Back"

Carnegie Mellon Robotics Institute Seminar.

Pittsburgh, PA. March 2017.

• "Learning and Multiagent Reasoning for Autonomous Robots"

Keynote at **Synopsis SNUG Conference**.

Santa Clara, CA. March 2017.

• "Artificial Intelligence and Life in 2030"

Keynote at IDG AGENDA17 Conference.

Ponte Vedra Beach, FL. March 2017.

• "Artificial Intelligence and Life in 2030"

IEEE Workshop on Advanced Robotics and its Social Applications (ARSO).

Austin, TX. March 2017.

• "Artificial Intelligence and Life in 2030"

Yale University CS Talk.

New Haven, CT. February 2017.

• "Robot Skill Learning: From the Real World to Simulation and Back"

Yale University CS Talk.

New Haven, CT. February 2017.

• "Artificial Intelligence and Life in 2030"

Sackler U.S.-U.K. Scientific Forum on the Frontiers of Machine Learning.

Washington, DC. January 2017.

• "Cerebellar Learning for Robotics and Deep Multi-Robot Learning"

NIPS 2016 workshop on Neurorobotics.

Barcelona, Spain. December 2016.

• "Robots that Learn to Communicate through Natural Human Dialog"

National Robotics Initiative 2016 PI Meeting.

Arlington, VA. November 2016.

• "Human-Aware Navigation in Populated Indoor Environments"

National Robotics Initiative 2016 PI Meeting.

Arlington, VA. November 2016.

• "Intersections of the Future: Possibilities for Autonomous Vehicles"

Keynote at Texas Wireless Summit.

Austin, TX. October 2016.

• "Autonomous Learning Agents"

Global Semiconductor Alliance (GSA) Executive Forum

Menlo Park, CA. October 2016.

• "Artificial Intelligence and Life in 2030"

UT Department of Computer Science Forum for AI.

Austin, TX. September 2016.

• "Artificial Intelligence and Life in 2030"

Gigaom Change.

Austin, TX. September 2016.

• "AI and Robotics Research in Sony: Past, Present, and Future"

Keynote at **IJCAI 2016 Industry Day** (with Hiroaki Kitano and Masahiro Fujita). New York, NY. July 2016.

• "Deep Multiagent RL for Partially Observable Parameterized Environments"

IJCAI 2016 Workshop on Deep Reinforcement Learning: Frontiers and Challenges.

New York, NY. July 2016.

• "IML for Building-Wide Intelligence"

IJCAI 2016 Workshop on Interactive Machine Learning.

New York, NY. July 2016.

• "Future Trends in Machine Learning"

AT&T Automation and Machine Learning Summit.

Middleton, NJ. June 2016.

• "Reinforcement Learning for Sequential Decision Making"

Exxon Mobil Cognitive Computing Event.

Houston, TX. March, 2016.

• "Practical RL: Representation, Interaction, Synthesis, and Mortality (PRISM)"

University of California at Berkeley EECS Seminar Series.

Berkeley, CA. January 2016.

• "AI as a Gradual, Long-term, Community-Wide Effort"

The NYU Future of AI Symposium.

New York, NY. January 2016.

• "Towards a Unification of Pardigmatic Realizations of Multiagent Systems"

NIPS 2015 Workshop on Learning, Inference and Control of Multi-Agent Systems.

Montreal, Canada. December 2015.

• "Towards a Greater Understanding of the Cerebellum via Experiments Motivated by Machine Learning"

AFOSR Program Meeting: Computational Cognition and Machine Intelligence Program.

Arlington, VA. November 2016.

• "HRI for Building-Wide Intelligence"

AAAI Fall Symposium on AI for Human-Robot Interaction.

Arlington, VA. November 2015.

• "Practical RL: Representation, Interaction, Synthesis, and Mortality (PRISM)"

AAAI Fall Symposium on Embedded Machine Learning.

Arlington, VA. November 2015.

• "Practical Reinforcement Learning for Robots and Autonomous Agents"

Neuroscience Seminar Series, UT Austin.

Austin, TX. October 2015.

• "Machine Learning Opportunities in Marketing"

Forbes CMO Excursions.

Palo Alto, CA. September 2015.

• "Making Friends on the Fly: Advances in Ad Hoc Teamwork"

ONR Science of Autonomy Meeting.

Washington, DC. August 2015.

• "Practical RL: Representation, Interaction, Synthesis, and Mortality (PRISM)"

Reinforcement Learning and Decision Making Conference (RLDM).

Edmonton, Alberta, Canada. June 2015.

• "Learning Agents for Sustainable Energy"

Plenary speaker at ExxonMobil's Longer Range Research Meeting.

Baltimore, MD. May 2015.

• "Autonomous Robots: from robot soccer to driverless cars"

Keynote at Harvard/USC Real Estate Summit.

Aspen, CO. February 2015.

• "A Neuroevolution Approach to Atari Game Playing"

AAAI Workshop on General Competency in Video Games.

Austin, TX. January 2015.

• "Robots at the Boundary of Robotics and AI"

AAAI NSF Sponsored Workshop: Research Issues at the Boundary of AI and Robotics.

Austin, TX. January 2015.

• "Transfer Learning for Autonomous Robots"

DARPA/ISAT workshop on Training of Things.

San Francisco, CA. November 2014.

• "Transfer Learning for Autonomous Robots"

AAAI Fall Symposium on Knowledge, Skill, and Behavior Transfer in Autonomous Robots.

Arlington, VA. November 2014.

• "Learning and Multiagent Reasoning for Autonomous Robots"

Center for Perceptual Systems Seminar Series, UT Austin.

Austin, TX. October 2014.

• "Autonomous Robots: from Robot Soccer to Driverless Cars"

PREA Institutional Investor Real Estate Conference.

Los Angeles, CA. Septeber 2014.

• "Learning and Multiagent Reasoning for Autonomous Robots"

Cognitive Science Seminar Series, UT Austin.

Austin, TX. September 2014.

• "Ad Hoc Autonomous Agent Teams: Collaboration without Pre-Coordination"

AAAI Workshop on Multiagent Interaction without Prior Coordination (MIPC).

Quebec City, Quebec, Canada. July 2014.

• "Learning and Multiagent Reasoning for Autonomous Robots"

Bar-Ilan University Computer Science colloquium.

Ramat Gan, Israel. July 2014.

• "Learning and Multiagent Reasoning for Autonomous Robots"

Word for Word Lecture Series.

Austin, TX. June 2014.

• "Learning Agents for Sustainable Energy"

3M.

Austin, TX. June 2014.

• "Learning and Multiagent Reasoning for Autonomous Robots"

Keynote at REFIC Spring Real Estate Conference.

Austin, TX. May 2014.

• "Learning Agents for Sustainable Energy"

AAMAS workshop on Trading Agent Design and Analysis (TADA).

Paris, France. May 2014.

• "Ad Hoc Autonomous Agent Teams: Collaboration without Pre-Coordination"

Game Theory and Human Behavior (GTHB) Symposium.

Los Angeles, USA. April 2014.

• "Learning Agents for Sustainable Energy"

Texas-Wisconsin-California Control Consortium.

Austin, TX. April 2014.

• "Learning and Multiagent Reasoning for Autonomous Robots"

University of Michigan CSE Colloquium.

Ann Arbor, MI. January 2014.

• "Intersections of the Future: Leveraging Fully Autonomous Vehicles"

Andreessen Horowitz Academic Round Table.

Menlo Park, CA. September 2013.

• "Intersections of the Future: Using Fully Autonomous Vehicles"

16th Annual Transportation and Infrastructure Summit.

Dallas, TX. August 2013.

• "Can Robots Play Soccer Better than People?"

TEDxYouth@Austin.

Austin, TX. March 2013.

• "Solve for Traffic Congestion and Fatalities"

Google Solve for [X] event at SXSW.

Austin, TX. March 2013.

• "UT Austin Villa: RoboCup 2012 Champions in the Standard Platform League"

Aldebaran Webinar.

February 2013.

• "Intersections of the Future: Using Fully Autonomous Vehicles"

IEEE CS Austin Section.

Austin, TX. August 2012.

• "UT Austin Villa: RoboCup 2012 Champions in the Standard Platform and 3D Simulation Leagues" **AAAI 2012 Workshop on Cognitive Robotics**.

Toronto, Ontario, Canada. July 2012.

• "Generalized Model Learning for Reinforcement Learning on a Humanoid Robot"

Aldebaran Tech Day.

Mexico City, Mexico. June 2012.

• "Intersections of the Future: Using Fully Autonomous Vehicles"

American Association for the Advancement of Science (AAAS).

Vancouver, Canada. February 2012.

• "Autonomous Robots Playing Soccer and Traversing Intersections"

UT Learning Activities for Mature People (LAMP).

Austin, TX. January 2012.

• "Autonomous Robots Playing Soccer and Traversing Intersections"

Lakeway Men's Breakfast Club.

Lakeway, TX. January 2012.

• "Active Learning for Sequential Sensing and Efficient Human Interactions in Collaborative Human-Robot Teams"

ONR Science of Autonomy Meeting.

Arlington, VA. December 2011.

• "Machine Learning and Multiagent Reasoning: from Robot Soccer to Autonomous Traffic"

Texas A&M Computer Science Department Symposium.

College Station, TX. December 2011.

• "Ad Hoc Autonomous Agent Teams: Collaboration without Pre-Coordination"

AAAI Fall Symposium on Multiagent Coordination under Uncertainty.

Washington, DC. November 2011.

• "Intersections of the Future: Using Fully Autonomous Vehicles"

AAMAS 2011 Workshop on Agents and Data Mining Interaction.

Taipei, Taiwan. May 2011.

• "Human-Assisted Reinforcement Learning"

ONR Science of Autonomy Meeting.

Arlington, VA. April 2011.

• "Machine Learning and Multiagent Reasoning: from robot soccer to autonomous traffic"

Texas State Computer Science Department Seminar.

San Marcos, TX. March 2011.

• "Machine Learning and Multiagent Reasoning: from robot soccer to autonomous traffic"

Virginia Tech Computer Science Department Seminar.

Blacksburg, VA. March 2011.

• "Machine Learning and Multiagent Reasoning: from robot soccer to autonomous traffic"

Baylor Physics Department Seminar.

Waco, TX. December 2010.

• "Machine Learning and Multiagent Reasoning: From Robot Soccer to Autonomous Traffic'

Johns Hopkins Computer Science Seminar.

Baltimore, MD. November 2010.

• "Autonomous Robots Playing Soccer and Traversing Intersections"

Hot Science - Cool Talks Outreach Lecture Series.

Austin, TX. October 2010.

• "Learning and Multiagent Reasoning for Autonomous Agents"

Blue Knot Austin.

Austin, TX. July 2010.

• "Intersections of the Future: Using Fully Autonomous Vehicles"

Taiwan Agent Summer School.

Hsinchu, Taiwan. June 2010.

• "Learning and Multiagent Reasoning for Autonomous Agents"

Taiwan Agent Summer School.

Hsinchu, Taiwan. June 2010.

• "Machine Learning on Physical Robots"

Seminar in Mechanical Engineering Department., UT Austin.

Austin, TX. April 2010.

• "Autonomous Robots Playing Soccer and Traversing Intersections"

TEDxUT.

Austin, TX. April 2010.

• "Progress in Artificial Intelligence: The Challenge Problem Approach"

SxSW Interactive, panel on AI 2010: Wall-e Or Rise Of The Machines?

Austin, TX. March 2010.

• "Learning and Multiagent Reasoning for Autonomous Agents"

National Instruments.

Austin, TX. January 2010.

• "Intersections of the Future: Using Fully Autonomous Vehicles"

IEEE Latin-American Summer School on Robotics.

Santiago, Chile. December 2009.

• "Learning and Multiagent Reasoning for Autonomous Agents"

IEEE Latin-American Summer School on Robotics.

Santiago, Chile. December 2009.

• "Intersections of the Future: Using Fully Autonomous Vehicles"

Robotics Seminar at University Polytecnica de Catalonya.

Barcelona, Spain. November 2009.

• "Teaching Teammates in Ad Hoc Teams"

UT Department of Computer Sciences Forum for AI and UTCS Colloquium.

Austin, TX. September 2009.

• "Learning and Multiagent Reasoning for Autonomous Agents"

Technion CS Department Pixel Club lecture.

Haifa, Israel. June 2009.

• "How Machines Learn: From Robot Soccer to Autonomous Traffic"

HEMDA Center for Science Education.

Tel Aviv, Israel. June 2009.

• "Teaching Teammates in Ad Hoc Teams"

Game theory seminar at Hebrew University Center for Rationality.

Jerusalem, Israel. May 2009.

• "Teaching Teammates in Ad Hoc Teams"

AAMAS 2009 workshop on Adaptive Learning Agents.

Budapest, Hungary. May 2009.

• "Generalization in Reinforcement Learning"

Hebrew University Machine Learning Club Talk.

Jerusalem, Israel. April 2009. 4/23/09

• "Machine Learning on Physical Robots"

Hebrew University CS Coloquium.

Jerusalem, Israel. April 2009.

• "Generalization in Reinforcement Learning"

Technion EE Guest Lecture.

Haifa, Israel, March 2009.

• "Machine Learning on Physical Robots"

Haifa Mini-Workshop on Machine Learning: Theory and Practice.

Haifa, Israel. March 2009.

• "Learning and Multiagent Reasoning for Autonomous Agents"

Ben Gurion University.

Beér Sheva, Israel. February 2009.

• "Learning and Multiagent Reasoning for Autonomous Agents"

IBM Haifa Research Lab.

Haifa, Israel. December 2009.

• "Learning and Multiagent Reasoning for Autonomous Agents"

The Israel Association for Artificial Intelligence Symposium.

Ashkelon, Israel. November 2008.

• "Learning and Multiagent Reasoning for Autonomous Agents"

International Workshop on Market-Based Control of Complex Computational Systems.

Liverpool, UK. September 2008.

• "Learning and Multiagent Reasoning for Autonomous Agents"

University of Alberta AI Seminar.

Edmonton, Alberta. March 2007.

• "Learning and Multiagent Reasoning for Autonomous Agents"

University of Southern Alabama.

Mobile, Alabama. February 2007.

• "Embracing Mobility"

DARPA kickoff meeting on Information Theory for Mobile Ad-Hoc Networks.

Chicago, IL. November 2006.

• "Layered Learning on Physical Robots"

University of Amsterdam.

Amsterdam, The Netherlands. November 2006.

• "Robust Autonomous Color Learning on a Mobile Robot"

Robotics Institute Seminar Series, Carnegie Mellon University.

Pittsburgh, PA. October 2006.

• "Robust Autonomous Color Learning on a Mobile Robot"

Center for Perceptual Systems Seminar Series, UT Austin.

Austin, TX. October 2006.

• "Machine Learning and Multiagent Systems: From robot soccer to autonomous traffic"

Lockheed Martin Aeronautics.

Fort Worth, TX. October 2006.

• "Machine Learning and Multiagent Systems: From robot soccer to autonomous traffic" IEEE MetroCon.

Arlington, TX. October 2006.

• "RoboCup: The Robot Soccer World Cup"

Department of Kinesiology and Health Education, UT Austin.

Austin, TX. September 2006.

• "State Abstraction Discovery, and Layered Learning on Physical Robots"

AAMAS workshop on Hierarchical Autonomous Agents and Multiagent Systems.

Hakodate, Japan. May 2006.

• "Machine Learning on Physical Robots"

GRASP seminar series at University of Pennsylvania.

Philadelphia, PA. March 2006.

• "Robot Learning"

National Academy of Sciences Frontiers of Science Symposium.

Irvine, CA. October 2005.

• "Reinforcement Learning for GGP"

DARPA kickoff meeting on Transfer Learning.

Palo Alto, CA. October 2005.

• "Behavior Transfer for Value-Function-Based Reinforcement Learning"

Cognition & Perception Seminar Series, Psychology Department, UT Austin.

Austin, TX. September 2005.

• "Machine Learning and Multiagent Systems: From robot soccer to autonomic computing" **IBM Technical Vitality Council.**

Austin, TX. September 2005.

• "The Trading Agent Competition: Two Champion Adaptive Bidding Agents"

Intelligent Systems Seminar Series, McCombs School of Business, UT Austin.

Austin, TX. September 2005.

• "Practical Vision-Based Monte Carlo Localization on a Legged Robot"

LJCAI Workshop on Reasoning with Uncertainty in Robotics.

Edinburgh, Scotland. July 2005.

• "RoboCup as an Introduction to Multiagent Systems and Research"

AAMAS Workshop on Teaching Multiagent Systems.

Utrecht, Netherlands. July 2005.

• "Multi-Robot Learning for Continuous Area Sweeping"

AAMAS Workshop on Learning and Adaptation in Multiagent Systems.

Utrecht, Netherlands. July 2005.

• "Speeding up Reinforcement Learning with Behavior Transfer"

RoboCup US Open.

Atlanta, GA. May 2005.

• "Scaling Up Reinforcement Learning via Task Transfer"

DARPA bidder's conference on Transfer Learning.

Washington, DC. March 2005.

• "Machine Learning on Physical Robots"

Bar Ilan University.

Ramat Gan, Israel. December 2004.

• "Adversarial Agents and Other Agent Topics"

Air Force Research Lab.

Rome, New York. July 2004.

• "Embodied Agents"

Americas' School on Agents and Multiagent Systems.

New York, New York. July 2004.

• "Coaching, Advising, and Task Transfer for Multiagent Learning"

Information Science and Technology (ISAT) study on Multiagent/Distributed Learning.

Boston, Massachusetts. June 2004.

• "Robot soccer: competitions and research"

National Instruments Scholarship for Excellence reception.

Austin, TX. June 2004.

• "The RoboCup Challenge: Progress and Research Results in Robot Soccer"

UT Mechanical Engineering Dept. Robotics Research Group.

Austin, TX. May 2004.

• "The RoboCup Challenge: Progress and Research Results in Robot Soccer"

UT Department of Computer Sciences Forum for AI.

Austin, TX. November 2003.

• "Policy Gradient Reinforcement Learning for Fast Quadrupedal Locomotion"

DARPA PI Meeting on Navigation, Locomotion, and Articulation.

Washington, DC. Nov. 2003.

• "Autonomous Learning Agents in Dynamic, Multiagent Environments"

IROS-2003 Workshop on Learning and Evolution in MAS.

Las Vegas, Nevada. October, 2003.

• "The RoboCup Challenge: Progress and Research Results in Robot Soccer"

University of Science and Technology China (USTC).

Hefei, China. October, 2003.

• "Layered Learning towards Autonomic Computing"

IJCAI-2003 Workshop on AI and Autonomic Computing.

Acapulco, Mexico. August 2003.

• "Autonomous Bidding Agents and the Power of Threats" (with Michael Littman)

CMU Machine Learning Lunch Seminar.

Pittsburgh, Pennsylvania. April 2003.

• "Machine Learning Research in the RoboCup Simulation League"

First RoboCup American Open Workshop.

Pittsburgh, Pennsylvania. April 2003.

• "Autonomous Learning Agents in Dynamic, Multiagent Environments"

UT College of Natural Sciences Advisory Council Meeting.

Austin, TX. April 2003.

• "Scaling Reinforcement Learning toward RoboCup Soccer"

NIPS 2002 workshop on Multi-Agent Learning.

Whistler, British Columbia, Canada. December 2002.

• "Autonomous Learning Agents in Dynamic, Multiagent Environments"

University of Texas School of Library and Information Science.

Austin, TX. September 2002.

• "The Trading Agent Competition: Two Champion Adaptive Bidding Agents"

UT Department of Computer Sciences Forum for AI.

Austin, TX. September 2002.

• "Multiagent Competitions and Research: Lessons from RoboCup and TAC"

Trading Agent Competition Workshop.

Edmonton, Alberta, Canada. July, 2002.

• "Autonomous Learning Agents in Dynamic, Multiagent Environments: Auctions and Soccer"

Santa Fe Institute Collective Cognition Workshop.

Santa Fe, New Mexico. January 2002.

• "Autonomous Learning Agents in Dynamic, Multiagent Environments: Auctions and Soccer" University of Alberta AI Seminar.

Edmonton, Alberta. November 2001.

• "Autonomous Bidding Agents"

Brookings Institution Workshop on Multi-Agent Computation in Natural and Artificial Economies.

Washington, DC. October 2001.

• "Layered Learning in Multi-Agent Systems: A Winning Approach to Robotic Soccer"

IEEE Computer Society and DigiPen Institute of Technology.

Seattle, Washington. July 2001.

• "Layered Learning in Multi-Agent Systems: A Winning Approach to Robotic Soccer"

The Boeing Company.

Seattle, Washington. July 2001.

• "Layered Learning in Multi-Agent Systems"

Multi-Strategy Learning Workshop.

Guimarães, Portugal. June 2000.

• "The RoboCup Challenge"

NASA Goddard Space Flight Center.

Greenbelt, Maryland. March 2000.

• "The RoboCup Challenge"

Ohio University.

Athens, Ohio. February 2000.

• "The RoboCup Challenge"

IEEE Real-Time Systems Symposium.

Phoenix, Arizona. December, 1999.

• "Layered Learning in Multi-Agent Systems"

Machines That Learn Workshop.

Snowbird, Utah. April 1998.

• "Layered Learning in Multi-Agent Systems"

SRI International.

Palo Alto, California. November 1997.

• "Layered Learning in Multi-Agent Systems"

University of Washington.

Seattle, Washington. November 1997.

• "Task Decomposition and Dynamic Role Assignment for Real-Time Strategic Teamwork"

Electrotechnical Laboratory (ETL).

Tsukuba, Japan. August 1997.

• "Machine Learning for Agent Control in Real-time Multi-Agent Domains"

Workshop on Intelligent Robotic Agents.

Porto Alegre, Brazil. March 1997.

• "Layered Learning in the Soccer Server"

Electrotechnical Laboratory (ETL).

Tsukuba, Japan. November 1996.

• "Layered Learning in the RoboCup Soccer Server"

Osaka University.

Osaka, Japan. November 1996.

• "Building a Dedicated Robotic Soccer System"

Korean Advanced Institute of Science and Technology (KAIST).

Taejon, Korea. August 1996.

• "Towards Collaborative and Adversarial Learning: A Case Study in Robotic Soccer"

Naval Research Labs (NRL).

Washington DC. July 1996.

• "FLECS: Planning with a Flexible Commitment Strategy"

NASA Jet Propulsion Laboratory (JPL).

Pasadena, California. July 1995.

• "FLECS: Planning with a Flexible Commitment Strategy"

USC Intelligent Software Institute (ISI).

Marina Del Rey, California. July 1995.

PUBLICATIONS

All listed publications are available and cross-listed by *type*, *date*, *topic*, and *co-author* at http://www.cs.utexas.edu/users/pstone/papers.html

Books

- 1. Xiaoping Chen, **Peter Stone**, Luis Enrique Sucar, and Tijn van der Zant, editors. *RoboCup-2012: Robot Soccer World Cup XVI*, volume 7500 of *Lecture Notes in Artificial Intelligence*. Springer Verlag, Berlin, 2013.
- Kagan Tumer, Pinar Yolum, Liz Sonenberg, and Peter Stone, editors. Proceedings of the Tenth International Conference on Autonomous Agents and Multiagent Systems. International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS), May 2011.
- 3. Michael P. Wellman, Amy Greenwald, and **Peter Stone**. *Autonomous Bidding Agents: Strategies and Lessons from the Trading Agent Competition*. MIT Press, 2007. (monograph)
- 4. **Peter Stone**. *Intelligent Autonomous Robotics: A Robot Soccer Case Study*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers, 2007. (monograph)
- 5. **Peter Stone** and Gerhard Weiss, editors. *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems*. Association for Computing Machinery (ACM), May 2006.
- 6. **Peter Stone**, Tucker Balch, and Gerhard Kraetzschmar, editors. *RoboCup-2000: Robot Soccer World Cup IV*. volume 2019 of Lecture Notes in Artificial Intelligence. Springer Verlag, Berlin, 2001.
- 7. **Peter Stone**. Layered Learning in Multiagent Systems: A Winning Approach to Robotic Soccer. MIT Press, 2000. (monograph)

Journal Articles

8. Xiaohan Zhang, Yan Ding, Yohei Hayamizu, Zainab Altaweel, Yifeng Zhu, Yuke Zhu, **Peter Stone**, Chris Paxton, and Shiqi Zhang. LLM-GROP: Visually grounded robot task and motion planning with large language models. *The International Journal of Robotics Research* (IJRR), 2025.

- 9. Rohan Chandra, Vrushabh Zinage, Efstathios Bakolas, **Peter Stone**, and Joydeep Biswas. Deadlockfree, safe, and decentralized multi-robot navigation in social mini-games via discrete-time control barrier functions. *Autonomous Robots*, 2025.
- 10. Hojoon Lee, Takuma Seno, Jun Jet Tai, Kaushik Subramanian, Kenta Kawamoto, **Peter Stone**, and Peter R. Wurman. A champion-level vision-based reinforcement learning agent for competitive racing in gran turismo 7. *IEEE Robotics and Automation Letters*, 10(6):5545–52, June 2025.
- 11. Luisa Mao, Garrett Warnell, **Peter Stone**, and Joydeep Biswas. PACER: Preference-conditioned all-terrain costmap generation. *Robotics and Automation Letters*, 2025.
- 12. Chen Tang, Ben Abbatematteo, Jiaheng Hu, Rohan Chandra, Roberto Martn-Martn, and **Peter Stone**. Deep reinforcement learning for robotics: A survey of real-world successes. *Annual Review of Control, Robotics, and Autonomous Systems (ARCRAS)*, 8:153–88, 2024.
- 13. Anthony Francis, Claudia Perez-D'Arpino, Chengshu Li, Fei Xia, Alexandre Alahi, Rachid Alami, Aniket Bera, Abhijat Biswas, Joydeep Biswas, Rohan Chandra, Hao-Tien Lewis Chiang, Michael Everett, Sehoon Ha, Justin Hart, Jonathan P. How, Haresh Karnan, Tsang-Wei Edward Lee, Luis J. Manso, Reuth Mirksy, Soren Pirk, Phani Teja Singamaneni, **Peter Stone**, Ada V. Taylor, Peter Trautman, Nathan Tsoi, Marynel Vazquez, Xuesu Xiao, Peng Xu, Naoki Yokoyama, Alexander Toshev, and Roberto Martin-Martin. Principles and guidelines for evaluating social robot navigation algorithms. *ACM Transactions on Human-Robot Interaction (THRI)*, 14(2), February 2025.
- 14. Josiah P. Hanna, Yash Chandak, Philip S. Thomas, Martha White, **Peter Stone**, and Scott Niekum. Data-efficient policy evaluation through behavior policy search. *Journal of Machine Learning Research*, 2024.
- 15. Alessandra Rossi, Maike Paetzel-Prüsmann, Merel Keijsers, Michael Anderson, Susan Leigh Anderson, Daniel Barry, Jan Gutsche, Justin Hart, Luca Iocchi, Ainse Kokkelmans, Wouter Kuijpers, Yun Liu, Daniel Polani, Caleb Roscon, Marcus Scheunemann, Peter Stone, Florian Vahl, René van de Molengraft, and Oskar von Stryk. The human in the loop: Perspectives and challenges for RoboCup 2050. Autonomous Robots, May 2024.
- 16. Shiqi Zhang, Piyush Khandelwal, and **Peter Stone**. iCORPP: interleaved commonsense reasoning and probabilistic planning on robots. *Robotics and Autonomous Systems*, 174, April 2024.
- 17. Andrea Soltoggio, Eseoghene Ben-Iwhiwhu, Vladimir Braverman, Eric Eaton, Benjamin Epstein, Yunhao Ge, Lucy Halperin, Jonathan How, Laurent Itti, Michael A. Jacobs, Pavan Kantharaju, Long Le, Steven Lee, Xinran Liu, Sildomar T. Monteiro, David Musliner, Saptarshi Nath, Priyadarshini Panda, Christos Peridis, Hamed Pirsiavash, Vishwa Parekh, Kaushik Roy, Shahaf Shperberg, Hava T. Siegelmann, **Peter Stone**, Kyle Vedder, Jingfeng Wu, Lin Yang, Guangyao Zheng, and Soheil Kolouri. A collective AI via lifelong learning and sharing at the edge. *Nature Machine Intelligence*, 6, pages 251–64, March 2024.
- 18. Reuth Mirsky, Xuesu Xiao, Justin Hart, and **Peter Stone**. Conflict avoidance in social navigation a survey. *ACM Transactions on Human-Robot Interaction* (THRI), 13:1, pages 1–36, March 2024.
- 19. W. Bradley Knox, Stephane Hatgis-Kessell, Serena Booth, Scott Niekum, **Peter Stone**, and Alessandro Allievi. Models of human preference for learning reward functions. *Transactions on Machine Learning Research* (TMLR), 2024.
- 20. Yulin Zhang, William Macke, Jiaxun Cui, Sharon Hornstein, Daniel Urieli, and **Peter Stone**. Learning a robust multiagent driving policy for traffic congestion reduction. *Neural Computing and Applications*, November 2023.
- 21. Ori Ossmy, Danyang Han, Patrick MacAlpine, Justine Hoch, **Peter Stone**, and Karen E. Adolph. Walking and falling: Using robot simulations to model the role of errors in infant walking. *Developmental Science*, 27:2, e13449, September 2023.

- 22. Peter R. Wurman, **Peter Stone**, and Michael Spranger. Improving artificial intelligence with games. *Science*, 381:147–8, July 2023.
- 23. Keya Ghonasgi, Reuth Mirsky, Nisha Bhargava, Adrian M Haith, **Peter Stone**, and Ashish D Deshpande. Kinematic coordinations capture learning during human-exoskeleton interaction. *Scientific Reports*, 13:10322, June 2023.
- 24. Varun Kompella, Thomas Walsh, Samuel Barrett, Peter Wurman, and **Peter Stone**. Event tables for efficient experience replay. *Transactions on Machine Learning Research* (TMLR), 2023.
- 25. Xiaohan Zhang, Saeid Amiri, Jivko Sinapov, Jesse Thomason, **Peter Stone**, and Shiqi Zhang. Multimodal embodied attribute learning by robots for object-centric action policies. *Autonomous Robots*, March 2023.
- 26. Megan M. Baker, Alexander New, Mario Aguilar-Simon, Ziad Al-Halah, Sébastien M. R. Arnold, Ese Ben-Iwhiwhu, Andrew P. Brna, Ethan Brooks, Ryan C. Brown, Zachary Daniels, Anurag Daram, Fabien Delattre, Ryan Dellana, Eric Eaton, Haotian Fu, Kristen Grauman, Jesse Hostetler, Shariq Iqbal, Cassandra Kent, Nicholas Ketz, Soheil Kolouri, George Konidaris, Erik Learned-Miller Dhireesha Kudithipudi, Seungwon Lee, Michael L. Littman, Jorge A. Mendez Sandeep Madireddy, Eric Q. Nguyen, Christine D. Piatko, Praveen K. Pilly, Aswin Raghavan, Abrar Rahman, Santhosh Kumar Ramakrishnan, Neale Ratzlaff, Andrea Soltoggio, Peter Stone, Indranil Sur, Zhipeng Tang, Saket Tiwari, Kyle Vedder, Felix Wang, Zifan Xu, Angel Yanguas-Gil, Harel Yedidsion, Shangqun Yu, and Gautam K. Vallabha. A domain-agnostic approach for characterization of lifelong learning systems. Neural Networks, pages 274–96, March 2023.
- 27. W. Bradley Knox, Alessandro Allievi, Holger Banzhaf, Felix Schmitt, and **Peter Stone**. Reward (mis)design for autonomous driving. *Artificial Intelligence* (AIJ), 316:103829, 2023.
- 28. Haresh Karnan, Anirudh Nair, Xuesu Xiao, Garrett Warnell, Soren Pirk, Alexander Toshev, Justin Hart, Joydeep Biswas, and **Peter Stone**. Socially compliant navigation dataset (SCAND): A large-scale dataset of demonstrations for social navigation. *Robotics and Automation Letters* (RA-L), 7:11807–14, presented at *International Conference on Intelligent Robots and Systems* (IROS), October 2022.
- 29. Kingsley Nweye, Bo Liu, Nagy Zoltan, and **Peter Stone**. Real-world challenges for multi-agent reinforcement learning in grid-interactive buildings. *Journal of Energy and AI*, September 2022.
- 30. Xuesu Xiao, Zizhao Wang, Zifan Xu, Bo Liu, abd Gauraang Dhamankar, Anirudh Nair, Garrett Warnell, and **Peter Stone**. APPL: Adaptive planner parameter learning. *Robotics and Autonomous Systems* (RAS), May 2022.
- 31. Yifeng Zhu, **Peter Stone**, and Yuke Zhu. Bottom-up skill discovery from unsegmented demonstrations for long-horizon robot manipulation. *IEEE Robotics and Automation Letters* (RA-L), 7:4126–33, April 2022.
- 32. Xuesu Xiao, Bo Liu, Garrett Warnell, and **Peter Stone**. Motion planning and control for mobile robot navigation using machine learning: a survey. *Autonomous Robots*, 46(5):569–97, March 2022.
- 33. Peter R. Wurman, Samuel Barrett, Kenta Kawamoto, James MacGlashan, Kaushik Subramanian, Thomas J. Walsh, Roberto Capobianco, Alisa Devlic, Franziska Eckert, Florian Fuchs, Leilani Gilpin, Varun Kompella, Piyush Khandelwal, HaoChih Lin, Patrick MacAlpine, Declan Oller, Craig Sherstan, Takuma Seno, Michael D. Thomure, Houmehr Aghabozorgi, Leon Barrett, Rory Douglas, Dion Whitehead, Peter Duerr, **Peter Stone**, Michael Spranger, , and Hiroaki Kitano. Outracing champion gran turismo drivers with deep reinforcement learning. *Nature*, 62:223–28, Feb. 2022.
- 34. Zizhao Wang, Xuesu Xiao, Bo Liu, Garrett Warnell, and **Peter Stone**. Apple: Adaptive planner parameter learning from evaluative feedback. *IEEE Robotics and Automation Letters* (RA-L), presented at International Conference on Intelligent Robots and Systems (IROS), October 2021.

- 35. Xuesu Xiao, Joydeep Biswas, and **Peter Stone**. Learning inverse kinodynamics for accurate high-speed off-road navigation on unstructured terrain. *IEEE Robotics and Automation Letters* (RA-L), presented at International Conference on Robotics and Automation(ICRA), July 2021.
- 36. Ruohan Zhang, Faraz Torabi, Garrett Warnell, and **Peter Stone**. Recent advances in leveraging human guidance for sequential decision-making tasks. *Autonomous Agents and Multi-Agent Systems* (JAA-MAS), 35:31, June 2021.
- 37. Josiah P. Hanna, Siddharth Desai, Haresh Karnan, Garrett Warnell, and **Peter Stone**. Grounded action transformation for sim-to-real reinforcement learning. *Machine Learning* (MLJ), May 2021. Special Issue on Reinforcement Learning for Real Life.
- 38. Roberto Capobianco, Varun Kompella, James Ault, Guni Sharon, Stacy Jong, Spencer Fox, Lauren Meyers, Peter R. Wurman, and **Peter Stone**. Agent-based markov modeling for improved COVID-19 mitigation policies. *The Journal of Artificial Intelligence Research* (JAIR), 71:953–92, August 2021. Based in part on earlier version in *Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2021.
- 39. Josiah P. Hanna, Scott Niekum, and **Peter Stone**. Importance sampling in reinforcement learning with an estimated behavior policy. *Machine Learning* (MLJ), 110:1267–1317, May 2021.
- 40. Yunshu Du, Garrett Warnell, Assefaw Gebremedhin, **Peter Stone**, and Matthew E. Taylor. Lucid dreaming for experience replay: Refreshing past states with the current policy. *Neural Computing and Applications*, May 2021.
- 41. Bo Liu, Xuesu Xiao, and **Peter Stone**. A lifelong learning approach to mobile robot navigation. *IEEE Robotics and Automation Letters* (RA-L), presented at International Conference on Robotics and Automation (ICRA), 6(2), April 2021.
- 42. Alec Koppel, Garrett Warnell, Ethan Stump, **Peter Stone**, and Alejandro Ribeiro. Policy evaluation in continuous MDPs with efficient kernelized gradient temporal difference. *IEEE Transactions on Automatic Control*, April 2021.
- 43. Michael Albert, Vincent Conitzer, Giuseppe Lopomo, and **Peter Stone**. Mechanism design for correlated valuations: efficient methods for revenue maximization. *Operations Research*, March 2021.
- 44. Xuesu Xiao, Bo Liu, Garrett Warnell, and **Peter Stone**. Toward agile maneuvers in highly constrained spaces: Learning from hallucination. *IEEE Robotics and Automation Letters* (RA-L), January 2021.
- 45. Shih-Yun Lo, Shiqi Zhang, and **Peter Stone**. The PETLON algorithm to plan efficiently for task-level-optimal navigation. *Journal of Artificial Intelligence Research* (JAIR), 69:471–520, October 2020. Based on earlier version in *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, July 2018. **Best Robotics Paper Award.**
- 46. Sanmit Narvekar, Bei Peng, Matteo Leonetti, Jivko Sinapov, Matthew E. Taylor, and **Peter Stone**. Curriculum learning for reinforcement learning domains: A framework and survey. *Journal of Machine Learning Research* (JMLR), 21(181):1–50, 2020.
- 47. Brahma Pavse, Faraz Torabi, Josiah Hanna, Garrett Warnell, and **Peter Stone**. RIDM: Reinforced inverse dynamics modeling for learning from a single observed demonstration. *IEEE Robotics and Automation Letters* (RA-L), presented at International Conference on Intelligent Robots and Systems (IROS), 5:6262–69, October 2020.
- 48. Xuesu Xiao, Bo Liu, Garrett Warnell, Jonathan Fink, and **Peter Stone**. APPLD: Adaptive planner parameter learning from demonstration. *IEEE Robotics and Automation Letters* (RA-L), presented at International Conference on Intelligent Robots and Systems (IROS), June 2020.
- 49. Felipe Leno Da Silva, Garrett Warnell, Anna Helena Reali Costa, and **Peter Stone**. Agents teaching agents: a survey on inter-agent transfer learning. *Autonomous Agents and Multi-Agent Systems* (JAAMAS), Jan 2020.

- 50. **Peter Stone**. A broader, more inclusive definition of AI. *Journal of Artificial General Intelligence* (JAGI), 11(2):63–65, 2020.
- 51. Jesse Thomason, Aishwarya Padmakumar, Jivko Sinapov, Nick Walker, Yuqian Jiang, Harel Yedidsion, Justin Hart, **Peter Stone**, and Raymond J. Mooney. Jointly improving parsing and perception for natural language commands through human-robot dialog. *The Journal of Artificial Intelligence Research* (JAIR), 67, February 2020.
- 52. Elad Liebman, Maytal Saar-Tsechansky, and **Peter Stone**. The Right Music at the Right Time: Adaptive Personalized Playlists Based on Sequence Modeling. *MIS Quarterly*, 43:3, pages 765–86, 2019.
- 53. Yuqian Jiang, Shi-qi Zhang, Piyush Khandelwal, and **Peter Stone**. Task planning in robotics: an empirical comparison of PDDL- and ASP-based systems. *Frontiers of Information Technology & Electronic Engineering*, April 2019.
- 54. Yuqian Jiang, Harel Yedidsion, Shiqi Zhang, Guni Sharon, and **Peter Stone**. Multi-robot planning with conflicts and synergies. *Autonomous Robots*, March 2019.
- 55. Ori Ossmy, Justine E. Hoch, Patrick MacAlpine, Shohan Hasan, **Peter Stone**, and Karen E. Adolph. Variety wins: Soccer-playing robots and infant walking. *Frontiers in Neurorobotics*, 12:19, 2018.
- 56. Decebal Constantin Mocanu, Elena Mocanu, **Peter Stone**, Phuong H. Nguyen, Madeleine Gibescu, and Antonio Liotta. Scalable training of artificial neural networks with adaptive sparse connectivity inspired by network science. *Nature Communications*, 9(2383), June 2018.
- 57. Stefano Albrecht and **Peter Stone**. Autonomous agents modelling other agents: A comprehensive survey and open problems. *Artificial Intelligence* (AIJ), 258:66–95, 2018.
- 58. Tarun Rambha, Stephen D. Boyles, Avinash Unnikrishnan, and **Peter Stone**. Marginal cost pricing for system optimal traffic assignment with recourse under supply-side uncertainty. *Transportation Research Part B: Methodological*, 110:104–21, 2018.
- 59. Patrick MacAlpine and **Peter Stone**. Overlapping layered learning. *Artificial Intelligence* (AIJ), 254:21–43, January 2018.
- 60. Guni Sharon, Michael W. Levin, Josiah P. Hanna, Tarun Rambha, Stephen D. Boyles, and **Peter Stone**. Network-wide adaptive tolling for connected and automated vehicles. *Transportation Research Part C*, 84:142–157, September 2017.
- 61. Corey White, Elad Liebman, and **Peter Stone**. Decision mechanisms underlying mood-congruent emotional classification. *Cognition and Emotion*, pages 1–10, 2017.
- 62. Piyush Khandelwal, Shiqi Zhang, Jivko Sinapov, Matteo Leonetti, Jesse Thomason, Fangkai Yang, Ilaria Gori, Maxwell Svetlik, Priyanka Khante, Vladimir Lifschitz, J. K. Aggarwal, Raymond Mooney, and **Peter Stone**. BWIbots: A platform for bridging the gap between ai and human–robot interaction research. *The International Journal of Robotics Research* (IJRR), 2017.
- 63. Stefano Albrecht, Somchaya Liemhetcharat, and **Peter Stone**. Special issue on multiagent interaction without prior coordination: Guest editorial. *Autonomous Agents and Multi-Agent Systems* (JAAMAS), 31:4, pages 765–66, July 2017.
- 64. Katie Genter, Tim Laue, and **Peter Stone**. Three years of the robocup standard platform league dropin player competition: Creating and maintaining a large scale ad hoc teamwork robotics competition. *Autonomous Agents and Multi-Agent Systems* JAAMAS), pages 31:4, pages 790–820, July 2017.
- 65. Matteo Leonetti, Luca Iocchi, and **Peter Stone**. A synthesis of automated planning and reinforcement learning for efficient, robust decision-making. *Artificial Intelligence* (AIJ), 241:103 130, September 2016.

- 66. Samuel Barrett, Avi Rosenfeld, Sarit Kraus, and **Peter Stone**. Making friends on the fly: Cooperating with new teammates. *Artificial Intelligence* (AIJ), October 2016.
- 67. Matthew Hausknecht, Wen-Ke Li, Michael Mauk, and **Peter Stone**. Machine learning capabilities of a simulated cerebellum. *IEEE Transactions on Neural Networks and Learning Systems*, Jan 2016.
- 68. Katie Genter and **Peter Stone**. Ad hoc teamwork behaviors for influencing a flock. *Acta Polytechnica*, 56(1), 2016.
- 69. Elad Liebman, Benny Chor, and **Peter Stone**. Representative selection in nonmetric datasets. *Applied Artificial Intelligence*, 29:807–838, 2015.
- 70. Todd Hester and **Peter Stone**. Intrinsically motivated model learning for developing curious robots. *Artificial Intelligence* (AIJ), May 2015.
- 71. W. Bradley Knox and **Peter Stone**. Framing reinforcement learning from human reward: Reward positivity, temporal discounting, episodicity, and performance. *Artificial Intelligence* (AIJ), 225, August 2015.
- 72. Matthew Hausknecht, Joel Lehman, Risto Miikkulainen, and **Peter Stone**. A neuroevolution approach to general atari game playing. In *IEEE Transactions on Computational Intelligence and AI in Games*, 2013.
- 73. **Peter Stone**, Gal A. Kaminka, Sarit Kraus, Jeffrey R. Rosenschein, and Noa Agmon. Teaching and leading an ad hoc teammate: Collaboration without pre-coordination. *Artificial Intelligence* (AIJ), 203:35–65, October 2013.
- 74. Doran Chakraborty and **Peter Stone**. Multiagent learning in the presence of memory-bounded agents. *Autonomous Agents and Multiagent Systems* (JAAMAS), 2013.
- 75. Todd Hester and **Peter Stone**. TEXPLORE: Real-time sample-efficient reinforcement learning for robots. *Machine Learning* (MLJ), 90(3), 2013.
- 76. Wen-Ke Li, Matthew J. Hausknecht, **Peter Stone**, and Michael D. Mauk. Using a million cell simulation of the cerebellum: Network scaling and task generality. *Neural Networks*, November 2012.
- 77. W. Bradley Knox, Brian D. Glass, Bradley C. Love, W. Todd Maddox, and **Peter Stone**. How humans teach agents: A new experimental perspective. *International Journal of Social Robotics* (IJSR), 4:409–421, October 2012.
- 78. W. Bradley Knox, A. Ross Otto, **Peter Stone**, and Bradley Love. The nature of belief-directed exploratory choice in human decision-making. *Frontiers in Psychology*, 2(398), January 2012.
- 79. Shivaram Kalyanakrishnan and **Peter Stone**. Characterizing reinforcement learning methods through parameterized learning problems. *Machine Learning* (MLJ), 84(1):205–247, June 2011.
- 80. David Fajardo, Tsz-Chiu Au, Travis Waller, **Peter Stone**, and David Yang. Automated intersection control: Performance of a future innovation versus current traffic signal control. *Transportation Research Record* (TRR), 2259:223–32, 2012.
- 81. Tobias Jung, Daniel Polani, and **Peter Stone**. Empowerment for continuous agent-environment systems. *Adaptive Behavior*, 19(1):16–39, 2011.
- 82. David Pardoe, **Peter Stone**, Maytal Saar-Tsechansky, Tayfun Keskin, and Kerem Tomak. Adaptive auction mechanism design and the incorporation of prior knowledge. *INFORMS Journal on Computing*, 22(3):353–370, Summer 2010.
- 83. Jonathan Wildstrom, **Peter Stone**, and Emmett Witchel. Autonomous return on investment analysis of additional processing resources. *International Journal on Autonomic Computing*, 1(3), 2010.

- 84. Matthew E. Taylor and **Peter Stone**. Transfer learning for reinforcement learning domains: A survey. *Journal of Machine Learning Research* (JMLR), 10(1):1633–1685, 2009.
- 85. Shimon Whiteson, Matthew E. Taylor, and **Peter Stone**. Critical factors in the empirical performance of temporal difference and evolutionary methods for reinforcement learning. *Journal of Autonomous Agents and Multi-Agent Systems* (JAAMAS), 21(1):1–27, 2010.
- 86. Mohan Sridharan and **Peter Stone**. Color learning and illumination invariance on mobile robots: A survey. *Robotics and Autonomous Systems (RAS) Journal*, 57(6-7):629–44, June 2009.
- 87. Juhyun Lee, W. Bradley Knox, and **Peter Stone**. Inter-classifier feedback for human-robot interaction in a domestic setting. *Journal of Physical Agents* (JoPhA), 2(2):41–50, July 2008. Special Issue on Human Interaction with Domestic Robots.
- 88. Patrick Beeson, Jack O'Quin, Bartley Gillan, Tarun Nimmagadda, Mickey Ristroph, David Li, and **Peter Stone**. Multiagent interactions in urban driving. *Journal of Physical Agents* (JoPhA), 2(1):15–30, March 2008. Special issue on Multi-Robot Systems.
- 89. Kurt Dresner and **Peter Stone**. A multiagent approach to autonomous intersection management. *Journal of Artificial Intelligence Research* (JAIR), 31:591–656, March 2008.
- 90. Daniel Stronger and **Peter Stone**. Polynomial regression with automated degree: A function approximator for autonomous agents. *International Journal on Artificial Intelligence Tools*, 17(1):159–174, 2008
 - Based on earlier version in *The 18th IEEE International Conference on Tools with Artificial Intelligence*, November 2006.
 - Nominee for Best Paper Award.
- 91. Matthew E. Taylor, **Peter Stone**, and Yaxin Liu. Transfer learning via inter-task mappings for temporal difference learning. *Journal of Machine Learning Research* (JMLR), 8(1):2125–2167, 2007.
- 92. Mohan Sridharan and **Peter Stone**. Planning actions to enable color learning on a mobile robot. *International Journal of Information and Systems Sciences*, 3(3):510–25, 2007.
- 93. Mohan Sridharan and **Peter Stone**. Structure-Based Color Learning on a Mobile Robot under Changing Illumination. *Autonomous Robots*, 23(3):161–182, 2007.
- 94. **Peter Stone**. Multiagent learning is not the answer it is the question. *Artificial Intelligence* (AIJ), 171:402–405, 2007.
- 95. Shimon Whiteson, Matthew E. Taylor, and **Peter Stone**. Empirical studies in action selection for reinforcement learning. *Adaptive Behavior*, 15(1):33–50, 2007.
- 96. **Peter Stone**, Mohan Sridharan, Daniel Stronger, Gregory Kuhlmann, Nate Kohl, Peggy Fidelman, and Nicholas K. Jong. From pixels to multi-robot decision-making: A study in uncertainty. *Robotics and Autonomous Systems* (RAS), 54(11):933–43, November 2006. Special issue on Planning Under Uncertainty in Robotics.
- 97. Daniel Stronger and **Peter Stone**. Towards autonomous sensor and actuator model induction on a mobile robot. *Connection Science Journal* (CSJ), 18(2):97–119, June 2006. Special Issue on Developmental Robotics.
 - Based on "Simultaneous calibration of action and sensor models on a mobile robot." In *IEEE International Conference on Robotics and Automation* (ICRA), April 2005.
- 98. Shimon Whiteson and **Peter Stone**. Evolutionary Function Approximation for Reinforcement Learning. *Journal of Machine Learning Research* (JMLR), 7:877–917, May 2006.

- 99. Charles Lee Isbell Jr., Michael Kearns, Dave Kormann, Satinder Singh, and **Peter Stone**. Cobot in LambdaMOO: an adaptive social statistics agent. *Journal of Autonomous Agents and Multi-Agent Systems* (JAAMAS), 13(3):327–354, November, 2006.

 Based on "Cobot in LambdaMOO: A social statistics agent." In *Proceedings of the Seventeenth National Conference on Artificial Intelligence* (AAAI), pages 36–41, 2001.
- 100. Peter Stone, Richard S. Sutton, and Gregory Kuhlmann. Reinforcement learning for RoboCup-soccer keepaway. Adaptive Behavior (AB), 13(2):165–188, 2005.
 Based on "Scaling reinforcement learning toward RoboCup soccer." In Proceedings of the Eighteenth International Conference on Machine Learning (ICML), 2001.
- 101. Shimon Whiteson, Nate Kohl, Risto Miikkulainen, and **Peter Stone**. Evolving keepaway soccer players through task decomposition. *Machine Learning* (MLJ), 59(1):5–30, May 2005. Based on earlier version in *Proceedings of the Genetic and Evolutionary Computation Conference* (GECCO), July 2003.
- 102. Michael Littman and Peter Stone. A Polynomial-time Nash Equilibrium Algorithm for Repeated Games. Decision Support Systems (DSS), 39:55–66, 2005.
 Based on earlier version in Proceedings of the fourth annual ACM Conference on Electronic Commerce (EC), June 2003.
- 103. **Peter Stone** and Amy Greenwald. The first international trading agent competition: Autonomous bidding agents. *Electronic Commerce Research* (EC), 5(2):229–65, April 2005.
- 104. Shimon Whiteson and **Peter Stone**. Towards autonomic computing: Adaptive job routing and scheduling. *Engineering Applications of Artificial Intelligence special issue on Autonomic Computing and Automation* (EAAI), 17(7):855–69, October 2004.

 Based on earlier version in *Proceedings of the Sixteenth Innovative Applications of AI Conference* (IAAI), San Jose, CA, July 2004.
- 105. Elizabeth Sklar, Simon Parsons, and **Peter Stone**. Using RoboCup in university-level computer science education. *Journal of Educational Resources in Computing* (JERIC), 4:2, June 2004. Special Issue on Robotics in Undergraduate Education, Part 1.
- 106. **Peter Stone**, Robert E. Schapire, Michael L. Littman, János A. Csirik, and David McAllester. Decision-theoretic bidding based on learned density models in simultaneous, interacting auctions. *Journal of Artificial Intelligence Research* (JAIR), 19:209–242, September 2003.
- 107. Itsuki Noda and **Peter Stone**. The RoboCup soccer server and CMUnited clients: Implemented infrastructure for MAS research. *Autonomous Agents and Multi-Agent Systems* (JAAMAS), 7(1–2):101–120, July–September 2003.
- 108. Michael P. Wellman, Amy Greenwald, **Peter Stone**, and Peter R. Wurman. The 2001 trading agent competition. *Electronic Markets* (EM), 13(1):4–12, May 2003.
- 109. **Peter Stone**, Michael L. Littman, Satinder Singh, and Michael Kearns. ATTac-2000: An adaptive autonomous bidding agent. *Journal of Artificial Intelligence Research* (JAIR), 15:189–206, June 2001. Based on earlier version in *Proceedings of the Fifth International Conference on Autonomous Agents*, 2001.
- 110. Peter Stone and Manuela Veloso. Multiagent systems: A survey from a machine learning perspective. Autonomous Robots, 8(3):345–383, July 2000.
 Also in Tucker Balch and Lynne E. Parker, editors, Robot Teams: From Diversity to Polymorphism. AK Peters Ltd, 2002.
- 111. Michael Bowling, Manuela Veloso, and **Peter Stone**. The CMUnited-98 champion small-robot team. *Advanced Robotics*, 13(8):753–766, 2000.

- 112. Manuela Veloso, **Peter Stone**, and Kwun Han. The CMUnited-97 robotic soccer team: Perception and multi-agent control. *Robotics and Autonomous Systems* (RAS), 29(2-3):133–143, 2000. Based on earlier version in *Proceedings of the Second International Conference on Autonomous Agents*, May 1998.
- 113. **Peter Stone** and Manuela Veloso. Task decomposition, dynamic role assignment, and low-bandwidth communication for real-time strategic teamwork. *Artificial Intelligence* (AIJ), 110(2):241–273, June 1999.
- 114. **Peter Stone** and Manuela Veloso. A layered approach to learning client behaviors in the RoboCup soccer server. *Applied Artificial Intelligence* (AAI), 12:165–188, 1998.
- 115. Minoru Asada, Yasuo Kuniyoshi, Alexis Drogoul, Hajime Asama, Maja Mataric, Dominique Duhaut, **Peter Stone**, and Hiroaki Kitano. The RoboCup physical agent challenge: Phase-I. *Applied Artificial Intelligence* (AAI), 12.2:251–263, March 1998.
- 116. **Peter Stone** and Manuela Veloso. Towards collaborative and adversarial learning: A case study in robotic soccer. *International Journal of Human-Computer Studies* (IJHCS), 48(1):83–104, January 1998.
- 117. Manuela Veloso and **Peter Stone**. FLECS: Planning with a flexible commitment strategy. *Journal of Artificial Intelligence Research* (JAIR), 3:25–52, June 1995.

Book Chapters

- 118. Sai Kiran Narayanaswami, Mauricio Tec, Ishan Durugkar, Siddharth Desai, Bharath Masetty, Sanmit Narvekar, and **Peter Stone**. Towards a real-time, low-resource, end-to-end object detection pipeline for robot soccer. In Amy Eguchi, Nuno Lau, Maike Paetzel-Prussman, and Thanapat Wanichanon, editors, *RoboCup 2022: Robot World Cup XXV*, pages 62–74. Springer International Publishing, 2023.
- 119. Reuth Mirsky, Ignacio Carlucho, Arrasy Rahman, Eliott Fosong, William Macke, Mohan Sridharan, **Peter Stone**, and Stefano Albrecht. A survey of ad hoc teamwork research. In Dorothea Baumeister and Jörg Rothe, editors, *Multi-Agent Systems*, pages 275–93. Springer International Publishing, 2022.
- 120. Patrick MacAlpine, Bo Liu, William Macke, Caroline Wang, and **Peter Stone**. UT Austin Villa: RoboCup 2021 3D simulation league competition champions. In Rachid Alami, Joydeep Biswas, Maya Cakmak, and Oliver Obst editors, *RoboCup 2021: Robot World Cup XXIV*, Lecture Notes in Artificial Intelligence, pages 314–26. Springer, 2022.
- 121. Patrick MacAlpine, Faraz Torabi, Brahma Pavse, and **Peter Stone**. UT Austin Villa: RoboCup 2019 3D simulation league competition and technical challenge champions. In Stephan Chalup, Tim Niemueller, Jackrit Suthakorn, and Mary-Anne Williams, editors, *RoboCup 2019: Robot World Cup XXIII*, Lecture Notes in Artificial Intelligence, pages 540–52. Springer, 2019.
- 122. Patrick MacAlpine, Faraz Torabi, Brahma Pavse, John Sigmon, and **Peter Stone**. UT Austin Villa: RoboCup 2018 3D simulation league champions. In Dirk Holz, Katie Genter, Maarouf Saad, and Oskar von Stryk, editors, *RoboCup 2018: Robot Soccer World Cup XXII*, Lecture Notes in Artificial Intelligence, pages 462–75. Springer, 2019.
- 123. Karl Tuyls and **Peter Stone**. Multiagent learning paradigms. In F. Belardinelli and E. Argente, editors, *Multi-Agent Systems and Agreement Technologies*, Volume 10767 of *Lecture Notes in Artificial Intelligence*, pages 1–19. Springer, 2018.
- 124. Patrick MacAlpine and **Peter Stone**. UT Austin Villa: RoboCup 2017 3D simulation league competition and technical challenges champions. In Claude Sammut, Oliver Obst, Flavio Tonidandel, and Hidehisa Akyama, editors, *RoboCup 2017: Robot Soccer World Cup XXI*, pages 473–85. Lecture Notes in Artificial Intelligence. Springer, 2018.

125. Jacob Menashe, Josh Kelle, Katie Genter, Josiah Hanna, Elad Liebman, Sanmit Narvekar, Ruohan Zhang and **Peter Stone**. Fast and precise black and white ball detection for RoboCup soccer. In Hidehisa Akyama, Oliver Obst, Claude Sammut, and Flavio Tonidandel, editors, *RoboCup 2017: Robot Soccer World Cup XXI*, pages 45–59. Lecture Notes in Artificial Intelligence. Springer Verlag, Berlin, 2018

Nominee for Best Paper Award.

- 126. Patrick MacAlpine and **Peter Stone**. Evaluating ad hoc teamwork performance in drop-in player challenges. In Gita Sukthankar and Juan A. Rodriguez-Aguilar, editors, *Autonomous Agents and Multiagent Systems*, *AAMAS 2017 Workshops*, *Best Papers*, volume 10642 of *Lecture Notes in Artificial Intelligence*, pages 168–86. Springer International Publishing, 2017.
- 127. Guni Sharon and **Peter Stone**. A protocol for mixed autonomous and human-operated vehicles at intersections. In Gita Sukthankar and Juan A. Rodriguez-Aguilar, editors, *Autonomous Agents and Multiagent Systems AAMAS 2017 Workshops, Best Papers*, volume 10642 of *Lecture Notes in Artificial Intelligence*, pages 151–67. Springer International Publishing, 2017.
- 128. Patrick MacAlpine and **Peter Stone**. Prioritized role assignment for marking. In Sven Behnke, Daniel D. Lee, Sanem Sariel, and Raymond Sheh, editors, *RoboCup 2016: Robot Soccer World Cup XX*, pages 306–18. Lecture Notes in Artificial Intelligence. Springer Verlag, Berlin, 2017.
- 129. Patrick MacAlpine and **Peter Stone**. UT Austin Villa roboCup 3D simulation base code release. In Sven Behnke, Daniel D. Lee, Sanem Sariel, and Raymond Sheh, editors, *RoboCup 2016: Robot Soccer World Cup XX*, pages 135–43. Lecture Notes in Artificial Intelligence. Springer Verlag, Berlin, 2017.
- 130. Patrick MacAlpine and **Peter Stone**. UT Austin Villa RoboCup 3D simulation league competition. In Sven Behnke, Daniel D. Lee, Sanem Sariel, and Raymond Sheh, editors, *RoboCup 2016: Robot Soccer World Cup XX*, pages 515–28. Lecture Notes in Artificial Intelligence. Springer Verlag, Berlin, 2017.
- 131. Kazunori Iwata, Elad Liebman, **Peter Stone**, Toyoshiro Nakashima, Yoshiyuki Anan, and Naohiro Ishii. Bin-Based Estimation of the Amount of Effort for Embedded Software Development Projects with Support Vector Machines. In Roger Lee, editor, *Computer and Information Science 2015*, Studies in Computational Intelligence. Springer Verlag, Berlin, 2016.
- 132. Tsz-Chiu Au, Shun Zhang, and **Peter Stone**. Autonomous intersection management for semi-autonomous vehicles. In *Handbook of Transportation*. Routledge, 2016.
- 133. Patrick MacAlpine, Josiah Hanna, Jason Liang, and **Peter Stone**. UT Austin Villa: RoboCup 2015 3D simulation league competition and technical challenges champions. In Luis Almeida, Jianmin Ji, Gerald Steinbauer, and Sean Luke, editors, *RoboCup-2015: Robot Soccer World Cup XIX*, Lecture Notes in Artificial Intelligence. Springer Verlag, Berlin, 2016.
- 134. David L. Leottau, Javier Ruiz del Solar, Patrick MacAlpine, and **Peter Stone**. A study of layered learning strategies applied to individual behaviors in robot soccer. In Luis Almeida, Jianmin Ji, Gerald Steinbauer, and Sean Luke, editors, *RoboCup-2015: Robot Soccer World Cup XIX*, Lecture Notes in Artificial Intelligence. Springer Verlag, Berlin, 2016.
- 135. Patrick MacAlpine, Mike Depinet, Jason Liang, and **Peter Stone**. UT Austin Villa: RoboCup 2014 3D simulation league competition and technical challenge champions. In Reinaldo A. C. Bianchi, H. Levent Akin, Subramanian Ramamoorthy, and Komei Sugiura, editors, *RoboCup-2014: Robot Soccer World Cup XVIII*, Lecture Notes in Artificial Intelligence. Springer Verlag, Berlin, 2015.
- 136. Mike Depinet, Patrick MacAlpine, and Peter Stone. Keyframe sampling, optimization, and behavior integration: Towards long-distance kicking in the robocup 3d simulation league. In H. Levent Akin, Reinaldo A. C. Bianchi, Subramanian Ramamoorthy, and Komei Sugiura, editors, RoboCup-2014: Robot Soccer World Cup XVIII, Lecture Notes in Artificial Intelligence. Springer Verlag, Berlin, 2015.

- 137. Todd Hester and **Peter Stone**. The Open-Source TEXPLORE Code Release for Reinforcement Learning on Robots In *RoboCup-2013: Robot Soccer World Cup XVII*, 2014.
- 138. Samuel Barrett, Katie Genter, Yuchen He, Todd Hester, Piyush Khandelwal, Jacob Menashe, and **Peter Stone**. The 2012 UT Austin Villa code release. In *RoboCup-2013: Robot Soccer World Cup XVII*, 2014.
- 139. Noa Agmon, Daniel Urieli, and **Peter Stone**. Ship Patrol: Multi Agent Patrol in Complex Environmental Conditions. In Oren Gal, editor, *Advanced in Marine Robotics*, pages 37–62. LAP LAMBERT Academic Publishing, 2013.
- 140. Katie Genter, Noa Agmon, and **Peter Stone**. Role-based ad hoc teamwork. In Gita Sukthankar, Robert P. Goldman, Christopher Geib, David V. Pyhadath, and Hung Hai Bui, editors, *Plan, Activity, and Intent Recognition: Theory and Practice*, pages 251–272. Elsevier, Philadelphia, PA, USA, 2013.
- 141. Samuel Barrett, Katie Genter, Yuchen He, Todd Hester, Piyush Khandelwal, Jacob Menashe, and Peter Stone. UT Austin Villa 2012: Standard platform league world champions. In Xiaoping Chen, Peter Stone, Luis Enrique Sucar, and Tijn Van der Zant, editors, RoboCup-2012: Robot Soccer World Cup XVI, Lecture Notes in Artificial Intelligence. Springer Verlag, Berlin, 2013.
- 142. Patrick MacAlpine, Nick Collins, Adrian Lopez-Mobilia, and **Peter Stone**. UT Austin Villa: RoboCup 2012 3D simulation league champion. In Xiaoping Chen, Peter Stone, Luis Enrique Sucar, and Tijn Van der Zant, editors, *RoboCup-2012: Robot Soccer World Cup XVI*, Lecture Notes in Artificial Intelligence. Springer Verlag, Berlin, 2013.
- 143. Patrick MacAlpine, Francisco Barrera, and **Peter Stone**. Positioning to win: A dynamic role assignment and formation positioning system. In Xiaoping Chen, Peter Stone, Luis Enrique Sucar, and Tijn Van der Zant, editors, *RoboCup-2012: Robot Soccer World Cup XVI*, Lecture Notes in Artificial Intelligence. Springer Verlag, Berlin, 2013.
- 144. Aijun Bai, Xiaoping Chen, Patrick MacAlpine, Daniel Urieli, Samuel Barrett, and **Peter Stone**. Wright Eagle and UT Austin Villa: RoboCup 2011 simulation league champions. In Thomas Roefer, Norbert Michael Mayer, Jesus Savage, and Uluc Saranli, editors, *RoboCup-2011: Robot Soccer World Cup XV*, Lecture Notes in Artificial Intelligence. Springer Verlag, Berlin, 2012.
- 145. Piyush Khandelwal and **Peter Stone**. A low cost ground truth detection system for RoboCup. In T. Roefer, N.M. Mayer, J. Savage, and U. Saranli, editors, *RoboCup-2011: Robot Soccer World Cup XV*, Lecture Notes in Artificial Intelligence. Springer Verlag, 2012.
- 146. Todd Hester and **Peter Stone**. The use of models. in Marco Wiering and Martijn van Otterlo, editors, *Reinforcement Learning: State of the Art*, Springer Verlag, 2011.
- 147. Matthew Hausknecht and **Peter Stone**. Learning powerful kicks on the aibo ers-7: The quest for a striker. In Javier Ruiz-del-Solar, Eric Chown and Paul G. Plöger, editors, *RoboCup-2010: Robot Soccer World Cup XIV*, pages 254–65. Lecture Notes in Artificial Intelligence. Springer Verlag, 2011.
- 148. **Peter Stone**, Michael Quinlan, and Todd Hester. Can robots play soccer? In Ted Richards, editor, *Soccer and Philosophy: Beautiful Thoughts on the Beautiful Game*, volume 51 of *Popular Culture and Philosophy*, pages 75–88. Open Court Publishing Company, 2010.
- 149. Itsuki Noda, **Peter Stone**, Tomohisa Yamashita, and Koichi Kurumatani. Multi-Agent Social Simulation. In Nakashima, H., Aghajan, H., & Augusto, J. C., editor, *Handbook of Ambient Intelligence and Smart Environments*, pages 703–729. 2010.
- 150. **Peter Stone**, Gal A. Kaminka, and Jeffrey S. Rosenschein. Leading a best-response teammate in an ad hoc team. In Esther David, Enrico Gerding, David Sarne, and Onn Shehory, editors, *Agent-Mediated Electronic Commerce: Designing Trading Strategies and Mechanisms for Electronic Markets*, pages 132–146, November 2010.

- 151. David Pardoe and **Peter Stone**. The 2007 TAC SCM prediction challenge. In Wolfgang Ketter, Han La Poutré, Norman Sadeh, Onn Shehory, and William Walsh, editors, *Agent-Mediated Electronic Commerce and Trading Agent Design and Alaysis*, volume 44 of *Lecture Notes in Business Information Processing (LNBIP)*, pages 175–89. 2010.
- 152. Shivaram Kalyanakrishnan and **Peter Stone**. Learning complementary multiagent behaviors: A case study. In *Proceedings of the RoboCup International Symposium 2009*. Springer Verlag, 2009. **Best Student Paper Award.**
- 153. Shivaram Kalyanakrishnan, Todd Hester, Michael Quinlan, Yinon Bentor, and **Peter Stone**. Three humanoid soccer platforms: Comparison and synthesis. In *Proceedings of the RoboCup International Symposium 2009*. Springer Verlag, 2009.
- 154. Kurt Dresner, Mark VanMiddlesworth, and **Peter Stone**. An unmanaged intersection protocol and improved intersection safety for autonomous vehicles. In Ana L. C. Bazzan and Franziska Klugl, editors, *Multi-Agent System for Traffic and Traffic Engineering*, pages 193–217. IGI Global, Hershey, PA, 2009.
- 155. W. Bradley Knox, Juhyun Lee, and **Peter Stone**. Domestic interaction on a segway base. In Luca Iocchi, Hitoshi Matsubara, Alfredo Weitzenfeld, and Changjiu Zhou, editors, *RoboCup-2008: Robot Soccer World Cup XII*, Lecture Notes in Artificial Intelligence. Springer Verlag, Berlin, 2008.
- 156. David Pardoe and **Peter Stone**. An autonomous agent for supply chain management. In Gedas Adomavicius and Alok Gupta, editors, *Handbooks in Information Systems Series: Business Computing*. Elsevier, 2007.
- 157. David Pardoe and **Peter Stone**. Adapting price predictions in TAC SCM. In John Collins, Peyman Faratin, Simon Parsons, Juan A. Rodriguez-Aguilar, Norman M. Sadeh, Onn Shehory, and Elizabeth Sklar, editors, *Agent-Mediated Electronic Commerce and Trading Agent Design and Analysis*, volume 13 of *Lecture Notes in Business Information Processing*, pages 30–45. 2009.
- 158. Mazda Ahmadi and **Peter Stone**. Instance-based action models for fast action planning. In Ubbo Visser, Fernando Ribeiro, Takeshi Ohashi, and Frank Dellaert, editors, *RoboCup-2007: Robot Soccer World Cup XI*. Springer Verlag, Berlin, 2008. **Best Paper Award.**
- 159. Shivaram Kalyanakrishnan, **Peter Stone**, and Yaxin Liu. Model-based reinforcement learning in a complex domain. In Ubbo Visser, Fernando Ribeiro, Takeshi Ohashi, and Frank Dellaert, editors, *RoboCup-2007: Robot Soccer World Cup XI*. Springer Verlag, Berlin, 2008.
- 160. Uli Grasemann, Daniel Stronger, and **Peter Stone**. A neural network-based approach to robot motion control. In Ubbo Visser, Fernando Ribeiro, Takeshi Ohashi, and Frank Dellaert, editors, *RoboCup-2007: Robot Soccer World Cup XI*. Springer Verlag, Berlin, 2008.
- 161. Daniel Stronger and **Peter Stone**. Selective visual attention for object detection on a legged robot. In Gerhard Lakemeyer, Elizabeth Sklar, Domenico Sorenti, and Tomoichi Takahashi, editors, *RoboCup-2006: Robot Soccer World Cup X*, pages 158–170. Volume 4434, Springer Verlag, Berlin, 2007.
- 162. Shivaram Kalyanakrishnan, Yaxin Liu, and Peter Stone. Half field offense in RoboCup soccer: A multiagent reinforcement learning case study. In Gerhard Lakemeyer, Elizabeth Sklar, Domenico Sorenti, and Tomoichi Takahashi, editors, *RoboCup-2006: Robot Soccer World Cup X*, pages 72–85. Volume 4434, Springer Verlag, Berlin, 2007.
 Best Student Paper Award.
- 163. Mohan Sridharan and **Peter Stone**. Autonomous planned color learning on a legged robot. In Gerhard Lakemeyer, Elizabeth Sklar, Domenico Sorenti, and Tomoichi Takahashi, editors, *RoboCup-2006: Robot Soccer World Cup X*, pages 270–278. Volume 4434, Springer Verlag, Berlin, 2007.

- 164. Peggy Fidelman and **Peter Stone**. The chin pinch: A case study in skill learning on a legged robot. In Gerhard Lakemeyer, Elizabeth Sklar, Domenico Sorenti, and Tomoichi Takahashi, editors, *RoboCup-2006: Robot Soccer World Cup X*, pages 59–71. Volume 4434, Springer Verlag, Berlin, 2007.
- 165. Manish Saggar, Thomas D'Silva, Nate Kohl, and **Peter Stone**. Autonomous learning of stable quadruped locomotion. In Gerhard Lakemeyer, Elizabeth Sklar, Domenico Sorenti, and Tomoichi Takahashi, editors, *RoboCup-2006: Robot Soccer World Cup X*, pages 98–109. Volume 4434, Springer Verlag, Berlin, 2007.
 - Nominee for Best Paper Award.
- 166. Kurt Dresner and **Peter Stone**. Multiagent traffic management: Opportunities for multiagent learning. In K. Tuyls et al., editor, *LAMAS 2005*, volume 3898 of *Lecture Notes in Artificial Intelligence*, pages 129–138. Springer Verlag, Berlin, 2006.
- 167. Mohan Sridharan and **Peter Stone**. Towards eliminating manual color calibration at RoboCup. In Itsuki Noda, Adam Jacoff, Ansgar Bredenfeld, and Yasutake Takahashi, editors, *RoboCup-2005: Robot Soccer World Cup IX*, pages 673–81. Volume 4020, Springer Verlag, Berlin, 2006.
- 168. **Peter Stone**, Gregory Kuhlmann, Matthew E. Taylor, and Yaxin Liu. Keepaway soccer: From machine learning testbed to benchmark. In Itsuki Noda, Adam Jacoff, Ansgar Bredenfeld, and Yasutake Takahashi, editors, *RoboCup-2005: Robot Soccer World Cup IX*, pages 93–105. Volume 4020, Springer Verlag, Berlin, 2006.
- 169. David Pardoe and **Peter Stone**. Bidding for customer orders in TAC SCM. In P. Faratin and J.A. Rodriguez-Aguilar, editors, *Agent Mediated Electronic Commerce VI: Theories for and Engineering of Distributed Mechanisms and Systems (AMEC 2004)*, volume 3435 of *Lecture Notes in Artificial Intelligence*, pages 143–157. Springer Verlag, Berlin, 2005.
- 170. Alexander A. Sherstov and **Peter Stone**. Three automated stock-trading agents: A comparative study. In P. Faratin and J.A. Rodriguez-Aguilar, editors, *Agent Mediated Electronic Commerce VI: Theories for and Engineering of Distributed Mechanisms and Systems (AMEC 2004)*, volume 3435 of *Lecture Notes in Artificial Intelligence*, pages 173–187. Springer Verlag, Berlin, 2005.
- 171. Alexander A. Sherstov and **Peter Stone**. Function approximation via tile coding: Automating parameter choice. In J.-D. Zucker and I. Saitta, editors, *SARA 2005*, volume 3607 of *Lecture Notes in Artificial Intelligence*, pages 194–205, Berlin, 2005. Springer Verlag.
- 172. Daniel Stronger and **Peter Stone**. A model-based approach to robot joint control. In Daniele Nardi, Martin Riedmiller, and Claude Sammut, editors, *RoboCup-2004: Robot Soccer World Cup VIII*, volume 3276 of *Lecture Notes in Artificial Intelligence*, pages 297–309. Springer Verlag, Berlin, 2005.
- 173. Gregory Kuhlmann, **Peter Stone**, and Justin Lallinger. The UT Austin Villa 2003 champion simulator coach: A machine learning approach. In Daniele Nardi, Martin Riedmiller, and Claude Sammut, editors, *RoboCup-2004: Robot Soccer World Cup VIII*, volume 3276 of *Lecture Notes in Artificial Intelligence*, pages 636–644. Springer Verlag, Berlin, 2005.
- 174. Mohan Sridharan and **Peter Stone**. Towards illumination invariance in the legged league. In Daniele Nardi, Martin Riedmiller, and Claude Sammut, editors, *RoboCup-2004: Robot Soccer World Cup VIII*, volume 3276 of *Lecture Notes in Artificial Intelligence*, pages 196–208. Springer Verlag, Berlin, 2005.
- 175. Yi Feng, Ronggang Yu, and **Peter Stone**. Two stock-trading agents: Market making and technical analysis. In *Agent Mediated Electronic Commerce V: Designing Mechanisms and Systems. Lecture Notes in Artificial Intelligence*. Springer Verlag, 2004.
- 176. **Peter Stone**. RoboCup as an introduction to CS research. In Daniel Polani, Brett Browning, Andrea Bonarini, and Kazuo Yoshida, editors, *RoboCup-2003: Robot Soccer World Cup VII*, volume 3020 of *Lecture Notes in Artificial Intelligence*, pages 284–95. Springer Verlag, Berlin, 2004.

- 177. **Peter Stone**. Multiagent competitions and research: Lessons from RoboCup and TAC. In Gal A. Kaminka, Pedro U. Lima, and Raul Rojas, editors, *RoboCup-2002: Robot Soccer World Cup VI*, volume 2752 of *Lecture Notes in Artificial Intelligence*, pages 224–237. Springer Verlag, Berlin, 2003.
- 178. **Peter Stone**, Robert E. Schapire, János A. Csirik, Michael L. Littman, and David McAllester. ATTac-2001: A learning, autonomous bidding agent. In *Agent Mediated Electronic Commerce IV: Designing Mechanisms and Systems*, pages 143–160. Volume 2531 of *Lecture Notes in Artificial Intelligence*. Springer Verlag, 2002.
- 179. Paul S. A. Reitsma, **Peter Stone**, János A. Csirik, and Michael L. Littman. Self-enforcing strategic demand reduction. In *Agent Mediated Electronic Commerce IV: Designing Mechanisms and Systems*, pages 289–306. Volume 2531 of *Lecture Notes in Artificial Intelligence*. Springer Verlag, 2002.
- 180. **Peter Stone** and Richard S. Sutton. Keepaway soccer: a machine learning testbed. In Andreas Birk, Silvia Coradeschi, and Satoshi Tadokoro, editors, *RoboCup-2001: Robot Soccer World Cup V*, volume 2377 of *Lecture Notes in Artificial Intelligence*, pages 214–223. Springer Verlag, Berlin, 2002.
- 181. **Peter Stone**, Richard S. Sutton, and Satinder Singh. Reinforcement learning for 3 vs. 2 keepaway. In Peter Stone, Tucker Balch, and Gerhard Kraetzschmar, editors, *RoboCup-2000: Robot Soccer World Cup IV*, volume 2019 of *Lecture Notes in Artificial Intelligence*, pages 249–258. Springer Verlag, Berlin, 2001.
- 182. David McAllester and **Peter Stone**. Keeping the ball from CMUnited-99. In Peter Stone, Tucker Balch, and Gerhard Kraetzschmar, editors, *RoboCup-2000: Robot Soccer World Cup IV*, volume 2019 of *Lecture Notes in Artificial Intelligence*, pages 333–338. Springer Verlag, Berlin, 2001.
- 183. **Peter Stone** and Manuela Veloso. Layered learning. In Ramon López de Mántaras and Enric Plaza, editors, *Machine Learning: ECML 2000*, pages 369–381. Springer Verlag, Barcelona, Catalonia, Spain, May/June 2000. Proceedings of the Eleventh European Conference on Machine Learning (ECML-2000).
- 184. **Peter Stone** and Manuela Veloso. Layered learning and flexible teamwork in RoboCup simulation agents. In M. Veloso, E. Pagello, and H. Kitano, editors, *RoboCup-99: Robot Soccer World Cup III*, volume 1856 of *Lecture Notes in Artificial Intelligence*, pages 495–508. Springer Verlag, Berlin, 2000.
- 185. **Peter Stone** and Manuela Veloso. Task decomposition and dynamic role assignment for real-time strategic teamwork. In J. P. Müller, M. P. Singh, and A. S. Rao, editors, *Intelligent Agents V Proceedings of the Fifth International Workshop on Agent Theories, Architectures, and Languages (ATAL-98), pages 293–308. Springer-Verlag, Heidelberg, 1999.*
- 186. **Peter Stone** and Manuela Veloso. Team-partitioned, opaque-transition reinforcement learning. In Minoru Asada and Hiroaki Kitano, editors, *RoboCup-98: Robot Soccer World Cup II*, volume 1604 of *Lecture Notes in Artificial Intelligence*, pages 261–72. Springer Verlag, Berlin, 1999. Also in *Proceedings of the Third International Conference on Autonomous Agents*, 1999.
- 187. **Peter Stone** and Manuela Veloso. Communication in domains with unreliable, single-channel, low-bandwidth communication. In Alexis Drogoul, Milind Tambe, and Toshio Fukuda, editors, *Collective Robotics*, pages 85–97. Springer Verlag, Berlin, July 1998.
- 188. **Peter Stone** and Manuela Veloso. Using decision tree confidence factors for multiagent control. In Hiroaki Kitano, editor, *RoboCup-97: Robot Soccer World Cup I*, volume 1395 of *Lecture Notes in Artificial Intelligence*, pages 99–111. Springer Verlag, Berlin, 1998.

 Also in *Proceedings of the Second International Conference on Autonomous Agents*, 1998.
- 189. **Peter Stone** and Manuela Veloso. User-guided interleaving of planning and execution. In M. Ghallab and A. Milani, editors, *New Directions in AI Planning*, pages 103–112. IOS Press, 1996.

Refereed Conference Papers

- 190. Krish Agarwal, Yuqian Jiang, Jiaheng Hu, Bo Liu, and **Peter Stone**. L3M+P: Lifelong planning with large language models. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, October 2025.
- 191. Linji Wang, Zifan Xu, **Peter Stone**, and Xuesu Xiao. GACL: Grounded adaptive curriculum learning with active task and performance monitoring. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, October 2025.
- 192. Rohan Chandra, Haresh Karnan, Negar Mehr, **Peter Stone**, and Joydeep Biswas. Multi-agent inverse reinforcement learning in real world unstructured pedestrian crowds. In *IEEE/RSJ International Conference on Intelligent Robots and Systems* (IROS), October 2025.
- 193. Saad Abdul Ghani, Zizhao Wang, **Peter Stone**, , and Xuesu Xiao. Dyna-LfLH: Learning agile navigation in dynamic environments from learned hallucination. In *International Conference on Intelligent Robots and Systems* (IROS), October 2025.
- 194. Zichao Hu, Chen Tang, Michael J. Munje, Yifeng Zhu, Alex Liu, Shuijing Liu, Garrett Warnell, **Peter Stone**, and Joydeep Biswas. ComposableNav: Instruction-following navigation in dynamic environments via composable diffusion. In *Conference on Robot Learning* (CoRL), September 2025.
- 195. Michael J. Munje, Chen Tang, Shuijing Liu, Zichao Hu, Yifeng Zhu, Jiaxun Cui, Garrett Warnell, Joydeep Biswas, and **Peter Stone**. SocialNav-SUB: Benchmarking VLMs for scene understanding in social robot navigation. In *Conference on Robot Learning* (CoRL), September 2025.
- 196. Yuxin Chen, Chen Tang, Jianglan Wei, Chenran Li, Ran Tian, Xiang Zhang, Wei Zhan, Peter Stone, and Masayoshi Tomizuka. MEReQ: Max-ent residual-Q inverse RL for sample-efficient alignment from intervention. In *Conference on Robot Learning* (CoRL), September 2025.
- 197. Jiaheng Hu, Peter Stone, and Roberto Martn-Martn. SLAC: Simulation-pretrained latent action space for whole-body real-world RL. In *Conference on Robot Learning* (CoRL), September 2025.
- 198. Viraj Joshi, Zifan Xu, Bo Liu, **Peter Stone**, and Amy Zhang. Benchmarking massively parallelized multi-task reinforcement learning for robotics tasks. In *Reinforcement Learning Conference* (RLC), August 2025.
- 199. Alexander Levine, **Peter Stone**, and Amy Zhang. Offline action-free learning of Ex-BMDPs by comparing diverse datasets. In *Reinforcement Learning Conference* (RLC), August 2025.
- 200. Michela Proietti, Peter R. Wurman, **Peter Stone**, and Roberto Capobianco. ProtoCRL: Prototype-based network for continual reinforcement learning. In *Reinforcement Learning Conference* (RLC), August 2025.
- 201. Hojoon Lee, Youngdo Lee, Takuma Seno, Donghu Kim, **Peter Stone**, and Jaegul Choo. Hyperspherical normalization for scalable deep reinforcement learning. In *International Conference on Machine Learning* (ICML), July 2025.
- 202. Siddhant Agarwal, Harshit Sikchi, **Peter Stone**, and Amy Zhang. Proto successor measure: Representing the behavior space of an RL agent. In *International Conference on Machine Learning* (ICML), July 2025.
- 203. Weiming Zhuang, Chen Chen, Zhizhong Li, Sina Sajadmanesh, Jingtao Li, Jiabo Huang, Vikash Sehwag, Vivek Sharma, Hirotaka Shinozaki, Felan Carlo Garcia, Yihao Zhan, Naohiro Adachi, Ryoji Eki, Michael Spranger, **Peter Stone**, and Lingjuan Lyu. Argus: A compact and versatile foundation model for vision. In *Conference on Computer Vision and Pattern Recognition* (CVPR), June 2025.
- 204. Mingyo Seo, Yoonyoung Cho, Yoonchang Sung, **Peter Stone**, Yuke Zhu, and Beomjoon Kim. PRESTO: Fast motion planning using diffusion models based on key-configuration environment representation. In *IEEE International Conference on Robotics and Automation* (ICRA), May 2025.

- 205. Adam Labiosa, Zhihan Wang, Siddhant Agarwal, William Cong, Geethika Hemkumar, Abhinav Narayan Harish, Benjamin Hong, Josh Kelle, Chen Li, Yuhao Li, Zisen Shao, **Peter Stone**, and Josiah Hanna. Reinforcement learning within the classical robotics stack: A case study in robot soccer. In *International Conference on Robotics and Automation (ICRA)*, May 2025.
- 206. Jiaheng Hu, Rose Hendrix, Ali Farhadi, Aniruddha Kembhavi, Roberto Martín-Martín, Peter Stone, Kuo-Hao Zeng, and Kiana Ehsani. Flare: Achieving masterful and adaptive robot policies with large-scale reinforcement learning fine-tuning. In *International Conference on Robotics and Automation (ICRA)*, May 2025.
- 207. Bo Liu, Rui Wang, Lemeng Wu, Yihao Feng, **Peter Stone**, and qiang liu. Longhorn: State space models are amortized online learners. In *International Conference on Learning Representations* (ICLR), April 2025.
- 208. Alexander Levine, **Peter Stone**, and Amy Zhang. Learning a fast mixing exogenous block mdp using a single trajectory. In *International Conference on Learning Representations* (ICLR), April 2025.
- 209. Hojoon Lee, Dongyoon Hwang, Donghu Kim, Hyunseung Kim, Jun Jet Tai, Kaushik Subramanian, Peter R. Wurman, Jaegul Choo, Peter Stone, and Takuma Seno. SimBa: Simplicity bias for scaling up parameters in deep reinforcement learning. In *International Conference on Learning Representations* (ICLR), April 2025.
- 210. Caroline Wang, Arrasy Rahman, Ishan Durugkar, Elad Liebman, and **Peter Stone**. N-agent ad hoc teamwork. In *Conference on Neural Information Processing Systems* (NeurIPS), December 2024.
- 211. Borja G. Leon, Francesco Riccio, Kaushik Subramanian, Peter R. Wurman, and **Peter Stone**. Discovering creative behaviors through DUPLEX: diverse universal features for policy exploration. In *Conference on Neural Information Processing Systems* (NeurIPS), December 2024.
- 212. Zizhao Wang, Jiaheng Hu, Caleb Chuck, Stephen Chen, Roberto Martin-Martin, Amy Zhang, Scott Niekum, and **Peter Stone**. SkiLD: unsupervised skill discovery guided by local dependencies. In *Conference on Neural Information Processing Systems* (NeurIPS), December 2024.
- 213. Jiaheng Hu, Zizhao Wang, Roberto Martn-Martn, and **Peter Stone**. Disentangled unsupervised skill discovery for efficient hierarchical reinforcement learning. In *Conference on Neural Information Processing Systems* (NeurIPS), December 2024.
- 214. Shivin Dass, Jiaheng Hu, Ben Abbatematteo, **Peter Stone**, and Roberto Martn-Martn. Learning to look: Seeking information for decision making via policy factorization. In *Conference on Robot Learning (CoRL)*, November 2024.
- 215. Zifan Xu, Haozhu Wang, Dmitriy Bespalov, **Peter Stone**, and Yanjun Qi. LaRS: Latent Reasoning Skill for Chain-of-Thought Reasoning. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, November 2024.
- 216. Carson Stark, Bohkyung Chun, Casey Charleston, Varsha Ravi, Luis Pabon, Surya Sunkari, Tarun Mohan, Peter Stone, and Justin Hart. Dobby: A conversational service robot driven by gpt-4. In International Symposium on Robot and Human Interactive Communication (RO-MAN), August 2024.
- 217. Miguel Vasco, Takuma Seno, Kenta Kawamoto, Kaushik Subramanian, Peter R. Wurman, and **Peter Stone**. A super-human vision-based reinforcement learning agent for autonomous racing in Gran Turismo. In *Reinforcement Learning Conference* (RLC), August 2024. **Outstanding Paper Award on Applications of RL.**
- 218. Alexander Levine, **Peter Stone**, and Amy Zhang. Multistep inverse is not all you need. In *Reinforcement Learning Conference* (RLC), August 2024.
- 219. William Yue, Bo Liu, and **Peter Stone**. t-DGR: A trajectory-based deep generative replay method for continual learning in decision making. In *Conference on Lifelong Learning Agents* (CoLLAs), July 2024.

- 220. Catherine Weaver, Roberto Capobianco, Peter Wurman, **Peter Stone**, and Masayoshi Tomizuka. Real-time trajectory generation via dynamic movement primitives for autonomous racing. In *American Control Conference* (ACC). IEEE, July 2024.
- 221. Amir Hossain Raj, Zichao Hu, Haresh Karnan, Rohan Chandra, Amirreza Payandeh, Luisa Mao, **Peter Stone**, Joydeep Biswas, , and Xuesu Xiao. Rethinking social robot navigation: Leveraging the best of two worlds. In *International Conference on Robotics and Automation* (ICRA), May 2024.
- 222. Haresh, Karnan; Elvin, Yang; Garrett, Warnell; Joydeep, Biswas; Peter, and Stone. Wait, that feels familiar: Learning to extrapolate human preferences for preference-aligned path planning. In *International Conference on Robotics and Automation* (ICRA), May 2024.
- 223. Yoonchang Sung, Rahul Shome, and **Peter Stone**. Asynchronous task plan refinement for multi-robot task and motion planning. In *IEEE International Conference on Robotics and Automation* (ICRA), May 2024.
- 224. Zifan Xu, Amir Hossain Raj, Xuesu Xiao, and **Peter Stone**. Dexterous legged locomotion in confined 3d spaces with reinforcement learning. In *IEEE International Conference on Robotics and Automation* (ICRA), May 2024.
- 225. Ziping Xu, Zifan Xu, Runxuan Jiang, **Peter Stone**, and Ambuj Tewari. Sample efficient myopic exploration through multitask reinforcement learning with diverse tasks. In *International Conference on Learning Representations* (ICLR), May 2024.
- 226. Shahaf S. Shperberg, Bo Liu, and **Peter Stone**. Relaxed exploration constrained reinforcement learning. In *Conference on Autonomous Agents and Multiagent Systems* (AAMAS), May 2024.
- 227. Arrasy Rahman, Jiaxun Cui, and **Peter Stone**. Minimum coverage sets for training robust ad hoc teamwork agents. In *The 38th Annual AAAI Conference on Artificial Intelligence* (AAAI), February 2024.
- 228. Zizhao Wang, Caroline Wang, Xuesu Xiao, Yuke Zhu, and **Peter Stone**. Building minimal and reusable causal state abstractions for reinforcement learning. In *The 38th Annual AAAI Conference on Artificial Intelligence* (AAAI), February 2024.
- 229. W. Bradley Knox, Stephane Hatgis-Kessell, Sigurdur Orn Adalgeirsson, Serena Booth, Anca Dragan, **Peter Stone**, and Scott Niekum. Learning optimal advantage from preferences and mistaking it for reward. In *The 38th Annual AAAI Conference on Artificial Intelligence* (AAAI), February 2024.
- 230. Bo Liu, Yihao Feng, **Peter Stone**, and Qiang Liu. Fast adaptive multitask optimization. In *Conference on Neural Information Processing Systems* (NeurIPS), December 2023.
- 231. Siddhant Agarwal, Ishan Durugkar, **Peter Stone**, and Amy Zhang. f-policy gradients: A general framework for goal conditioned rl using f-divergences. In *Conference on Neural Information Processing Systems*, December 2023.
- 232. Bo Liu, Yifeng Zhu, Chongkai Gao, Yihao Feng, Qiang Liu, Yuke Zhu, and Peter Stone. LIBERO: Benchmarking knowledge transfer in lifelong robot learning. In 37th Conference on Neural Information Processing Systems (NeurIPS) Track on Datasets and Benchmarks, December 2023.
- 233. Swathi Mannem, William Macke, **Peter Stone**, and Reuth Mirsky. Exploring the cost of interruptions in human-robot teaming. In *IEEE-RAS Humanoids*, December 2023.
- 234. Zizhao Wang, Jiaheng Hu, **Peter Stone**, and Roberto Martn-Martn. Elden: Exploration via local dependencies. In *Conference on Neural Information Processing Systems*, December 2023.
- 235. Bo Liu, Yihao Feng, **Peter Stone**, and Qiang Liu. FAMO: Fast Adaptive Multitask Optimization. In *Conference on Neural Information Processing Systems*, December 2023.

- 236. Yifeng Zhu, Zhenyu Jiang, **Peter Stone**, and Yuke Zhu. Learning generalizable manipulation policies with object-centric 3D representations. In *Proceedings of the 6th Conference on Robot Learning* (CoRL), November 2023.
- 237. Haresh Karnan, Elvin Yang, Daniel Furkash, Garrett Warnell, Joydeep Biswas, and **Peter Stone**. STERLING: Self-supervised terrain representation learning from unconstrained robot experience. In *Proceedings of the 6th Conference on Robot Learning* (CoRL), November 2023.
- 238. Xiaohan Zhang, Yifeng Zhu, Yan Ding, Yuqian Jiang, Yuke Zhu, **Peter Stone**, and Shiqi Zhang. Symbolic state space optimization for long horizon mobile manipulation planning. In *International Conference on Intelligent Robots and Systems* (IROS), October 2023.
- 239. Keya Ghonasgi, Reuth Mirsky, Adrian M. Haith, **Peter Stone**, and Ashish Deshpande. A novel control law for multi-joint human-robot interaction tasks while maintaining postural coordination. In *International Conference on Intelligent Robots and Systems* (IROS), October 2023.
- 240. Zifan Xu, Yulin Zhang, Shahaf S. Shperberg, Reuth Mirsky, Yuqian Jiang, Bo Liu, and **Peter Stone**. Model-based meta automatic curriculum learning. In *The Second Conference on Lifelong Learning Agents* (CoLLAs), August 2023.
- 241. Dustin Morrill, Thomas J. Walsh, Daniel Hernandez, Peter R. Wurman, and Peter Stone. Composing efficient, robust tests for policy selection. In *The 39th Conference on Uncertainty in Artificial Intelligence* (UAI), August 2023.
- 242. Jiaheng Hu, **Peter Stone**, and Roberto Martin-Martin. Causal policy gradient for whole-body mobile manipulation. In *Robotics: Science and Systems* (RSS), July 2023.
- 243. Yoonchang Sung and **Peter Stone**. Motion planning (in)feasibility detection using a prior roadmap via path and cut search. In *Robotics: Science and Systems* (RSS), July 2023.
- 244. Vaibhav Bajaj, Guni Sharon, and **Peter Stone**. Task phasing: Automated curriculum learning from demonstrations. In *Proceedings of the 33rd International Conference on Automated Planning and Scheduling* (ICAPS), July 2023.
- 245. Caroline Wang, Garrett Warnell, and **Peter Stone**. D-shape: Demonstration-shaped reinforcement learning via goal conditioning. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems* (AAMAS), May 2023.
- 246. Shahaf S. Shperberg, Bo Liu, and **Peter Stone**. Relaxed exploration constrained reinforcement learning. In *International Conference on Autonomous Agents and Multiagent Systems*, May 2023.
- 247. Jiaxun Cui, Xiaomeng Yang, Mulong Luo, Geunbae Lee, **Peter Stone**, Hsien-Hsin S. Lee, Benjamin Lee, G. Edward Suh, Wenjie Xiong, and Yuandong Tian. Macta: A multi-agent reinforcement learning approach for cache timing attacks and detection. In *The Eleventh International Conference on Learning Representations* (ICLR), May 2023.
- 248. Jinsoo Park, Xuesu Xiao, Garrett Warnell, Harel Yedidsion, and **Peter Stone**. Learning perceptual hallucination for multi-robot navigation in narrow hallways. In *Proceedings of the 2023 IEEE International Conference on Robotics and Automation* (ICRA), May 2023.
- 249. Zifan Xu, Bo Liu, Xuesu Xiao, Anirudh Nair, and **Peter Stone**. Benchmarking reinforcement learning techniques for autonomous navigation. In *Proceedings of the 2023 IEEE International Conference on Robotics and Automation* (ICRA), May 2023.
- 250. Caroline Wang, Ishan Durugkar, Elad Liebman, and **Peter Stone**. Dm²: Decentralized multi-agent reinforcement learning via distribution matching. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence* (AAAI), February 2023.
- 251. Serena Booth, W Bradley Knox, Julie Shah, Scott Niekum, **Peter Stone**, and Alessandro Allievi. The perils of trial-and-error reward design: Misdesign through overfitting and invalid task specifications. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence* (AAAI), Feb 2023.

- 252. Bo Liu, Yihao Feng, Qiang Liu, and **Peter Stone**. Metric residual networks for sample efficient goal-conditioned reinforcement learning. In *Thirty-Seventh AAAI Conference on Artificial Intelligence* (AAAI), Februray 2023.
- 253. Yifeng Zhu, Abhishek Joshi, **Peter Stone**, and Yuke Zhu. VIOLA: Object-centric imitation learning for vision-based robot manipulation. In *Proceedings of the 6th Conference on Robot Learning* (CoRL), December 2022.
- 254. Yoonchang Sung, Zizhao Wang, and **Peter Stone**. Learning to correct mistakes: Backjumping in long-horizon task and motion planning. In *Proceedings of the 6th Conference on Robot Learning* (CoRL), December 2022.
- 255. James MacGlashan, Evan Archer, Alisa Devlic, Takuma Seno, Craig Sherston, Peter R. Wurman, and **Peter Stone**. Value function decomposition for iterative design of reinforcement learning agents. In *Conference on Neural Information Processing Systems* (NeurIPS), December 2022.
- 256. Mao Ye, Bo Liu, Stephen Wright, and **Peter Stone**. BOME! Bilevel optimization made easy: A simple first-order approach. In *Conference on Neural Information Processing Systems* (NeurIPS), December 2022.
- 257. Keya Ghonasgi, Reuth Mirsky, Adrian M Haith, **Peter Stone**, and Ashish D Deshpande. Quantifying changes in kinematic behavior of a human-exoskeleton interactive system. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems* (IROS), October 2022.
- 258. Haresh Karnan, Kavan Sikand, Pranav Atreya, Sadegh Rabiee, Xuesu Xiao, Garrett Warnell, **Peter Stone**, and Joydeep Biswas. VI-IKD: High-speed accurate off-road navigation using learned visual-inertial inverse kinodynamics. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems* (IROS), October 2022.
- 259. Shahaf Shperberg, Bo Liu, Allessandro Allievi, and **Peter Stone**. A rule-based shield: Accumulating safety rules from catastrophic action effects. In *Proceedings of the 1st Conference on Lifelong Learning Agents* (CoLLA), August 2022.
- 260. Bo Liu, Qiang Liu, and **Peter Stone**. Continual learning and private unlearning. In *Proceedings of the 1st Conference on Lifelong Learning Agents* (CoLLA), August 2022.
- 261. Akarsh Kumar, Bo Liu, Risto Miikkulainen, and **Peter Stone**. Effective mutation rate adaptation through group elite selection. In *Proceedings of the Genetic and Evolutionary Computation Conference* (GECCO), July 2022.
- 262. Zizhao Wang, Xuesu Xiao, Zifan Xu, Yuke Zhu, and **Peter Stone**. Causal dynamics learning for task-independent state abstraction. In *Proceedings of the 39th International Conference on Machine Learning* (ICML), July 2022.
- 263. Ghada Sokar, Elena Mocanu, Decebal Constantin Mocanu, Mykola Pechenizkiy, and Peter Stone. Dynamic sparse training for deep reinforcement learning. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence* (IJCAI), July 2022.
 Based on earlier version in *AAMAS workshop on Adaptive Learning Agents* (ALA), May 2022.
 Best Paper Award.
- 264. Jiaxun Cui, Hang Qiu, Dian Chen, **Peter Stone**, and Yuke Zhu. Coopernaut: End-to-end driving with cooperative perception for networked vehicles. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (CVPR), June 2022.
- 265. Haresh Karnan, Garrett Warnell, Xuesu Xiao, and **Peter Stone**. VOILA: Visual-observation-only imitation learning for autonomous navigation. In *International Conference on Robotics and Automation* (ICRA), May 2022.

- 266. Haresh Karnan, Garrett Warnell, Faraz Torabi, and **Peter Stone**. Adversarial imitation learning from video using a state observer. In *International Conference on Robotics and Automation* (ICRA), May 2022.
- 267. Eddy Hudson, Garrett Warnell, Faraz Torabi, and **Peter Stone**. Skeletal feature compensation for imitation learning with embodiment mismatch. In *International Conference on Robotics and Automation* (ICRA), May 2022.
- 268. Xiaohan Zhang, Yifeng Zhu, Yan Ding, Yuke Zhu, **Peter Stone**, and Shiqi Zhang. Visually grounded task and motion planning for mobile manipulation. In *International Conference on Robotics and Automation (ICRA)*, May 2022.
- 269. Ishan Durugkar, Mauricio Tec, Scott Niekum, and **Peter Stone**. Adversarial Intrinsic Motivation for Reinforcement Learning. In *Conference on Neural Information Processing Systems* (NeurIPS), December 2021.
- 270. Bo Liu, Xingchao Liu, **Peter Stone**, and Qiang Liu. Conflict-averse gradient descent for multi-task learning. In *Conference on Neural Information Processing Systems* (NeurIPS), December 2021.
- 271. Sihang Guo, Ruohan Zhang, Bo Liu, Yifeng Zhu, Dana Ballard, Mary Hayhoe, and **Peter Stone**. Machine versus human attention in deep reinforcement learning tasks. In *Conference on Neural Information Processing Systems* (NeurIPS), December 2021.
- 272. Faraz Torabi, Garrett Warnell, and **Peter Stone**. Dealio: Data-efficient adversarial learning for imitation from observation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems* (IROS), September 2021.
- 273. Keya Ghonasgi, Reuth Mirsky, Sanmit Narvekar, Bharath Masetty, Adrian M. Haith, **Peter Stone**, and Ashish D. Deshpande. Capturing skill state in curriculum learning for human skill acquisition. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems* (IROS), September 2021.
- 274. Zizhao Wang, Xuesu Xiao, Alexander J Nettekoven, Kadhiravan Umasankar, Anika Singh, Sriram Bommakanti, Ufuk Topcu, and **Peter Stone**. From agile ground to aerial navigation: Learning from learned hallucination. In *Proceedings of the International Conference on Intelligent Robots and Systems* (IROS), October 2021.
- 275. Bo Liu, Xuesu Xiao, and **Peter Stone**. Team orienteering coverage planning with uncertain reward. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems* (IROS), September 2021.
- 276. Bo Liu, Qiang Liu, **Peter Stone**, Animesh Garg, Yuke Zhu, and Animashree Anandkumar. Coachplayer multi-agent reinforcement learning for dynamic team composition. In *Proceedings of the 38th International Conference on Machine Learning*, *PMLR 139*, 2021 (ICML), July 2021.
- 277. Shih-Yun Lo, Benito R. Fernandez, **Peter Stone**, and Andrea Lockerd Thomaz. Towards Safe Motion Planning in Human Workspaces: A Robust Multi-agent Approach. In *Proceedings of the International Conference on Robotics and Automation* (ICRA), June 2021.
- 278. Blake Holman, Abrar Anwar, Akash Singh, Mauricio Tec, Justin Hart, and **Peter Stone**. Watch where you're going! gaze and head orientation as predictors for social robot navigation. In *Proceedings of the International Conference on Robotics and Automation* (ICRA), June 2021.
- 279. Zizhao Wang, Xuesu Xiao, Bo Liu, Garrett Warnell, and **Peter Stone**. APPLI: Adaptive planner parameter learning from interventions. In *Proceedings of the International Conference on Robotics and Automation* (ICRA), June 2021.
- 280. Zifan Xu, Gauraang Dhamankar, Anirudh Nair, Xuesu Xiao, Garrett Warnell, Bo Liu, Zizhao Wang, and **Peter Stone**. APPLR: Adaptive planner parameter learning from reinforcement. In *Proceedings of the 2021 IEEE International Conference on Robotics and Automation* (ICRA), June 2021.

- 281. Piyush Khandelwal, James MacGlashan, Peter Wurman, and **Peter Stone**. Efficient real-time inference in temporal convolution networks. In *Proceedings of the 2021 International Conference on Robotics and Automation* (ICRA), June 2021.
- 282. Harel Yedidsion, Jennifer Suriadinata, Zifan Xu, Stefan Debruyn, and **Peter Stone**. A scavenger hunt for service robots. In *Proceedings of the 2021 International Conference on Robotics and Automation* (ICRA), June 2021.
- 283. Xuesu Xiao, Bo Liu, and **Peter Stone**. Agile robot navigation through hallucinated learning and sober deployment. In *Proceedings of the 2021 IEEE International Conference on Robotics and Automation (ICRA)*, June 2021.
- 284. Jiaxun Cui, William Macke, Harel Yedidsion, Aastha Goyal, Daniel Urieli, and **Peter Stone**. Scalable multiagent driving policies for reducing traffic congestion. In *Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems* (AAMAS), May 2021.
- 285. Reuth Mirsky and **Peter Stone**. The seeing-eye robot grand challenge: Rethinking automated care. In *Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems* (AAMAS), May 2021.
- 286. Yu-Sian Jiang, Garrett Warnell, and **Peter Stone**. Goal blending for responsive shared autonomy in a navigating vehicle. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence* (AAAI), Feb 2021.
- 287. Yuqian Jiang, Suda Bharadwaj, Bo Wu, Rishi Shah, Ufuk Topcu, and **Peter Stone**. Temporal-logic-based reward shaping for continuing reinforcement learning tasks. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence* (AAAI), February 2021.
- 288. William Macke, Reuth Mirsky, and **Peter Stone**. Expected value of communication for planning in ad hoc teamwork. In *Proceedings of the 35th Conference on Artificial Intelligence* (AAAI), February 2021.
- 289. Siddharth Desai, Ishan Durugkar, Haresh Karnan, Garrett Warnell, Josiah Hanna, and **Peter Stone**. An imitation from observation approach to transfer learning with dynamics mismatch. In *Conference on Neural Information Processing Systems* (NeurIPS), December 2020.
- 290. Lemeng Wu, Bo Liu, Qiang Liu, and **Peter Stone**. Firefly steepest descent: a general approach for growing neural networks. In *Conference on Neural Information Processing Systems* (NeurIPS), December 2020.
- 291. Yuchen Cui, Qiping Zhang, Alessandro Allievi, **Peter Stone**, Scott Niekum, and W. Bradley Knox. The EMPATHIC framework for task learning from implicit human feedback. In *Proceedings of the 4th Conference on Robot Learning* (CoRL), November 2020.
- 292. Jin-Soo Park, Brian Tsang, Harel Yedidsion, Garrett Warnell, Daehyun Kyoung, and **Peter Stone**. Learning to improve multi-robot hallway navigation. In *Proceedings of the 4th Conference on Robot Learning* (CoRL), November 2020.
- 293. Justin Hart, Reuth Mirsky, Xuesu Xiao, Stone Tejeda, Bonny Mahajan, Jamin Goo, Kathryn Baldauf, Sydney Owen, and **Peter Stone**. Using human-inspired signals to disambiguate navigational intentions. In *Proceedings of the International Conference on Social Robotics (ICSR)*, November 2020.
- 294. Rishi Shah, Yuqian Jiang, Justin Hart, and **Peter Stone**. Deep r-learning for continual area sweeping. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems* (IROS), October 2020.
- 295. Siddharth Desai, Haresh Karnan, Josiah P. Hanna, Garrett Warnell, and **Peter Stone**. Stochastic grounded action transformation for robot learning in simulation. In *Proceedigs of the IEEE/RSJ International Conference on Intelligent Robots and Systems* (IROS), October 2020.

- 296. Haresh Karnan, Siddharth Desai, Josiah P. Hanna, Garrett Warnell, and **Peter Stone**. Reinforced grounded action transformation for sim-to-real transfer. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems* (IROS), October 2020.
- 297. Keting Lu, Shiqi Zhang, **Peter Stone**, and Xiaoping Chen. Learning and reasoning for robot dialog and navigation tasks. In *Proceedings of the 21th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 107–117, 1st virtual meeting, July 2020. Association for Computational Linguistics.
- 298. Brahma Pavse, Ishan Durugkar, Josiah Hanna, and **Peter Stone**. Reducing sampling error in batch temporal difference learning. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, July 2020.
- 299. Reuth Mirsky, William Macke, Andy Wang, Harel Yedidsion, and **Peter Stone**. A penny for your thoughts: The value of communication in ad hoc teamwork. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, July 2020.
- 300. Ishan Durugkar, Elad Liebman, and **Peter Stone**. Balancing individual preferences and shared objectives in multiagent reinforcement learning. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, July 2020.
- 301. Harel Yedidsion, Jacqueline Deans, Connor Sheehan, Mahathi Chillara, Justin Hart, **Peter Stone**, and Raymond Mooney. Optimal use of verbal instructions for multi-robot human navigation guidance. In *International Conference on Social Robotics (ICSR)*, pages 133–143. Springer, November 2019.
- 302. Yuqian Jiang, Fangkai Yang, Shiqi Zhang, and **Peter Stone**. Task-Motion Planning with Reinforcement Learning for Adaptable Mobile Service Robots. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems* (IROS), November 2019.
- 303. Felipe Leno Da Silva, Anna Helena Reali Costa, and **Peter Stone**. Building self-play curricula online by playing with expert agents in adversarial games. In *Proceedings of the 8th Brazilian Conference on Intelligent Systems (BRACIS)*, October 2019.
- 304. Shani Alkobi Manish Ravula and **Peter Stone**. Ad hoc teamwork with behavior switching agents. In *International Joint Conference on Artificial Intelligence (IJCAI)*, August 2019.
- 305. Faraz Torabi, Garrett Warnell, and **Peter Stone**. Imitation learning from video by leveraging proprioception. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, August 2019.
- 306. Faraz Torabi, Garrett Warnell, and **Peter Stone**. Recent advances in imitation learning from observation. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, August 2019.
- 307. Ruohan Zhang, Faraz Torabi, Lin Guan, Dana Ballard, and **Peter Stone**. Leveraging Human Guidance for Deep Reinforcement Learning Tasks. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, August 2019.
- 308. Josiah Hanna, Scott Niekum, and **Peter Stone**. Importance sampling policy evaluation with an estimated behavior policy. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, June 2019.
- 309. Yuqian Jiang, Nick Walker, Justin Hart, and **Peter Stone**. Open-world reasoning for service robots. In *Proceedings of the 29th International Conference on Automated Planning and Scheduling (ICAPS)*, July 2019.
- 310. Josiah Hanna and **Peter Stone**. Reducing sampling error in policy gradient learning. In *Proceedings* of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), May 2019.

- 311. Sanmit Narvekar and **Peter Stone**. Learning curriculum policies for reinforcement learning. In *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems (AA-MAS)*, May 2019.
- 312. Guni Sharon, Stephen D. Boyles, Shani Alkoby, and **Peter Stone**. Marginal cost pricing with a fixed error factor in traffic networks. In *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2019.
- 313. Jesse Thomason, Aishwarya Padmakumar, Jivko Sinapov, Nick Walker, Yuqian Jiang, Harel Yedidsion, Justin Hart, **Peter Stone**, and Raymond J. Mooney. Improving Grounded Natural Language Understanding through Human-Robot Dialog. In *IEEE International Conference on Robotics and Automation* (ICRA), May 2019.
- 314. Josiah Hanna, Guni Sharon, Stephen Boyles, and **Peter Stone**. Selecting compliant agents for opt-in micro-tolling. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, January 2019.
- 315. Aishwarya Padmakumar, **Peter Stone**, and Raymond J. Mooney. Learning a policy for opportunistic active learning. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Brussels, Belgium, November 2018.
- 316. Justin W. Hart, Rishi Shah, Sean Kirmani, Nick Walker, Kathryn Baldauf, Nathan John, and **Peter Stone**. PRISM: Pose registration for integrated semantic mapping. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems* (IROS), October 2018.
- 317. Yu-Sian Jiang, Garrett Warnell, and **Peter Stone**. Inferring user intention using gaze in vehicles. In *The 20th ACM International Conference on Multimodal Interaction (ICMI)*, October 2018.
- 318. Elad Liebman and **Peter Stone**. On the impact of music on decision making in cooperative tasks In *Proceedings of the 19th International Society for Music Information retrieval Conference (ISMIR)*, September 2018.
- 319. Rolando Fernandez, Nathan John, Sean Kirmani, Justin Hart, Jivko Sinapov, and **Peter Stone**. Passive demonstrations of light-based robot signals for improved human interpretability. In *Proceedings of the 27th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, August 2018.
- 320. Yu-Sian Jiang, Garrett Warnell, Eduardo Munera, and **Peter Stone**. A study of human-robot copilot systems for en-route destination changing. In *Proceedings of the 27th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, August 2018.
- 321. Elad Liebman, Eric Zavesky, and **Peter Stone**. A stitch in time Autonomous model management via reinforcement learning. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, July 2018.
- 322. Saeid Amiri, Suhua Wei, Shiqi Zhang, Jivko Sinapov, Jesse Thomason, and **Peter Stone**. Multi-modal predicate identification using dynamically learned robot controllers. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, July 2018.
- 323. Faraz Torabi, Garrett Warnell, and **Peter Stone**. Behavioral cloning from observation. In *Proceedings* of the 27th International Joint Conference on Artificial Intelligence (IJCAI), July 2018.
- 324. Jesse Thomason, Jivko Sinapov, Raymond J. Mooney, and **Peter Stone**. Guiding exploratory behaviors for multi-modal grounding of linguistic descriptions. In *Proceedings of the 32nd Conference on Artificial Intelligence (AAAI)*, February 2018.
- 325. Haipeng Chen, Bo An, Guni Sharon, Josiah P. Hanna, **Peter Stone**, Chunyan Miao, and Yeng Chai Soh. Dyetc: Dynamic electronic toll collection for traffic congestion alleviation. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, February 2018.

- 326. Guni Sharon, Michael Albert, Tarun Rambha, Stephen Boyles, and **Peter Stone**. Traffic optimization for a mixture of self-interested and compliant agents. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, February 2018.
- 327. Garrett Warnell, Nicholas Waytowich, Vernon Lawhern, and **Peter Stone**. Deep TAMER: Interactive agent shaping in high-dimensional state spaces. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, February 2018.
- 328. Jesse Thomason, Aishwarya Padmakumar, Jivko Sinapov, Justin Hart, **Peter Stone**, and Raymond J. Mooney. Opportunistic active learning for grounding natural language descriptions. In Sergey Levine, Vincent Vanhoucke, and Ken Goldberg, editors, *Proceedings of the 1st Annual Conference on Robot Learning* (CoRL), pages 67–76, Mountain View, California, November 2017. PMLR.
- 329. Dongcai Lu, Shiqi Zhang, **Peter Stone**, and Xiaoping Chen. Leveraging commonsense reasoning and multimodal perception for robot spoken dialog systems. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems* (IROS), September 2017.
- 330. Josiah Hanna, Philip Thomas, **Peter Stone**, and Scott Niekum. Data-efficient policy evaluation through behavior policy search. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, August 2017.
- 331. Sanmit Narvekar, Jivko Sinapov, and **Peter Stone**. Autonomous task sequencing for customized curriculum design in reinforcement learning. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, August 2017.
- 332. Piyush Khandelwal and **Peter Stone**. Multi-robot human guidance: Human experiments and multi-ple concurrent requests. In *International Conference on Autonomous Agents and Multiagent Systems* (AAMAS), May 2017.
- 333. Guni Sharon, Josiah P. Hanna, Tarun Rambha, Michael W. Levin, Michael Albert, Stephen D. Boyles, and **Peter Stone**. 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), May 2017.
- 334. Josiah Hanna, **Peter Stone**, and Scott Niekum. Bootstrapping with models: Confidence intervals for off-policy evaluation. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2017.
- 335. Shiqi Zhang, Yuqian Jiang, Guni Sharon, and **Peter Stone**. Multirobot symbolic planning under temporal uncertainty. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Sytems (AAMAS)*, May 2017.
- 336. Michael Albert, Vincent Conitzer, and **Peter Stone**. Mechanism Design with Unknown Correlated Distributions: Can We Learn Optimal Mechanisms? In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2017.
- 337. Stefano Albrecht and **Peter Stone**. Reasoning about hypothetical agent behaviours and their parameters. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2017.
- 338. Elad Liebman, Piyush Khandelwal, Maytal Saar-Tsechansky, and **Peter Stone**. Designing better playlists with monte carlo tree search. In *Proceedings of the Twenty-Ninth Conference On Innovative Applications Of Artificial Intelligence (IAAI)*, February 2017.
- 339. Shiqi Zhang, Piyush Khandelwal, and **Peter Stone**. Dynamically constructed (PO)MDPs for adaptive robot planning. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, February 2017.
- 340. Michael Albert, Vincent Conitzer, and **Peter Stone**. Automated Design of Robust Mechanisms. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, February 2017.

- 341. Maxwell Svetlik, Matteo Leonetti, Jivko Sinapov, Rishi Shah, Nick Walker, and **Peter Stone**. Automatic curriculum graph generation for reinforcement learning agents. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, February 2017.
- 342. Josiah Hanna and **Peter Stone**. Grounded action transformation for robot learning in simulation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, February 2017.
- 343. Donghyun Kim, Steven Jens Jorgensen, **Peter Stone**, and Luis Sentis. Dynamic behaviors on the NAO robot with closed-loop whole body operational space control. In *IEEE-RAS International Conference on Humanoid Robots*, 2016.
- 344. Jesse Thomason, Jivko Sinapov, Max Svetlik, **Peter Stone** and Raymond Mooney. Learning multimodal grounded linguistic semantics by playing I, Spy. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, July 2016.
- 345. Jivko Sinapov, Priyanka Khante, Maxwell Svetlik and **Peter Stone**. Learning to order objects using haptic and proprioceptive exploratory behaviors. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, July 2016.
- 346. Shiqi Zhang, Dongcai Lu, Xiaoping Chen, and **Peter Stone**. Robot scavenger hunt: A standardized framework for evaluating intelligent mobile robots. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI) Demonstration Track*, July 2016.
- 347. Elad Liebman, **Peter Stone**, and Corey N. White. Impact of music on decision making in quantitative tasks. In *Proceedings of the 17th International Society for Music Information retrieval Conference (ISMIR)*, August 2016.
- 348. Piyush Khandelwal, Elad Liebman, Scott Niekum, and **Peter Stone**. On the analysis of complex backup strategies in Monte Carlo Tree Search. In *Proceedings of the 33rd International Conference on Machine Learning* (ICML), June 2016.
- 349. Katie Genter and **Peter Stone**. Adding influencing agents to a flock. In *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2016.
- 350. Sanmit Narvekar, Jivko Sinapov, Matteo Leonetti, and **Peter Stone**. Source task creation for curriculum learning. In *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2016.
- 351. Daniel Urieli and **Peter Stone**. An MDP-Based Winning Approach to Autonomous Power Trading: Formalization and Empirical Analysis. In *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2016.
- 352. Matthew Hausknecht and **Peter Stone**. Deep reinforcement learning in parameterized action space. In *Proceedings of the International Conference on Learning Representations (ICLR)*, May 2016.
- 353. Daniel Urieli and **Peter Stone**. Autonomous Electricity Trading using Time-Of-Use Tariffs in a Competitive Market. In *Proceedings of the 30th Conference on Artificial Intelligence (AAAI)*, February 2016.
- 354. **Peter Stone**. What's hot at RoboCup (extended abstract). In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI)*, February 2016.
- 355. Ilaria Gori, Jivko Sinapov, Priyanka Khante, **Peter Stone**, and J.K. Aggarwal. Robot-centric activity recognition "in the wild." In *International Conference on Social Robotics (ICSR)*, October 2015.
- 356. Elad Liebman, **Peter Stone**, and Corey N. White. How music alters decision making: Impact of music stimuli on emotional classification. In *Proceedings of the 16th International Society for Music Information retrieval Conference (ISMIR)*, October 2015.

- 357. Katie Genter, Tim Laue, and **Peter Stone**. Benchmarking robot cooperation without pre-coordination in the robocup standard platform league drop-in player competition. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, September 2015. **Finalist for RoboCup Best Paper Award.**
- 358. Shiqi Zhang, Fangkai Yang, Piyush Khandelwal, and **Peter Stone**. Mobile robot planning using action language BC with an abstraction hierarchy. In *Proceedings of the 13th International Conference on Logic Programming and Non-monotonic Reasoning (LPNMR)*, September 2015.
- 359. Jesse Thomason, Shiqi Zhang, Raymond Mooney, and **Peter Stone**. Learning to interpret natural language commands through human-robot dialog. In *Proceedings of the 2015 International Joint Conference on Artificial Intelligence (IJCAI)*, July 2015.
- 360. Fei Fang, **Peter Stone**, and Milind Tambe. When security games go green: Designing defender strategies to prevent poaching and illegal fishing. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, July 2015.
 - Computational Sustainability Outstanding Paper Award.
- 361. Jivko Sinapov, Sanmit Narvekar, Matteo Leonetti, and **Peter Stone**. Learning inter-task transferability in the absence of target task samples. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. ACM, 2015.
- 362. Piyush Khandelwal, Samuel Barrett, and **Peter Stone**. Leading the way: An efficient multi-robot guidance system. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2015.
- 363. Jacob Menashe and **Peter Stone**. Monte carlo heirarchical model learning. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2015.
- 364. Elad Liebman, Maytal Saar-Tsechansky, and **Peter Stone**. DJ-MC: A reinforcement-learning agent for music playlist recommendation. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2015.
- 365. Katie Genter, Shun Zhang, and **Peter Stone**. Determining placements of influencing agents in a flock. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems (AAMAS-15)*, May 2015.
- 366. Samuel Barrett and **Peter Stone**. Cooperating with unknown teammates in complex domains: A robot soccer case study of ad hoc teamwork. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, January 2015.
- 367. Patrick MacAlpine, Mike Depinet, and **Peter Stone**. UT Austin Villa 2014: RoboCup 3D simulation league champion via overlapping layered learning. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI)*, January 2015.
- 368. Patrick MacAlpine, Eric Price, and **Peter Stone**. SCRAM: Scalable collision-avoiding role assignment with minimal-makespan for formational positioning. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI)*, January 2015.
- 369. Shiqi Zhang and **Peter Stone**. Corpp: Commonsense reasoning and probabilistic planning, as applied to dialog with a mobile robot. In *Proceedings of the 29th Conference on Artificial Intelligence (AAAI)*, January 2015.
- 370. Samuel Barrett, Noa Agmon, Noam Hazon, Sarit Kraus, and **Peter Stone**. Communicating with Unknown Teammates. In *Proceedings of the 21st European Conference on Artificial Intelligence* (ECAI), August 2014.
- 371. Patrick MacAlpine, Katie Genter, Samuel Barrett, and **Peter Stone**. The RoboCup 2013 drop-in player challenges: Experiments in ad hoc teamwork. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems* (IROS), September 2014.

- 372. Katie Genter and **Peter Stone**. Influencing a flock via ad hoc teamwork. In *Proceedings of the Ninth International Conference on Swarm Intelligence* (ANTS), September 2014.
- 373. Daniel Urieli and **Peter Stone**. Tactex'13: A champion adaptive power trading agent. In *Proceedings* of the Twenty-Eighth Conference on Artificial Intelligence (AAAI), July 2014.
- 374. Piyush Khandelwal, Fangkai Yang, Matteo Leonetti, Vladimir Lifschitz, and **Peter Stone**. Planning in action language *BC* while learning action costs for mobile robots. In *International Conference on Automated Planning and Scheduling* (ICAPS), June 2014.
- 375. Noa Agmon, Samuel Barrett, and **Peter Stone**. Modeling uncertainty in leading ad hoc teams. In *Proc.* of 12th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS), May 2014.
- 376. W. Bradley Knox, **Peter Stone**, and Cynthia Breazeal. Training a robot via human feedback: A case study. In *Proceedings of the International Conference on Social Robotics* (ICSR), October 2013. **Best Paper Award.**
- 377. Dustin Carlino, Stephen D. Boyles, and **Peter Stone**. Auction-based autonomous intersection management. In *Proceedings of the 16th IEEE Intelligent Transportation Systems Conference* (ITSC), September 2013.
- 378. Daniel Urieli and **Peter Stone**. Model-selection for non-parametric function approximation in continuous control problems: A case study in a smart energy system. In *Proceedings of the European Conference on Machine Learning and Knowledge Discovery in Databases* (ECML), September 2013.
- 379. Samuel Barrett, **Peter Stone**, Sarit Kraus, and Avi Rosenfeld. Teamwork with limited knowledge of teammates. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence* (AAAI), July 2013.
- 380. Doran Chakraborty and **Peter Stone**. Cooperating with a markovian ad hoc teammate. In *Proceedings* of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), May 2013.
- 381. Todd Hester, Manuel Lopes, and **Peter Stone**. Learning exploration strategies in model-based reinforcement learning. In *The Twelfth International Conference on Autonomous Agents and Multiagent Systems* (AAMAS), May 2013.
- 382. Alon Farchy, Samuel Barrett, Patrick MacAlpine, and **Peter Stone**. Humanoid robots learning to walk faster: From the real world to simulation and back. In *Proc. of 12th Int. Conf. on Autonomous Agents and Multiagent Systems* (AAMAS), May 2013.
- 383. Katie Genter, Noa Agmon, and **Peter Stone**. Ad hoc teamwork for leading a flock. In *Proceedings* of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), May 2013.
- 384. Daniel Urieli and **Peter Stone**. A learning agent for heat-pump thermostat control. In *Proceedings* of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), May 2013.
- 385. Tsz-Chiu Au, Chien-Liang Fok, Sriram Vishwanath, Christine Julien, and **Peter Stone**. Evasion Planning for Autonomous Vehicles at Intersections. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems* (IROS), October 2012.
- 386. W. Bradley Knox and **Peter Stone**. Reinforcement learning from human reward: Discounting in episodic tasks. In *In Proceedings of the 21st IEEE International Symposium on Robot and Human Interactive Communication* (Ro-Man), September 2012.

 Nominee for Best Paper Award.
- 387. Todd Hester and **Peter Stone**. Intrinsically motivated model learning for a developing curious agent. In *The Eleventh International Conference on Development and Learning* (ICDL), November 2012. **Paper of Excellence Award.**

- 388. Dustin Carlino, Mike Depinet, Piyush Khandelwal, and **Peter Stone**. Approximately orchestrated routing and transportation analyzer: Large-scale traffic simulation for autonomous vehicles. In *Proceedings* of the 15th IEEE Intelligent Transportation Systems Conference (ITSC), September 2012.
- 389. Matthew Hausknecht, Piyush Khandelwal, Risto Miikkulainen, and **Peter Stone**. Hyperneat-ggp: A hyperneat-based atari general game player. In *Genetic and Evolutionary Computation Conference* (GECCO), July 2012.
- 390. Shivaram Kalyanakrishnan, Ambuj Tewari, Peter Auer, and **Peter Stone**. PAC subset selection in stochastic multi-armed bandits. In John Langford and Joelle Pineau, editors, *In proceedings of the 29th International Conference on Machine Learning* (ICML), pages 655–662, New York, NY, USA, June-July 2012. Omnipress.
- 391. Patrick MacAlpine, Samuel Barrett, Daniel Urieli, Victor Vu, and **Peter Stone**. Design and optimization of an omnidirectional humanoid walk: A winning approach at the RoboCup 2011 3D simulation competition. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence* (AAAI), July 2012.
- 392. Chien-Liang Fok, Maykel Hanna, Seth Gee, Tsz-Chiu Au, **Peter Stone**, Christine Julien, and Sriram Vishwanath. A platform for evaluating autonomous intersection management policies. In *Proceedings* of the ACM/IEEE Third International Conference on Cyber-Physical Systems (ICCPS), April 2012.
- 393. W. Bradley Knox and **Peter Stone**. Reinforcement learning from simultaneous human and MDP reward. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems* (AAMAS), June 2012.
- 394. W. Bradley Knox, Cynthia Breazeal, and **Peter Stone**. Learning from feedback on actions past and intended. In *In Proceedings of 7th ACM/IEEE International Conference on Human-Robot Interaction, Late-Breaking Reports Session* (HRI), March 2012.
- 395. Patrick MacAlpine, Daniel Urieli, Samuel Barrett, Shivaram Kalyanakrishnan, Francisco Barrera, Adrian Lopez-Mobilia, Nicolae Ştiurcă, Victor Vu, and **Peter Stone**. UT Austin Villa 2011: A champion agent in the RoboCup 3D soccer simulation competition. In *Proc. of 11th Int. Conf. on Autonomous Agents and Multiagent Systems* (AAMAS), June 2012.
- 396. Noa Agmon and **Peter Stone**. Leading ad hoc agents in joint action settings with multiple teammates. In *Proc. of 11th Int. Conf. on Autonomous Agents and Multiagent Systems* (AAMAS), June 2012.
- 397. Noa Agmon, Chien-Liang Fok, Yehuda Emaliah, **Peter Stone**, Christine Julien, and Sriram Vishwanath. On coordination in practical multi-robot patrol. In *IEEE International Conference on Robotics and Automation* (ICRA), May 2012.
- 398. Samuel Barrett and **Peter Stone**. An analysis framework for ad hoc teamwork tasks. In *Proceedings* of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS), June 2012.
- 399. Tsz-Chiu Au, Michael Quinlan, and **Peter Stone**. Setpoint scheduling for autonomous vehicle controllers. In *Proceedings of IEEE International Conference on Robotics and Automation* (ICRA), May 2012.
- 400. Todd Hester, Michael Quinlan, and **Peter Stone**. RTMBA: A real-time model-based reinforcement learning architecture for robot control. In *IEEE International Conference on Robotics and Automation* (ICRA), May 2012.
- 401. Matthew Hausknecht, Tsz-Chiu Au, **Peter Stone**, David Fajardo, and Travis Waller. Dynamic Lane Reversal in Autonomous Traffic Management. In *Proceedings of the IEEE Intelligent Transportation systems Society Conference* (ITSC), October 2011.

- 402. Matthew Hausknecht, Tsz-Chiu Au, and **Peter Stone**. Autonomous Intersection Management: Multi-Intersection Optimization. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems* (IROS), September 2011.
- 403. Noa Agmon, Daniel Urieli, and **Peter Stone**. Multiagent patrol generalized to complex environmental conditions. In *Proceedings of the Twenty-Fifth Conference on Artificial Intelligence* (AAAI), August 2011.
- 404. Tsz-Chiu Au, Neda Shahidi, and **Peter Stone**. Enforcing Liveness in Autonomous Traffic Management. In *Proceedings of the Twenty-Fifth Conference on Artificial Intelligence* (AAAI), August 2011.
- 405. Raz Lin, Noa Agmon, Sarit Kraus, Samuel Barrett, and **Peter Stone**. Comparing Agents' Success against People in Security Domains. In *Proceedings of the Twenty-Fifth Conference on Artificial Intelligence* (AAAI), August 2011.
- 406. Doran Chakraborty and **Peter Stone**. Structure learning in ergodic factored MDPs without knowledge of the transition function's in-degree. In *Proceedings of the Twenty Eighth International Conference on Machine Learning* (ICML), June 2011.
- 407. Samuel Barrett, **Peter Stone**, and Sarit Kraus. Empirical evaluation of ad hoc teamwork in the pursuit domain. In *Proc. of 11th Int. Conf. on Autonomous Agents and Multiagent Systems* (AAMAS), May 2011.
- 408. Daniel Urieli, Patrick MacAlpine, Shivaram Kalyanakrishnan, Yinon Bentor, and **Peter Stone**. On optimizing interdependent skills: A case study in simulated 3d humanoid robot soccer. In *Proc. of 10th Int. Conf. on Autonomous Agents and Multiagent Systems* (AAMAS), May 2011.
- 409. David Pardoe and **Peter Stone**. A Particle Filter for Bid Estimation in Ad Auctions with Periodic Ranking Observations. In *Proc. of 10th Int. Conf. on Autonomous Agents and Multiagent Systems* (AAMAS), May 2011.
- 410. Michael Quinlan, Tsz-Chiu Au, Jesse Zhu, Nicolae Stiurca, and **Peter Stone**. Bringing simulation to life: A mixed reality autonomous intersection. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems* (IROS), October 2010.
- 411. Tobias Jung and **Peter Stone**. Gaussian processes for sample efficient reinforcement learning with RMAX-like exploration. In *The European Conference on Machine Learning* (ECML), September 2010.
- 412. Todd Hester and **Peter Stone**. Real time targeted exploration in large domains. In *The Ninth International Conference on Development and Learning* (ICDL), August 2010.
- 413. Doran Chakraborty and **Peter Stone**. Convergence, targeted optimality and safety in multiagent learning. In *Proceedings of the Twenty-seventh International Conference on Machine Learning* (ICML), June 2010.
- 414. Shivaram Kalyanakrishnan and **Peter Stone**. Efficient selection of multiple bandit arms: Theory and practice. In *Proceedings of the Twenty-seventh International Conference on Machine Learning* (ICML), 2010.
- 415. David Pardoe and **Peter Stone**. Boosting for regression transfer. In *Proceedings of the 27th International Conference on Machine Learning* (ICML), June 2010.
- 416. **Peter Stone**, Gal A. Kaminka, Sarit Kraus, and Jeffrey S. Rosenschein. Ad hoc autonomous agent teams: Collaboration without pre-coordination. In *Proceedings of the Twenty-Fourth Conference on Artificial Intelligence* (AAAI), July 2010.
- 417. W. Bradley Knox and **Peter Stone**. Combining manual feedback with subsequent MDP reward signals for reinforcement learning. In *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems* (AAMAS), May 2010.
 - **Best Student Paper Award.**

- 418. David Pardoe, Doran Chakraborty, and **Peter Stone**. TacTex09: A champion bidding agent for ad auctions. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems* (AAMAS), May 2010.
- 419. **Peter Stone** and Sarit Kraus. To teach or not to teach? decision making under uncertainty in ad hoc teams. In *The Ninth International Conference on Autonomous Agents and Multiagent Systems* (AAMAS). International Foundation for Autonomous Agents and Multiagent Systems, May 2010.
- 420. Todd Hester, Michael Quinlan, and **Peter Stone**. Generalized model learning for reinforcement learning on a humanoid robot. In *IEEE International Conference on Robotics and Automation* (ICRA), May 2010.
- 421. Peter Djeu, Michael Quinlan, and **Peter Stone**. Improving particle filter performance using sse instructions. In *Proceedings of the IEEE/RSJ International Conference on Intelligent RObots and Systems* (IROS), October 2009.
- 422. W. Bradley Knox and **Peter Stone**. Interactively shaping agents via human reinforcement: The TAMER framework. In *The Fifth International Conference on Knowledge Capture*, September 2009.
- 423. Tobias Jung and **Peter Stone**. Feature selection for value function approximation using bayesian model selection. In *The European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases* (ECML), September 2009.
- 424. Nicholas K. Jong and **Peter Stone**. Compositional models for reinforcement learning. In *The European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases* (ECML), September 2009.
- 425. Shivaram Kalyanakrishnan and **Peter Stone**. An empirical analysis of value function-based and policy search reinforcement learning. In *The Eighth International Conference on Autonomous Agents and Multiagent Systems* (AAMAS), pages 749–56, May 2009.
- 426. Todd Hester and **Peter Stone**. Generalized model learning for reinforcement learning in factored domains. In *The Eighth International Conference on Autonomous Agents and Multiagent Systems* (AA-MAS), May 2009.
- 427. W. Bradley Knox and **Peter Stone**. TAMER: Training an Agent Manually via Evaluative Reinforcement. In *IEEE 7th International Conference on Development and Learning* (ICDL), August 2008.
- 428. Doran Chakraborty and **Peter Stone**. Online multiagent learning against memory bounded adversaries. In *Machine Learning and Knowledge Discovery in Databases* (ECML), volume 5212 of *Lecture Notes in Artificial Intelligence*, pages 211–26, September 2008.
- 429. Matthew E. Taylor, Nicholas K. Jong, and **Peter Stone**. Transferring instances for model-based reinforcement learning. In *Machine Learning and Knowledge Discovery in Databases* (ECML), volume 5212 of *Lecture Notes in Artificial Intelligence*, pages 488–505, September 2008.
- 430. Nicholas K. Jong and **Peter Stone**. Hierarchical model-based reinforcement learning: R-MAX + MAXQ. In *Proceedings of the Twenty-Fifth International Conference on Machine Learning* (ICML), July 2008.
- 431. Joseph Reisinger, **Peter Stone**, and Risto Miikkulainen. Online kernel selection for bayesian reinforcement learning. In *Proceedings of the Twenty-Fifth International Conference on Machine Learning* (ICML), July 2008.
- 432. Jonathan Wildstrom, **Peter Stone**, and Emmett Witchel. CARVE: A Cognitive Agent for Resource Value Estimation. In *International Conference on Autonomic Computing* (ICAC), June 2008.
- 433. Mohan Sridharan and **Peter Stone**. Long-term vs. greedy action planning for color learning on a mobile robot. In *International Conference on Computer Vision Theory and Applications*, January 2008.

- 434. Nicholas K. Jong, Todd Hester, and **Peter Stone**. The utility of temporal abstraction in reinforcement learning. In *The Seventh International Joint Conference on Autonomous Agents and Multiagent Systems* (AAMAS), May 2008.
- 435. Matthew E. Taylor, Gregory Kuhlmann, and **Peter Stone**. Autonomous transfer for reinforcement learning. In *The Seventh International Joint Conference on Autonomous Agents and Multiagent Systems* (AAMAS), May 2008.
- 436. Juhyun Lee and **Peter Stone**. Person tracking on a mobile robot with heterogeneous inter-characteristic feedback. In *IEEE International Conference on Robotics and Automation* (ICRA), May 2008.
- 437. Daniel Stronger and **Peter Stone**. Maximum likelihood estimation of sensor and action model functions on a mobile robot. In *IEEE International Conference on Robotics and Automation* (ICRA), May 2008.
- 438. Todd Hester and **Peter Stone**. Negative information and line observations for monte carlo localization. In *IEEE International Conference on Robotics and Automation* (ICRA), May 2008.
- 439. Matthew E. Taylor, Gregory Kuhlmann, and **Peter Stone**. Transfer learning and intelligence: an argument and approach. In *Proceedings of the First Conference on Artificial General Intelligence*, March 2008.
- 440. M. Sridharan and **Peter Stone**. Color learning on a mobile robot: Towards full autonomy under changing illumination. In *Proceedings of the IEEE International Conference on Intelligent Robots and Systems* (IROS), 2007.
- 441. Gregory Kuhlmann and **Peter Stone**. Graph-based domain mapping for transfer learning in general games. In *Proceedings of The Eighteenth European Conference on Machine Learning* (ECML), September 2007.
- 442. Matthew E. Taylor, Shimon Whiteson, and **Peter Stone**. Temporal difference and policy search methods for reinforcement learning: An empirical comparison. In *Proceedings of the Twenty-Second Conference on Artificial Intelligence* (AAAI), July 2007. Nectar Track.
- 443. Matthew E. Taylor and **Peter Stone**. Cross-domain transfer for reinforcement learning. In *Proceedings* of the Twenty-Fourth International Conference on Machine Learning (ICML), June 2007.
- 444. D Stronger and P Stone. A comparison of two approaches for vision and self-localization on a mobile robot. In *IEEE International Conference on Robotics and Automation* (ICRA), April 2007.
- 445. Matthew E. Taylor, Shimon Whiteson, and **Peter Stone**. Transfer via inter-task mappings in policy search reinforcement learning. In *The Sixth International Joint Conference on Autonomous Agents and Multiagent Systems* (AAMAS), May 2007.
- 446. Mazda Ahmadi, Matthew E. Taylor, and **Peter Stone**. IFSA: Incremental feature-set augmentation for reinforcement learning tasks. In *The Sixth International Joint Conference on Autonomous Agents and Multiagent Systems* (AAMAS), May 2007.
 - Nominee for Best Student Paper Award.
- 447. Shivaram Kalyanakrishnan and **Peter Stone**. Batch reinforcement learning in a complex domain. In *The Sixth International Joint Conference on Autonomous Agents and Multiagent Systems* (AAMAS), May 2007.
 - Nominee for Best Student Paper Award.
- 448. Nicholas K. Jong and **Peter Stone**. Model-based function approximation for reinforcement learning. In *The Sixth International Joint Conference on Autonomous Agents and Multiagent Systems* (AAMAS), May 2007.
- 449. **Peter Stone**. Learning and multiagent reasoning for autonomous agents. In *The 20th International Joint Conference on Artificial Intelligence* (IJCAI), pages 13–30, January 2007.

- 450. Mohan Sridharan and **Peter Stone**. Color learning on a mobile robot: Towards full autonomy under changing illumination. In *The 20th International Joint Conference on Artificial Intelligence* (IJCAI), pages 2212–2217, January 2007.
- 451. Kurt Dresner and **Peter Stone**. Sharing the road: Autonomous vehicles meet human drivers. In *The 20th International Joint Conference on Artificial Intelligence* (IJCAI), pages 1263–68, January 2007.
- 452. Bikramjit Banerjee and **Peter Stone**. General game learning using knowledge transfer. In *The 20th International Joint Conference on Artificial Intelligence* (IJCAI), pages 672–677, January 2007.
- 453. Jonathan Wildstrom, **Peter Stone**, Emmett Witchel, and Mike Dahlin. Machine learning for on-line hardware reconfiguration. In *The 20th International Joint Conference on Artificial Intelligence* (IJCAI), pages 1113–1118, January 2007.
- 454. Mohan Sridharan and **Peter Stone**. Autonomous planned color learning on a mobile robot without labeled data. In *The Ninth International Conference on Control, Automation, Robotics and Vision* (ICARCV), December 2006.
- 455. David Pardoe, **Peter Stone**, Maytal Saar-Tsechansky, and Kerem Tomak. Adaptive mechanism design: A metalearning approach. In *The Eighth International Conference on Electronic Commerce* (ICEC), pages 92–102, August 2006.
- 456. Mazda Ahmadi and **Peter Stone**. Keeping in touch: Maintaining biconnected structure by homogeneous robots. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence* (AAAI), pages 580–85, July 2006.
- 457. Shimon Whiteson and **Peter Stone**. Sample-efficient evolutionary function approximation for reinforcement learning. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence* (AAAI), pages 518–23, July 2006.
- 458. David Pardoe and **Peter Stone**. TacTex-2005: A champion supply chain management agent. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence* (AAAI), pages 1489–94, July 2006.
- 459. Gregory Kuhlmann, Kurt Dresner, and **Peter Stone**. Automatic heuristic construction in a complete general game player. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence* (AAAI), pages 1457–62, July 2006.
- 460. Gregory Kuhlmann, William B. Knox, and **Peter Stone**. Know thine enemy: A champion RoboCup coach agent. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence* (AAAI), pages 1463–68, July 2006.
- 461. Yaxin Liu and **Peter Stone**. Value-function-based transfer for reinforcement learning using structure mapping. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence* (AAAI), pages 415–20, July 2006.
- 462. Shimon Whiteson and **Peter Stone**. On-line evolutionary computation for reinforcement learning in stochastic domains. In *Proceedings of the Genetic and Evolutionary Computation Conference* (GECCO), pages 1577–84, July 2006.
- 463. Harish Subramanian, Subramanian Ramamoorthy, **Peter Stone**, and Benjamin Kuipers. Designing safe, profitable automated stock trading agents using evolutionary algorithms. In *Proceedings of the Genetic and Evolutionary Computation Conference* (GECCO), July 2006.
- 464. Matthew Taylor, Shimon Whiteson, and Peter Stone. Comparing evolutionary and temporal difference methods for reinforcement learning. In *Proceedings of the Genetic and Evolutionary Computation Conference* (GECCO), pages 1321–28, July 2006. Best Paper Award, GA Track.

- 465. David Pardoe and **Peter Stone**. Predictive planning for supply chain management. In *International Conference on Automated Planning and Scheduling* (ICAPS), June 2006.
- 466. Mazda Ahmadi and **Peter Stone**. A multi-robot system for continuous area sweeping tasks. In *IEEE International Conference on Robotics and Automation* (ICRA), pages 1724–1729, May 2006.
- 467. Mohan Sridharan and **Peter Stone**. Real-time vision on a mobile robot platform. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems* (IROS), August 2005.
- 468. Nicholas K. Jong and **Peter Stone**. State Abstraction Discovery from Irrelevant State Variables. In *The Nineteenth International Joint Conference on Artificial Intelligence* (IJCAI), pages 752–757, August 2005.
- 469. Mohan Sridharan and **Peter Stone**. Autonomous color learning on a mobile robot. In *Proceedings of the Twentieth National Conference on Artificial Intelligence* (AAAI), July 2005.
- 470. Alexander A. Sherstov and **Peter Stone**. Improving action selection in MDP's via knowledge transfer. In *Proceedings of the Twentieth National Conference on Artificial Intelligence* (AAAI), July 2005.
- 471. Matthew E. Taylor, **Peter Stone**, and Yaxin Liu. Value functions for RL-based behavior transfer: A comparative study. In *Proceedings of the Twentieth National Conference on Artificial Intelligence* (AAAI), July 2005.
- 472. Matthew E. Taylor and **Peter Stone**. Behavior transfer for value-function-based reinforcement learning. In *The Fourth International Joint Conference on Autonomous Agents and Multiagent Systems* (AAMAS), pages 53–59, July 2005.
- 473. Kurt Dresner and **Peter Stone**. Multiagent traffic management: An improved intersection control mechanism. In *The Fourth International Joint Conference on Autonomous Agents and Multiagent Systems* (AAMAS), July 2005.
- 474. Mazda Ahmadi and **Peter Stone**. Continuous area sweeping: A task definition and initial approach. In *The 12th International Conference on Advanced Robotics* (ICAR), July 2005.
- 475. Jonathan Wildstrom, **Peter Stone**, Emmett Witchel, Raymond J. Mooney, and Mike Dahlin. Towards self-configuring hardware for distributed computer systems. In *The Second International Conference on Autonomic Computing* (ICAC), pages 241–249, June 2005.
- 476. Shimon Whiteson, **Peter Stone**, Kenneth O. Stanley, Risto Miikkulainen, and Nate Kohl. Automatic feature selection via neuroevolution. In *The Genetic and Evolutionary Computation Conference* (GECCO), June 2005.
- 477. Mohan Sridharan, Gregory Kuhlmann, and **Peter Stone**. Practical vision-based monte carlo localization on a legged robot. In *IEEE International Conference on Robotics and Automation* (ICRA), April 2005.
- 478. Nate Kohl and **Peter Stone**. Machine Learning for Fast Quadrupedal Locomotion. In *Proceedings of the Nineteenth National Conference on Artificial Intelligence* (AAAI), pages 611–616, San Jose, CA, July 2004.
- 479. Kurt Dresner and **Peter Stone**. Multiagent Traffic Management: A Reservation-Based Intersection Control Mechanism. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems* (AAMAS), New York, NY, July 2004.

 Also in *The 5th IFAC Symposium on Intelligent Autonomous Vehicles* (IAV2004), July 2004.
- 480. Mohan Sridharan and **Peter Stone**. Towards On-Board Color Constancy on Mobile Robots. In *the First Canadian Conference on Computer and Robot Vision* (CRV), London, Ontario, Canada, May 2004.
- 481. Nate Kohl and **Peter Stone**. Policy Gradient Reinforcement Learning for Fast Quadrapedal Locomotion. In *Proceedings of the IEEE International Conference on Robotics and Automation* (ICRA), New Orleans, LA, May 2004.

- 482. Ronggang Yu and **Peter Stone**. Performance analysis of a counter-intuitive automated stock-trading strategy. In *Proceedings of the Fifth International Conference on Electronic Commerce* (ICEC), Pittsburgh, PA, October 2003.
- 483. Satinder Singh, Michael L. Littman, Nicholas K. Jong, David Pardoe, and **Peter Stone**. Learning predictive state representations. In *Proceedings of the Twentieth International Conference on Machine Learning* (ICML), August 2003.
- 484. Shimon Whiteson and **Peter Stone**. Concurrent layered learning. In *Second International Joint Conference on Autonomous Agents and Multiagent Systems* (AAMAS), pages 193–200, July 2003.
- 485. Michael P. Wellman, Amy Greenwald, **Peter Stone**, and Peter R. Wurman. The 2001 trading agent competition. In *Proceedings of the Fourteenth Innovative Applications of Artificial Intelligence Conference* (IAAI), July 2002.
- 486. Robert E. Schapire, **Peter Stone**, David McAllester, Michael L. Littman, and János A. Csirik. Modeling auction price uncertainty using boosting-based conditional density estimation. In *Proceedings of the Nineteenth International Conference on Machine Learning* (ICML), July 2002.
- 487. **Peter Stone** and David McAllester. An architecture for action selection in robotic soccer. In *Proceedings of the Fifth International Conference on Autonomous Agents* (AA), 2001.
- 488. Charles Lee Isbell Jr., Christian R. Shelton, Michael Kearns, Satinder Singh, and Peter Stone. A social reinforcement learning agent. In *Proceedings of the Fifth International Conference on Autonomous Agents* (AA), 2001.
 Best Paper Award.
- 489. **Peter Stone**, Patrick Riley, and Manuela Veloso. Defining and using ideal teammate and opponent agent models. In *Proceedings of the Twelfth Annual Conference on Innovative Applications of Artificial Intelligence* (IAAI), 2000.
- 490. **Peter Stone**. TPOT-RL applied to network routing. In *Proceedings of the Seventeenth International Conference on Machine Learning* (ICML), pages 935–942, 2000.
- 491. Manuela Veloso, Michael Bowling, Sorin Achim, Kwun Han, and **Peter Stone**. CMUnited-98: A team of robotic soccer agents. In *Proceedings of the Eleventh Annual Conference on Innovative Applications of Artificial Intelligence* (IAAI), 1999.
- 492. Manuela Veloso, **Peter Stone**, and Michael Bowling. Anticipation as a key for collaboration in a team of agents: A case study in robotic soccer. In *Proceedings of SPIE Sensor Fusion and Decentralized Control in Robotic Systems II*, volume 3839, pages 134–143, Boston, September 1999.
- 493. Manuela Veloso and **Peter Stone**. Individual and collaborative behaviors in a team of homogeneous robotic soccer agents. In *Proceedings of the Third International Conference on Multi-Agent Systems* (ICMAS), pages 309–316, 1998.
- 494. Hiroaki Kitano, Milind Tambe, **Peter Stone**, Manuela Veloso, Silvia Coradeschi, Eiichi Osawa, Hitoshi Matsubara, Itsuki Noda, and Minoru Asada. The RoboCup synthetic agent challenge 97. In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence* (IJCAI), pages 24–29, San Francisco, CA, 1997. Morgan Kaufmann.
- 495. Gregg Rabideau, Steve Chien, **Peter Stone**, Jason Willis, Curt Eggemeyer, and Tobias Mann. Interactive, repair-based planning and scheduling for shuttle payload operations. In *Proceedings of the 1997 IEEE Aerospace Conference*, pages 325–341, Aspen, CO, February 1997.
- 496. Manuela Veloso, **Peter Stone**, and Sorin Achim. A layered approach for an autonomous robotic soccer system. In *Proceedings of the First International Conference on Autonomous Agents* (AA), pages 530–531, Marina del Rey, CA, February 1997.

- 497. **Peter Stone** and Manuela Veloso. Beating a defender in robotic soccer: Memory-based learning of a continuous function. In David S. Touretzky, Michael C. Mozer, and Michael E. Hasselmo, editors, *Advances in Neural Information Processing Systems* 8 (NIPS), pages 896–902, Cambridge, MA, 1996. MIT Press.
- 498. **Peter Stone**, Manuela Veloso, and Jim Blythe. The need for different domain-independent heuristics. In *Proceedings of the Second International Conference on AI Planning Systems* (AIPS), pages 164–169, June 1994.

Refereed Workshop and Symposium Papers

- 499. Joydeep Biswas, Don Fussell, **Peter Stone**, Kristin Patterson, Kristen Procko, Lea Sabatini, and Zifan Xu. The essentials of AI for life and society: An AI literacy course for the university community. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 28973–8, April 2025.
- 500. Caleb Chuck, Carl Qi, Michael J. Munje, Shuozhe Li, Max Rudolph, Chang Shi, Siddhant Agarwal, Harshit Sikchi, Abhinav Peri, Sarthak Dayal, Evan Kuo, Kavan Mehta, Anthony Wang, **Peter Stone**, Amy Zhang, and Scott Niekum. Robot air hockey: A manipulation testbed for robot learning with reinforcement learning. In *ICRA Workshop on Manipulation Skills*, May 2024.
- 501. Shivin Dass, Wensi Ai, Yuqian Jiang, Samik Singh, Jiaheng Hu, Ruohan Zhang, **Peter Stone**, Ben Abbatematteo, and Roberto Martin-Martin. Telemoma: A modular and versatile teleoperation system for mobile manipulation. In *ICRA Workshop on Mobile Manipulation and Embodied Intelligence*, May 2024.
- 502. Zifan Xu, Anirudh Nair, Xuesu Xiao, and **Peter Stone**. Learning real-world autonomous navigation by self-supervised environment synthesis. In *IROS Workshop on Photorealistic Image and Environment Synthesis for Robotics (PIES-Rob)*, January 2023.
- 503. Juan Huerta, Bo Liu, and **Peter Stone**. VaryNote: a method to automatically vary the number of notes in symbolic music. In *The 16th International Symposium on Computer Music Multidisciplinary Research*, (CMMR), Lecture Notes in Computer Science. Springer, November 2023.
- 504. Shravan Ravi, Gary Wang, Shreyas Satewar, Xuesu Xiao, Garrett Warnell, Joydeep Biswas, and **Peter Stone**. Visually adaptive geometric navigation. In *IEEE International Symposium on Safety, Security, and Rescue Robotics* (SSRR), November 2023.
- 505. Elliott Hauser, Yao-Cheng Chan, Geethika Hemkumar, Daksh Dua, Parth Chonkar, Efren Mendoza Enriquez, Tiffany Kao, Shikhar Gupta, Huihai Wang, Justin Hart, Reuth Mirsky, Joydeep Biswas, Junfeng Jiao, and **Peter Stone**. "what's that robot doing here?": Factors influencing perceptions of incidental encounters with autonomous quadruped robots. In *Proceedings of the First International Symposium on Trustworthy Autonomous Systems* (TAS), pages 1–15, July 2023.
- 506. Sai Kiran Narayanaswami, David Fridovich-Keil, Swarat Chaudhuri, and **Peter Stone**. Program embeddings for rapid mechanism evaluation. In *ICRA workshop on Multi-Robot Learning*, May 2023.
- 507. Swathi Mannem, William Macke, **Peter Stone** and Reuth Mirsky. Exploring the cost of interruptions in human-robot teaming. In *AAMAS workshop on Rebellion and disobedience in AI* (RaD-AI23), May 2023.

Best Paper Award.

508. Anthony Francis, Claudia Perez-D'Arpino, Chengshu Li, Fei Xia, Alexandre Alahi, Rachid Alami1, Aniket Bera, Abhijat Biswas, Joydeep Biswas, Rohan Chandra, Hao-Tien Lewis Chiang, Michael Everett, Sehoon Ha, Justin Hart, Jonathan P. How, Haresh Karnan, Tsang-Wei Edward Lee, Luis J. Manso, Reuth Mirksy, Soren Pirk, Phani Teja Singamaneni, **Peter Stone**, Ada V. Taylor, Peter Trautman, Nathan Tsoi, Marynel Vazquez, Xuesu Xiao, Peng Xu, Naoki Yokoyama, Alexander Toshev, and Roberto Martin-Martin Benchmarking social robot navigation across academia and industry In *AAAI Spring Symposium on "HRI in Academia and Industry: Bridging the Gap*, March 2023. arXiv e-Prints 2306.16740. **Nominee for Best Paper Award.**

- 509. Eddy Hudson, Ishan Durugkar, Garrett Warnell, and **Peter Stone**. ABC: Adversarial Behavioral Cloning for offline mode-seeking imitation learning In *NeurIPS 2022 Workshop on Deep RL*, December 2022. *arXiv e-Prints 2211.04005*.
- 510. Anirudh Nair, Fulin Jiang, Kang Hou, Zifan Xu, Shuozhe Li, Xuesu Xiao, and **Peter Stone**. Dynabarn: Benchmarking metric ground navigation in dynamic environments. In *Proceedings of the 2022 IEEE International Symposium on Safety, Security, and Rescue Robotics* (SSRR), November 2022.
- 511. Yair Hanina, Reuth Mirsky, William Macke, and **Peter Stone**. Quantifying human rationality in ad-hoc teamwork. In *AAMAS workshop on Autonomous Robots and Multirobot Systems* (ARMS), May 2022.
- 512. Yulin Zhang, William Macke, Jiaxun Cui, Daniel Urieli, and **Peter Stone**. Learning a robust multiagent driving policy for traffic congestion reduction. In *AAMAS workshop on Adaptive and Learning Agents* (*ALA*), May 2022.
- 513. Zifan Xu, Yulin Zhang, Shahaf S. Shperberg, Reuth Mirsky, Yulin Zhan, Yuqian Jiang, Bo Liu, and **Peter Stone**. Model-based meta automatic curriculum learning. In *ICML workshop on Decision Awareness in Reinforcement Learning (DARL)*, July 2022.
- 514. Reuth Mirsky, Shahaf S. Shperberg, Yulin Zhang, Zifan Xu, Yuqian Jiang, Jiaxun Cui, and **Peter Stone**. Task factorization in curriculum learning. In *ICML workshop on Decision Awareness in Reinforcement Learning (DARL)*, July 2022.
- 515. Sai Kiran Narayanaswami, Swarat Chauduri, Moshe Vardi, and **Peter Stone**. Automating mechanism design with program synthesis. In *AAMAS workshop on Adaptive and Learning Agents* (ALA), May 2022.
- 516. Caroline Wang, Ishan Durugkar, Elad Liebman, and **Peter Stone**. Decentralized multi-agent reinforcement learning via distribution matching. In *AAMAS workshop on Adaptive and Learning Agents* (ALA), May 2022.
- 517. Parth Chonkar, Geethika Hemkumar, Huihai Wang, Daksh Dua, Shikhar Gupta, Yao-Cheng Chan, Justin Hart, Elliott Hauser, Reuth Mirsky, Joydeep Biswas, Junfeng Jiao, and **Peter Stone**. Look to my Lead: How does a leash affect perceptions of a quadruped robot? In *ICRA workshop on Social Navigation*, May 2022.
- 518. Hager Radi, Josiah P. Hanna, **Peter Stone**, and Matthew E. Taylor. Leveraging information about background music in human-robot interaction. In *NeurIPS Workshop on Deployable Decision Making in Embodied Systems* (DDM), Decmeber 2021.
- 519. Elad Liebman and **Peter Stone**. Leveraging information about background music in human-robot interaction. In *NeurIPS Workshop on Human and Machine Decisions*, Decmeber 2021.
- 520. Reuth Mirsky and **Peter Stone**. Intelligent disobedience and AI rebel agents in assistive robotics. In *ICSR workshop on Adaptive Social Interaction and MOVement for assistive and rehabilitation robotics* (ASIMOV), November 2021.
- 521. Zifan Xu, Xuesu Xiao, Garrett Warnell, Anirudh Nair, and **Peter Stone**. Machine learning methods for local motion planning: A study of end-to-end vs. parameter learning. In *IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*, October 2021.
- 522. Justin Hart, Reuth Mirsky, Xuesu Xiao, and **Peter Stone**. Incorporating gaze into social navigation. In *RSS Workshop on Social Robot Navigation*, July 2021.
- 523. Yuchen Cui, Akanksha Saran, Bo Liu, Scott Niekum, Stephen Giguere, and **Peter Stone**. Aux-AIRL: End-to-End Self-Supervised Reward Learning for Extrapolating beyond Suboptimal Demonstrations. In *ICML workshop on Self-Supervised Learning for Reasoning and Perception*, July 2021.
- 524. Bharath Masetty, Reuth Mirsky, Ashish D. Deshpande, Michael Mauk, and **Peter Stone**. Is the cerebellum a model-based reinforcement learning agent? In *AAMAS workshop on Adaptive and Learning Agents (ALA)*, May 2021.

- 525. Jennifer Suriadinata, William Macke, Reuth Mirsky, and **Peter Stone**. Reasoning about human behavior in ad hoc teamwork. In *AAMAS workshop on Adaptive and Learning Agents (ALA)*, May 2021.
- 526. Eddy Hudson, Garrett Warnell, and **Peter Stone**. Rail: A modular framework for reinforcement-learning-based adversarial imitation learning. In *AAMAS workshop on Autonomous Robots and Multi-robot Systems (ARMS)*, May 2021.
- 527. Yuchen Cui, Qiping Zhang, Sahil Jain, Alessandro Allievi, **Peter Stone**, Scott Niekum, and W. Bradley Knox. Reaction Modeling for Deriving General Task Information from Implicit Human Feedback. In *HRI workshop on Exploring Applications for Autonomous Nonverbal Human-Robot Interaction*, March 2021.
- 528. Haresh Karnan, Garrett Warnell, Xuesu Xiao, and **Peter Stone**. VOILA: Visual-Observation-only Imitation Learning for Autonomous navigation. In *AAAI Spring Symposium on Machine Learning for Mobile Robot Navigation in the Wild*, March 2021.
- 529. Minkyu Kim, Miguel Arduengo, Nick Walker, Yuqian Jiang, Hart W. Justin, **Peter Stone**, and Luis Sentis. Active Target Search for Autonomous Person Following. In *AAAI Spring Symposium on Machine Learning for Mobile Robot Navigation in the Wild*, March 2021.
- 530. Xuesu Xiao, Bo Liu, Garrett Warnell, and **Peter Stone**. Safe Learning from Hallucination for Navigation in the Wild. In *AAAI Spring Symposium on Machine Learning for Mobile Robot Navigation in the Wild*, March 2021.
- 531. William Macke, Reuth Mirsky, and **Peter Stone**. Expected Divergence Point of Plans in Ad Hoc Teamwork. In *Cooperative AI Workshop at NeurIPS 2020 (CoopAI)*, December 2020.
- 532. Brahma S. Pavse, Josiah P. Hanna, Ishan Durugkar, and **Peter Stone**. On sampling error in batch action-value prediction algorithms. In *In the Offline Reinforcement Learning Workshop at Neural Information Processing Systems (NeurIPS)*, *December 2020*., December 2020.
- 533. Daniel Perille, Abigail Truong, Xuesu Xiao, and **Peter Stone**. Benchmarking metric ground navigation. In *special session on Performance Evaluation & Benchmarking for Disaster Robotics at International Symposium on Safety, Security, and Rescue Robotics (SSRR)*, November 2020.
- 534. Xuesu Xiao, Bo Liu, and **Peter Stone**. Motion planners learned from geometric hallucination. In *IROS* workshop on bringing geometric methods to robot learning, optimization and control, October 2020.
- 535. Keya Ghonasgi, Reuth Mirsky, Bharath Masetty, Sanmit Narvekar, Adrian Haith, **Peter Stone**, and Ashish Deshpande. Leveraging reinforcement learning for human motor skill acquisition. In *13th International Workshop on Human Friendly Robotics (HFR)*, October 2020.
- 536. William Macke, Reuth Mirsky and **Peter Stone.** Query content in sequential one-shot multi-agent limited inquiries when communicating in ad hoc teamwork. In *Workshop on Distributed and Multi-Agent Planning (DMAP)*, October 2020.
- 537. Sanmit Narvekar and **Peter Stone**. Generalizing curricula for reinforcement learning. In *ICML Workshop on Lifelong Learning*, July 2020.
- 538. Yunshu Du, Garrett Warnell, Assefaw Gebremedhin, **Peter Stone**, and Matthew E. Taylor. Work-in-progress: Corrected Self Imitation learning via Demonstrations. In *AAMAS workshop on Adaptive and Learning Agents (ALA)*, May 2020.
- 539. Harel Yedidsion, Alkoby Shani, and **Peter Stone**. Sequential online chore division for autonomous vehicle convoy formation. In *AAMAS workshop on Games*, *Agents and Incentives Workshop (GAIW)*, May 2020.
- 540. Rishi Shah, Yuqian Jiang, Haresh Karnan, Gilberto Briscoe-Martinez, Dominick Mulder, Ryan Gupta, Rachel Schlossman, Marika Murphy, Justin Hart, Luis Sentis and **Peter Stone**. Solving Service Robot Tasks: UT Austin VillaHome 2019 Team Report. In *Proceedings of the AAAI Fall Symposium on Artificial Intelligence for Human-Robot Interaction*, November 2019.

- 541. Nick Walker, Yuqian Jiang, Maya Cakmak and **Peter Stone**. Desiderata for Planning Systems in General Purpose Service Robots. In *ICAPS workshop on Planning for Robotics (PlanRob)*, July 2019.
- 542. Faraz Torabi, Garrett Warnell, and **Peter Stone**. Generative adversarial imitation from observation. In *ICML Workshop on Imitation, Intent, and Interaction (I3) Workshop*, June 2019.
- 543. Faraz Torabi, Sean Geiger, Garrett Warnell, and **Peter Stone**. Sample-efficient adversarial imitation learning from observation. In *ICML Workshop on Imitation, Intent, and Interaction (I3) Workshop*, June 2019.
- 544. Felipe Leno Da Silva and **Peter Stone**. Distributional Reinforcement Learning Applied to Robot Soccer Simulation. In *AAMAS Workshop on Adaptive and Learning Agents (ALA)*, May 2019.
- 545. Shih-Yun Lo, Shani Alkoby and **Peter Stone**. Robust Motion Planning and Safety Benchmarking in Human Workspaces. In *AAAI workshop on Safe AI*, January 2019.
- 546. Yuqian Jiang, Nick Walker, Minkyu Kim, Nicolas Brissonneau, Daniel S. Brown, Justin W. Hart, Scott Niekum, Luis Sentis, and **Peter Stone**. LAAIR: A layered architecture for autonomous interactive robots. In *Proceedings of the AAAI Fall Symposium on Reasoning and Learning in Real-World Systems for Long-Term Autonomy*, October 2018.
- 547. Prabhat Nagarajan, Garrett Warnell, and **Peter Stone**. Deterministic implementations for reproducibility in deep reinforcement learning. In *2nd ICML Workshop on Reproducibility in Machine Learning*, July 2018.
- 548. Jake Menashe and **Peter Stone**. State abstraction synthesis for discrete models of continuous domains. In *Proceedings of the AAAI Spring Symposium on Data Efficient Reinforcement Learning (DERL)*, March 2018.
- 549. Ishan Durugkar and **Peter Stone** TD learning with constrained gradients. In *NIPS Deep Reinforcement Learning Symposium*, December 2017.
- 550. Josiah Hanna and **Peter Stone**. Towards a data efficient off-policy policy gradient. In *AAAI Spring Symposium on Data Efficient Reinforcement Learning*, March 2018.
- 551. Santiago Gonzalez, Vijay Chidambaram, Jivko Sinapov, and **Peter Stone**. CC-Log: Drastically reducing storage requirements for robots using classification and compression. In *The 9th USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage)*, July 2017.
- 552. Shih-Yun Lo, Benito Fernandez, and **Peter Stone**. Iterative Human-Aware Mobile Robot Navigation. In *RSS Workshop on Human-Centered Robotics: Interaction, Physiological Integration, and Autonomy*, July 2017.
- 553. Shiqi Zhang, Jivko Sinapov, Suhua Wei, and **Peter Stone**. Robot behavioral exploration and multimodal perception using POMDPs. In *Proceedings of the AAAI Spring Symposium on Interactive Multi-Sensory Perception for Embodied Agents*, March 2017.
- 554. Matthew Hausknecht and **Peter Stone**. Grounded semantic networks for learning shared communication protocols. In *NIPS workshop on Deep Reinforcement Learning*, December 2016.
- 555. Patrick MacAlpine, Elad Liebman, and **Peter Stone**. Adaptation of surrogate tasks for bipedal walk optimization. In *GECCO Surrogate-Assisted Evolutionary Optimisation (SAEOpt) Workshop*, July 2016.
- 556. Matthew Hausknecht and **Peter Stone**. On-policy vs. off-policy updates for deep reinforcement learning. In *IJCAI 2016 Workshop on Deep Reinforcement Learning*, July 2016.
- 557. Ginevra Gaudioso, Matteo Leonetti, and **Peter Stone**. State aggregation through reasoning in answer set programming. In *Proceedings of the IJCAI Workshop on Autonomous Mobile Service Robots* (WSR), July 2016.

- 558. Josiah P. Hanna, Michael Albert, Donna Chen, and **Peter Stone**. Minimum cost matching for autonomous carsharing. In *Proceedings of the 9th IFAC Symposium on Intelligent Autonomous Vehicles (IAV)*, June 2016.
- 559. Matthew Hausknecht, Prannoy Mupparaju, Sandeep Subramanian, Shivaram Kalyanakrishnan, and **Peter Stone**. Half field offense: An environment for multiagent learning and ad hoc teamwork. In *AAMAS Workshop on Adaptive and Learning Agents (ALA)*, May 2016.
- 560. Matthew Hausknecht, Yilun Chen, and **Peter Stone**. Deep imitation learning for parameterized action spaces. In *AAMAS Workshop on Adaptive and Learning Agents (ALA)*, May 2016.
- 561. Jonathan Grizou, Samuel Barrett, Manuel Lopes, and **Peter Stone**. Collaboration in ad hoc teamwork: ambiguous tasks, roles, and communication. In *AAMAS Workshop on Adaptive and Learning Agents* (*ALA*), May 2016.
- 562. Matthew Hausknecht and **Peter Stone**. Deep recurrent Q-learning for partially observable MDPs. In *AAAI Fall Symposium on Sequential Decision Making for Intelligent Agents (SDMIA)*, November 2015.
- 563. Matthew Hausknecht and **Peter Stone**. The impact of determinism on learning atari 2600 games. In *AAAI Workshop on Learning for General Competency in Video Games*, January 2015.
- 564. Piyush Khandelwal and **Peter Stone**. Multi-robot human guidance using topological graphs. In *AAAI* Spring 2014 Symposium on Qualitative Representations for Robots, March 2014.
- 565. Fangkai Yang, Piyush Khandelwal, Matteo Leonetti, and **Peter Stone**. Planning in answer set programming while learning action costs for mobile robots. In *AAAI Spring 2014 Symposium on Knowledge Representation and Reasoning in Robotics*, March 2014.
- 566. Jacob Menashe, Katie Genter, Samuel Barrett, and **Peter Stone**. UT Austin Villa 2013: Advances in vision, kinematics, and strategy. In *The Eighth Workshop on Humanoid Soccer Robots at Humanoids* 2013, October 2013.
- 567. Katie Genter, Noa Agmon, and **Peter Stone**. Improving efficiency of leading a flock in ad hoc teamwork settings. In *AAMAS Workshop on Autonomous Robots and Multirobot Systems* (ARMS), May 2013.
- 568. Patrick MacAlpine, Elad Liebman, and **Peter Stone**. Simultaneous learning and reshaping of an approximated optimization task. In *AAMAS Workshop on Adaptive and Learning Agents* (ALA), May 2013.
- 569. Doran Chakraborty, Noa Agmon, and **Peter Stone**. Targeted opponent modeling of memory-bounded agents. In *AAMAS Workshop on Adaptive and Learning Agents* (ALA), May 2013.
- 570. Patrick MacAlpine and **Peter Stone**. Using dynamic rewards to learn a fully holonomic bipedal walk. In *AAMAS Workshop on Adaptive and Learning Agents* (ALA), June 2012.
- 571. Samuel Barrett, **Peter Stone**, Sarit Kraus, and Avi Rosenfeld. Learning teammate models for ad hoc teamwork. In *AAMAS Workshop on Adaptive and Learning Agents* (ALA), June 2012.
- 572. W. Bradley Knox, Matthew Taylor, and **Peter Stone**. Understanding Human Teaching Modalities in Reinforcement Learning Environments: A Preliminary Report In *IJCAI workshop on Agents Learning Interactively from Human Teachers (ALIHT)*, July 2011.
- 573. W. Bradley Knox and **Peter Stone**. Augmenting Reinforcement Learning with Human Feedback. In *ICML Workshop on Imitation Learning*, July 2011.
- 574. Paul Scerri, Balajee Kannan, Pras Velagapudi, Kate Macarthur, **Peter Stone**, Matthew E. Taylor, John Dolan, Alessandro Farinelli, Archie Chapman, Bernadine Dias, and George Kantor. Flood disaster mitigation: A real-world challenge problem for multi-agent unmanned surface vehicles. In *Proceedings of the AAMAS workshop on Autonomous Robots and Multirobot Systems (ARMS)*, May 2011.

- 575. Samuel Barrett and **Peter Stone**. Ad Hoc Teamwork Modeled with Multi-armed Bandits: An Extension to Discounted Infinite Rewards. In *Proceedings of the AAMAS workshop on Adaptive Learning Agents (ALA)*, May 2011.
- 576. Shivaram Kalyanakrishnan and **Peter Stone**. On learning with imperfect representations. In *Proceedings of the 2011 IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning*, April 2011.
- 577. Shimon Whiteson, Brian Tanner, Matthew E. Taylor, and **Peter Stone**. Protecting against evaluation overfitting in empirical reinforcement learning. In *IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning (ADPRL)*, April 2011.
- 578. W. Bradley Knox and **Peter Stone**. Reinforcement learning with human feedback in mountain car. In *AAAI Spring 2011 Symposium on Bridging the Gaps in Human-Agent Collaboration*, March 2011.
- 579. Samuel Barrett, Katie Genter, Todd Hester, Michael Quinlan, and **Peter Stone**. Controlled kicking under uncertainty. In *The Fifth Workshop on Humanoid Soccer Robots at Humanoids 2010*, December 2010.
- 580. Piyush Khandelwal, Matthew Hausknecht, Juhyun Lee, Aibo Tian, and **Peter Stone**. Vision calibration and processing on a humanoid soccer robot. In *The Fifth Workshop on Humanoid Soccer Robots at Humanoids 2010*, December 2010.
- 581. Adam Setapen, Michael Quinlan, and **Peter Stone**. MARIOnET: motion acquisition for robots through iterative online evaluative training. In *AAMAS 2010 Workshop on Agents Learning Interactively from Human Teachers (ALIHT)*, May 2010.
- 582. Samuel Barret, Matthew E. Taylor, and **Peter Stone**. Transfer learning for reinforcement learning on a physical robot In *AAMAS 2010 Workshop on Adaptive and Learning Agents (ALA)*, May 2010.
- 583. Tobias Jung, Mazda Ahmadi, and **Peter Stone**. Connectivity-based localization in robot networks. In *International Workshop on Robotic Wireless Sensor Networks (WSN)*, June 2009.
- 584. Todd Hester and **Peter Stone**. An empirical comparison of abstraction in models of markov decision processes. In *Proceedings of the ICML/UAI/COLT Workshop on Abstraction in Reinforcement Learning*, June 2009.
- 585. Ian Fasel, Michael Quinlan, and **Peter Stone**. A task specification language for bootstrap learning. In *AAAI Spring 2009 Symposium on Agents that Learn from Human Teachers*, March 2009.
- 586. W. Bradley Knox, Ian Fasel, and **Peter Stone**. Design principles for creating human-shapable agents. In *AAAI Spring 2009 Symposium on Agents that Learn from Human Teachers*, March 2009.
- 587. Mohan Sridharan and **Peter Stone**. Comparing two action planning approaches for color learning on a mobile robot. In *VISAPP Workshop on Robotic Perception*, January 2008.
- 588. Kurt Dresner and **Peter Stone**. Learning policy selection for autonomous intersection management. In *AAMAS workshop on Adaptive and Learning Agents (ALA)*, May 2007.
- 589. Matthew E. Taylor and **Peter Stone**. Representation transfer for reinforcement learning. In *AAAI* 2007 Fall Symposium on Computational Approaches to Representation Change during Learning and Development, November 2007.
- 590. Matthew E. Taylor, Gregory Kuhlmann, and **Peter Stone**. Accelerating search with transferred heuristics. In *ICAPS workshop on AI Planning and Learning*, September 2007.
- 591. Nicholas K. Jong and **Peter Stone**. Model-based exploration in continuous state spaces. In *The Seventh Symposium on Abstraction, Reformulation, and Approximation*, July 2007.
- 592. Nicholas Jong and **Peter Stone**. Kernel-based models for reinforcement learning in continuous state spaces. In *ICML workshop on Kernel Machines and Reinforcement Learning*, June 2006.

- 593. Matthew E. Taylor, Shimon Whiteson, and **Peter Stone**. Transfer learning for policy search methods. In *ICML workshop on Structural Knowledge Transfer for Machine Learning*, June 2006.
- 594. Bikramjit Banerjee, Gregory Kuhlmann, and **Peter Stone**. Value function transfer for general game playing. In *ICML workshop on Structural Knowledge Transfer for Machine Learning*, June 2006.
- 595. Daniel Stronger and **Peter Stone**. Expectation-based vision for precise self-localization on a mobile robot. In *AAAI workshop on Cognitive Robotics*, July 2006.
- 596. Mazda Ahmadi and **Peter Stone**. Instance-based action models for fast action planning. In *AAAI* workshop on Cognitive Robotics, July 2006.
- 597. **Peter Stone**, David Pardoe, and Mark VanMiddlesworth. TacTex-05: An adaptive agent for TAC SCM. In *AAMAS 2006 Joint workshop on Trading Agent Design and Analysis & Agent Mediated Electronic Commerce VIII*, May 2006.
- 598. Kurt Dresner and **Peter Stone**. Human-usable and emergency vehicle-aware control policies for autonomous intersection management. In *AAMAS 2006 Workshop on Agents in Traffic and Transportation*, May 2006.
- 599. Mazda Ahmadi and **Peter Stone**. Keeping in touch: A distributed check for biconnected structure by homogeneous robots. In *The 8th International Symposium on Distributed Autonomous Robotic Systems*, July 2006.
- 600. Subramanian Ramamoorthy, Harish Subramanian, **Peter Stone**, and Benjamin Kuipers. A qualitative multiple-model approach to the design of autonomous financial trading agents. In *NIPS 2005 workshop on Machine Learning in Finance*, December 2005.
- 601. Nicholas K. Jong and **Peter Stone**. Bayesian models of nonstationary markov decision problems. In *IJCAI 2005 workshop on Planning and Learning in A Priori Unknown or Dynamic Domains*, August 2005.
- 602. Nicholas K. Jong and **Peter Stone**. Towards learning to ignore irrelevant state variables. In *The AAAI-2004 Workshop on Learning and Planning in Markov Processes Advances and Challenges*, July 2004.
- 603. Gregory Kuhlmann, **Peter Stone**, Raymond Mooney, and Jude Shavlik. Guiding a reinforcement learner with natural language advice: Initial results in RoboCup soccer. In *The AAAI-2004 Workshop on Supervisory Control of Learning and Adaptive Systems*, July 2004.
- 604. János A. Csirik, Michael L. Littman, Satinder Singh, and **Peter Stone**. FAucS: An FCC spectrum auction simulator for autonomous bidding agents. In Ludger Fiege, Gero Mühl, and Uwe Wilhelm, editors, *Electronic Commerce: Proceedings of the Second International Workshop*. Springer Verlag, Heidelberg, 2001.
- 605. Michael L. Littman and **Peter Stone**. Implicit Negotiation in Repeated Games. In *Proceedings of The Eighth International Workshop on Agent Theories, Architectures, and Languages* (ATAL), 2001.
- 606. M. Asada, A. Birk, E. Pagello, M. Fujita, I. Noda, S. Tadokoro D. Duhaut, P. Stone, M. Veloso, T. Balch, H. Kitano, and B. Thomas. Progress in RoboCup soccer research in 2000. In *Proceedings of the 2000 International Symposium on Experimental Robotics*, Honolulu, 2000.
- 607. Patrick Riley, **Peter Stone**, and Manuela Veloso. Layered disclosure: Revealing agents' internals. In *The Seventh International Workshop on Agent Theories, Architectures, and Languages* (ATAL), 2000.
- 608. **Peter Stone**. Multiagent learning for autonomous spacecraft constellations. In *Proceedings of the NASA Workshop on Planning and Scheduling for Space*, 1997.
- 609. Sorin Achim, **Peter Stone**, and Manuela Veloso. Building a dedicated robotic soccer system. In *Proceedings of the IROS-96 Workshop on RoboCup*, pages 41–48, Osaka, Japan, November 1996.

- 610. Mike Bowling, **Peter Stone**, and Manuela Veloso. Predictive memory for an inaccessible environment. In *Proceedings of the IROS-96 Workshop on RoboCup*, pages 28–34, Osaka, Japan, November 1996.
- 611. **Peter Stone** and Manuela Veloso. Using machine learning in the soccer server. In *Proceedings of the IROS-96 Workshop on RoboCup*, pages 19–27, Osaka, Japan, November 1996.
- 612. **Peter Stone**, Manuela Veloso, and Sorin Achim. Collaboration and learning in robotic soccer. In *Proceedings of the Micro-Robot World Cup Soccer Tournament*, pages 26–37, Taejon, Korea, November 1996. IEEE Robotics and Automation Society.
- 613. **Peter Stone** and Manuela Veloso. Using testing to iteratively improve training. In *Working Notes of the AAAI 1995 Fall Symposium on Active Learning*, pages 110–111, Boston, MA, November 1995.
- 614. **Peter Stone** and Manuela Veloso. Learning to solve complex planning problems finding useful auxiliary problems. In *Technical Report of the AAAI 1994 Fall Symposium on Planning and Learning: On to Real Applications*, pages 137–141, New Orleans, LA, November 1994.

Magazine Articles

- 615. Peter Stone. Artificial intelligence: Looking forward 15 years. IEEE Intelligent Systems, 2025.
- 616. Xuesu Xiao, Zifan Xu, Aniket Datar, Garrett Warnell, Peter Stone, Joshua Julian Damanik, Jaewon Jung, Chala Adane Deresa, Than Duc Huy, Chen Jinyu, Chen Yichen, Joshua Adrian Cahyono, Jingda Wu, Longfei Mo, Mingyang Lv, Bowen Lan, Qingyang Meng, Weizhi Tao, and Li Cheng. Autonomous ground navigation in highly constrained spaces: Lessons learned from the third BARN challenge at ICRA 2024 [competitions]. IEEE Robotics & Automation Magazine, 2024.
- 617. Eric Horvitz, Vincent Conitzer, Sheila McIlraith, and **Peter Stone**. Now, later, and lasting: 10 priorities for AI research, policy, and practice. *Communications of the ACM*, 67(6):39–40, May 2024.
- 618. **Peter Stone**. Will robots triumph over world cup winners by 2050? *IEEE Spectrum*, 60(7):40–9, July 2023.
- 619. Xuesu Xiao, Zifan Xu, Garrett Warnell, **Peter Stone**, Ferran Bebelli Guinjoan, Romulo T. Rodrigues, Herman Bryunickx, Hanjaya Mandala, Guilherme Christmann, Jose Luis Blanco-Claraco, and Shravan Somashekara Rai. Autonomous ground navigation in highly constrained spaces: lessons learned from the 2nd BARN challenge at ICRA 2023. *IEEE Robotics & Automation Magazine*, 30(4):91–97, Dec. 2023.
- 620. Xuesu Xiao, Zifan Xu, Zizhao Wang, Yunlong Song, Garrett Warnell, **Peter Stone**, Tingnan Zhang, Shravan Ravi, Gary Wang, Haresh Karnan, Joydeep Biswas, Nicholas Mohammad, Lauren Bramblett, Rahul Peddi, Nicola Bezzo, Zhanteng Xie, and Philip Dames. Autonomous ground navigation in highly constrained spaces: Lessons learned from the barn challenge at icra 2022. *IEEE Robotics & Automation Magazine*, 29(4):148–56, Dec. 2022.
- 621. Peter R. Wurman, **Peter Stone**, and Michael Sprannger. Challenges and opportunities of applying reinforcement learning to autonomous racing. *IEEE Intelligent Systems*, 37(3):20–3, May-June 2022.
- 622. **Peter Stone**, Luca Iocchi, Flavio Tonidandel, and Changjiu Zhou. RoboCup 2021 worldwide: a successful robotics competition during a pandemic. *IEEE Robotics & Automation Magazine*, 28(4):114–19, December 2021.
- 623. Barbara J. Grosz and **Peter Stone**. A century long commitment to assessing artificial intelligence and its impact on society. *Communications of the ACM*, 61(12), December 2018.
- 624. Tsz-Chiu Au, Bikramjit Banerjee, Prithviraj Dasgupta, **Peter Stone**. Guest Editors' Introduction: Multirobot Systems. In *IEEE Intelligent Systems*, pages 3–5, November/December 2017.

- 625. , Katie Genter, Patrick MacAlpine, Jacob Menashe, Josiah Hannah, Elad Liebman, Sanmit Narvekar, Ruohan Zhang, and **Peter Stone**. UT Austin Villa: project-driven research in AI and Robotics. *IEEE Intelligent Systems*, 31:2:94–101, 2016.
- 626. Eric Eaton, Tom Dietterich, Maria Gini, Barbara J. Grosz, Charles L. Isbell, Subbarao Kambhamp, Michael Littman, Francesca Rossi, Stuart Russell, **Peter Stone**, Toby Walsh, and Michael Wooldridge. Who Speaks for AI? in *AI Matters*, 2(2), December 2015.
- 627. Daniele Nardi, Itsuki Noda, Fernando Ribeiro, **Peter Stone**, Oskar von Stryk, and Manuela Veloso. RoboCup soccer leagues. *AI Magazine*, 35(3):77–85, 2014.
- 628. David Pardoe and **Peter Stone**. Designing Adaptive Trading Agents. *ACM SIGecom Exchanges*, 10(2), June 2011.
- 629. Matthew E. Taylor and **Peter Stone**. An Introduction to Inter-task Transfer for Reinforcement Learning. *AI Magazine*, 32(1):15–20, March 2011.
- 630. David Pardoe and **Peter Stone**. Developing Adaptive Auction Mechanisms. *ACM SIGecom Exchanges*, 5(3):1–10, April 2005.
- 631. David Pardoe and **Peter Stone**. TacTex-03: A Supply Chain Management Agent. *ACM SIGecom Exchanges*, 4(3):19–28, Winter 2004.
- 632. Amy Greenwald, Nicholas R. Jennings, and **Peter Stone**. Guest Editors' Introduction: Agents and Markets. *IEEE Intelligent Systems*, 18(6):12–14, November/December 2003.
- 633. Manuela Veloso, Tucker Balch, Peter Stone, Hiroaki Kitano, Fuminori Yamasaki, Ken Endo, Minoru Asada, M. Jamzad, B. S. Sadjad, V. S. Mirrokni, M. Kazemi, H. Chitsaz, A. Heydarnoori, M. T. Hajiaghai, and E. Chiniforooshan. RoboCup-2001: The fifth robotic soccer world championships. AI Magazine, 23(1):55–68, 2002.
- 634. **Peter Stone** (editor), Minoru Asada, Tucker Balch, Raffaelo D'Andrea, Masahiro Fujita, Bernhard Hengst, Gerhard Kraetzschmar, Pedro Lima, Nuno Lau, Henrik Lund, Daniel Polani, Paul Scerri, Satoshi Tadokoro, Thilo Weigel, and Gordon Wyeth. RoboCup-2000: The fourth robotic soccer world championships. *AI Magazine*, 22(1), 2001.
- 635. Amy Greenwald and **Peter Stone**. Autonomous bidding agents in the trading agent competition. *IEEE Internet Computing*, 5(2), 2001.
- 636. Silvia Coradeschi, Lars Karlsson, **Peter Stone**, Tucker Balch, Gerhard Kraetzschmar, and Minoru Asada. Overview of RoboCup-99. *AI Magazine*, 21(3), 2000.
- 637. Manuela Veloso, **Peter Stone**, Kwun Han, and Sorin Achim. CMUnited: A team of robotic soccer agents collaborating in an adversarial environment. *Crossroads*, 4.3, February 1998.

Multimedia

- 638. Manuela Veloso and **Peter Stone**. Video: Robocup robot soccer history 1997 2011, October 2012. Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), October 2012. Available from https://www.youtube.com/watch?v=WLOv2AFAZhc **Finalist for Best Video Award.**
- 639. W. Bradley Knox, Juhyun Lee, and **Peter Stone**. Person recognition on a segway robot: A video of UT Austin Villa Robocup@Home 2007 finals demonstration. In *IEEE International Conference on Robotics and Automation*, May 2008.

 Available from www.cs.utexas.edu/~AustinVilla/?p=athome

Unrefereed Publications

- 640. Elad Liebman and Peter Stone. Utilizing mood-inducing background music in human-robot interaction. *arXiv e-Prints* 2308.14269, September 2023.
- 641. Bo Liu, Yuqian Jiang, Xiaohan Zhang, Qiang Liu, Shiqi Zhang, Joydeep Biswas, and **Peter Stone**. LLM+P: empowering large language models with optimal planning proficiency. *arXiv e-Prints* 2304.11477, April 2023.
- 642. Justin Hart, Harel Yedidsion, Yuqian Jiang, Nick Walker, Rishi Shah, Jesse Thomason, Aishwarya Padmakumar, Rolando Fernandez, Jivko Sinapov, Raymond Mooney, and **Peter Stone**. Interaction and autonomy in RoboCupHome and Building-Wide Intelligence. In *Proceedings of the AAAI Fall Symposium on Interactive learning in Artificial Intelligence for Human-Robot Interaction (AI-HRI)*, October 2018.
- 643. Kara Kockelman, Stephen Boyles, Purser Sturgeon, Christian Claudel, Lisa Loftus-Otway, Wendy Wagner, Duncan Stewart, Guni Sharon, Michael Albert, **Peter Stone**, Josiah Hanna, Yantao Huang, Krishna Murthy Gurumurthy, Dongxu He, Abduallah Mohamed, Rahul Patel, Tian Lei, Michele Simoni, and Sadegh Yarmohammadisatri. Bringing smart transport to texans: Ensuring the benefits of a connected and autonomous transport system in texas final report. Technical Report 0-6838-3, The University of Texas at Austin Center for Transportation Research, July 2018.
- 644. Kara Kockelman, Stephen Boyles, **Peter Stone**, Dan Fagnant, Rahul Patel, Michael W. Levin, Guni Sharon, Michael Simoni, Michael Albert, Hagen Fritz, Rebecca Hutchinson, Prateek Bansal, Gelb Domnenko, Pavle Bujanovic, Bumsik Kim, Elaham Pourrahmani, Sudesh Agrawal, Tianxin Li, Josiah Hanna, Aqshems Nichols, and Jia Li. An assessment of autonomous vehicles: traffic impacts and infrastructure needs final report. Technical Report 0-6847-1, The University of Texas at Austin Center for Transportation Research, March 2017.
- 645. Kara Kockelman, Stephen Boyles, Paul Avery, Christian Claudel, Lisa Loftus-Otway, Daniel Fagnant, Prateek Bansal, Michael Levin, Yong Zhao, Jun Liu, Lewis Clements, Wendy Wagner, Duncan Stewart, Guni Sharon, Michael Albert, **Peter Stone**, Josiah Hanna, Rahul Patel, Hagen Fritz, Tejas Choudhary, Tianxin Li, Aqshems Nichols, Kapil Sharma, and Michele Simoni. Bringing smart transport to texans: Ensuring the benefits of a connected and autonomous transport system in texas final report. Technical Report 0-6838-2, The University of Texas at Austin Center for Transportation Research, November 2016.
- 646. Decebal Constantin Mocanu, Maria Torres Vega, Eric Eaton, **Peter Stone**, and Antonio Liotta. Online contrastive divergence with generative replay: Experience replay without storing data. *arXiv e-Prints* 1610.05555, October 2016.
- 647. **Peter Stone**, Rodney Brooks, Erik Brynjolfsson, Ryan Calo, Oren Etzioni, Greg Hager, Julia Hirschberg, Shivaram Kalyanakrishnan, Ece Kamar, Sarit Kraus, Kevin Leyton-Brown, David Parkes, William Press, AnnaLee Saxenian, Julie Shah, Milind Tambe, and Astro Teller. Artificial Intelligence and life in 2030. *One Hundred Year Study on Artificial Intelligence: Report of the 2015-2016 Study Panel*, Stanford University, Stanford, CA, September 2016.
- 648. Patrick MacAlpine, Daniel Urieli, Samuel Barrett, Shivaram Kalyanakrishnan, Francisco Barrera, Adrian Lopez-Mobilia, Nicolae Ştiurcă, Victor Vu, and **Peter Stone**. UT Austin Villa 2011 3D Simulation Team report. Technical Report AI11-10, The University of Texas at Austin, Department of Computer Science, AI Laboratory, December 2011.
- 649. Samuel Barrett, Katie Genter, Matthew Hausknecht, Todd Hester, Piyush Khandelwal, Juhyun Lee, Michael Quinlan, Aibo Tian, **Peter Stone**, and Mohan Sridharan. Austin Villa 2010 standard platform team report. Technical Report UT-AI-TR-11-01, The University of Texas at Austin, Department of Computer Sciences, AI Laboratory, January 2011.

- 650. Todd Hester, Michael Quinlan, **Peter Stone**, and Mohan Sridharan. TT-UT Austin Villa 2009: Naos across Texas. Technical Report UT-AI-TR-09-08, The University of Texas at Austin, Department of Computer Science, AI Laboratory, December 2009.
- 651. Shivaram Kalyanakrishnan, Yinon Bentor, and **Peter Stone**. The UT Austin Villa 3D Simulation Soccer Team 2008. Technical Report AI09-01, The University of Texas at Austin, Department of Computer Sciences, AI Laboratory, February 2009.
- 652. Todd Hester, Michael Quinlan, and **Peter Stone**. UT Austin Villa 2008: Standing on Two Legs. Technical Report UT-AI-TR-08-8, The University of Texas at Austin, Department of Computer Sciences, AI Laboratory, November 2008.
- 653. Shivaram Kalyanakrishnan and **Peter Stone**. UT Austin Villa 3D simulation soccer team 2008. In *RoboCup Symposium CD Proceedings*, 2008.
- 654. Peter Stone Todd Hester, Michael Quinlan. UT Austin Villa 2008 team description paper for the standard platform league. In *RoboCup Symposium CD Proceedings*, 2008.
- 655. **Peter Stone** and Neda Shahidi. The UT Austin Villa 2008 mixed reality team. In *RoboCup Symposium CD Proceedings*, 2008.
- 656. **Peter Stone**, Patrick Beeson, Tekin Mericli, and Ryan Madigan. DARPA urban challenge technical report: Austin Robot Technology, June 2007.
- 657. Shivaram Kalyanakrishnan and **Peter Stone**. The UT Austin Villa 3D Simulation Soccer Team 2007. Technical Report AI-07-348, The University of Texas at Austin, Department of Computer Sciences, AI Laboratory, September 2007.
- 658. Shimon Whiteson, Matthew E. Taylor, and **Peter Stone**. Adaptive tile coding for value function approximation. Technical Report AI-TR-07-339, University of Texas at Austin, 2007.
- 659. **Peter Stone**, William B. Knox, Juhyun Lee, Youngin Shin, and Mohan Sridharan. The UT Austin Villa 2007 team description paper. In *RoboCup Symposium CD Proceedings*, 2007.
- 660. **Peter Stone**, Todd Hester, Tekin Meriçli, Neda Shahidi, and Mohan Sridharan. The UT Austin Villa 2007 four-legged team description paper. In *RoboCup Symposium CD Proceedings*, 2007.
- 661. **Peter Stone**, Peggy Fidelman, Nate Kohl, Gregory Kuhlmann, Tekin Mericli, Mohan Sridharan, and Shao en Yu. The UT Austin Villa 2006 RoboCup four-legged team. Technical Report UT-AI-TR-06-337, The University of Texas at Austin, Department of Computer Sciences, AI Laboratory, December 2006.
- 662. Daniel Stronger and **Peter Stone**. Polynomial regression with automated degree: A function approximator for autonomous agents. Technical Report UT-AI-TR-06-329, The University of Texas at Austin, Department of Computer Sciences, AI Laboratory, April 2006.
- 663. **Peter Stone**, Kurt Dresner, Peggy Fidelman, Nate Kohl, Gregory Kuhlmann, Mohan Sridharan, and Daniel Stronger. The UT Austin Villa 2005 RoboCup four-legged team. Technical Report UT-AI-TR-05-325, The University of Texas at Austin, Department of Computer Sciences, AI Laboratory, November 2005.
- 664. Kurt Dresner and **Peter Stone**. Multiagent traffic management: Driver agent improvements and a protocol for intersection control. Technical Report UT-AI-TR-04-315, The University of Texas at Austin, Department of Computer Sciences, AI Laboratory, May 2005.
- 665. **Peter Stone**, Kurt Dresner, Peggy Fidelman, Nicholas K. Jong, Nate Kohl, Gregory Kuhlmann, Mohan Sridharan, and Daniel Stronger. The UT Austin Villa 2004 RoboCup four-legged team: Coming of age. Technical Report UT-AI-TR-04-313, The University of Texas at Austin, Department of Computer Sciences, AI Laboratory, October 2004.

- 666. Nicholas K. Jong and **Peter Stone**. Towards employing PSRs in a continuous domain. Technical Report UT-AI-TR-04-309, The University of Texas at Austin, Department of Computer Sciences, AI Laboratory, February 2004.
- 667. Peter Stone, Kurt Dresner, Selim T. Erdogan, Peggy Fidelman, Nicholas K. Jong, Nate Kohl, Gregory Kuhlmann, Ellie Lin, Mohan Sridharan, Daniel Stronger, and Gurushyam Hariharan. The UT Austin Villa 2003 Four-Legged Team. In Daniel Polani, Brett Browning, Andrea Bonarini, and Kazuo Yoshida, editors, RoboCup-2003: Robot Soccer World Cup VII. Springer Verlag, Berlin, 2004.
- 668. **Peter Stone**. ATTUnited-2001: Using Heterogeneous Players. In Andreas Birk, Silvia Coradeschi, and Satoshi Tadokoro, editors, *RoboCup-2001: Robot Soccer World Cup V*, volume 2377 of *Lecture Notes in Artificial Intelligence*, pages 495–98. Springer Verlag, Berlin, 2002.
- 669. Patrick Riley, **Peter Stone**, David McAllester, and Manuela Veloso. ATT-CMUnited-2000: Third place finisher in the RoboCup-2000 simulator league. In P. Stone, T. Balch, and G. Kraetzschmar, editors, *RoboCup-2000: Robot Soccer World Cup IV*, volume 2019 of *Lecture Notes in Artificial Intelligence*. Springer Verlag, Berlin, 2001.
- 670. **Peter Stone**. Layered learning in multi-agent systems. Ph.D. thesis, Computer Science Department, Carnegie Mellon University, Pittsburgh, PA, December 1998. Technical Report CMU-CS-98-187.
- 671. Manuela Veloso, Hiroaki Kitano, Enrico Pagello, Gerhard Kraetzschmar, Peter Stone, Tucker Balch, Minoru Asada, Silvia Coradeschi, Lars Karlsson, and Masahiro Fujita. Overview of RoboCup-99. In M. Veloso, E. Pagello, and H. Kitano, editors, RoboCup-99: Robot Soccer World Cup III, volume 1856 of Lecture Notes in Artificial Intelligence, pages 1–34. Springer Verlag, Berlin, 2000.
- 672. Peter Stone, Patrick Riley, and Manuela Veloso. The CMUnited-99 champion simulator team. In M. Veloso, E. Pagello, and H. Kitano, editors, *RoboCup-99: Robot Soccer World Cup III*, volume 1856 of *Lecture Notes in Artificial Intelligence*, pages 35–48. Springer Verlag, Berlin, 2000. Also in *AI Magazine*, 21(3), 2000.
- 673. **Peter Stone**, Manuela Veloso, and Patrick Riley. The CMUnited-98 champion simulator team. In M. Asada and H. Kitano, editors, *RoboCup-98: Robot Soccer World Cup II*, volume 1604 of *Lecture Notes in Artificial Intelligence*, pages 61–76. Springer Verlag, 1999.
- 674. **Peter Stone** and Manuela Veloso. The CMUnited-97 simulator team. In Hiroaki Kitano, editor, *RoboCup-97: Robot Soccer World Cup I*, volume 1395 of *Lecture Notes in Artificial Intelligence*, pages 387–397. Springer Verlag, Berlin, 1998.
- 675. Manuela Veloso, **Peter Stone**, Kwun Han, and Sorin Achim. The CMUnited-97 small-robot team. In Hiroaki Kitano, editor, *RoboCup-97: Robot Soccer World Cup I*, volume 1395 of *Lecture Notes in Artificial Intelligence*, pages 242–256. Springer Verlag, Berlin, 1998.

 Also in *AI Magazine*, 19(3):61–69, 1998.
- 676. Emiel Corten, Klaus Dorer, Fredrik Heintz, Kostas Kostiadis, Johan Kummeneje, Helmut Myritz, Itsuki Noda, Patrick Riley, **Peter Stone**, and Travlex Yeap. Soccer server manual, version 5.0. Technical Report RoboCup-1999-001, RoboCup, 1999.
- 677. Gregg Rabideau, Steve Chien, Tobias Mann, Curt Eggemeyer, and **Peter Stone**. DCAPS User's Manual. Technical Document D-13741, Jet Propulsion Laboratory (JPL), 1996.

PROFESSIONAL MEMBERSHIPS

- Association for the Advancement of Artificial Intelligence, Fellow (AAAI)
- Institute of Electrical and Electronic Engineers, Fellow (IEEE)
- American Association for the Advancement of Science, Fellow (AAAS)
- Association for Computing Machinery, Fellow (ACM)

PRESS

Interviewed and quoted regarding research several times on television, on radio, and in magazines and newspapers including CNN, NPR, BBC, The New York Times, Wall Street Journal, USA Today, Pittsburgh Post-Gazette, Scientific American, and Austin American Statesman. Appeared on PBS *Scientific American Frontiers* hosted by Alan Alda.

Links to stories available at http://www.cs.utexas.edu/users/pstone/press.shtml

PERSONAL

Married, three children — born 1998, 2000, 2002. Citizenship: U.S.

- Violin performed with the CMU philharmonic in Carnegie Hall, NY.
- Soccer played in a semi-professional league, tried out for Major League Soccer.
- Languages English (native), French and Hebrew (conversational).