Peter Herald Stone

Professor
Dept. of Computer Science
University of Texas at Austin
Austin, TX 78712
(512) 471–9796

7208 Lamplight Lane
Austin, TX 78731
(512) 346–3809

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

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, robotics, and e-commerce. Application domains have included robot soccer, autonomous bidding agents, intelligent traffic management, social agents, 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**, September 2015 – Present.
  Associate Chair of the Department of Computer Science.
- **The University of Texas at Austin**, September 2014 – Present.
  David Bruton, Jr. Centennial Professor in the Department of Computer Science.
- **Cogitai, Inc.**, September 2015 – Present.
  President, COO, and co-Founder.
- **The University of Texas at Austin**, September 2012 – Present.
  Professor in the Department of Computer Science and Center for Perceptual Systems.
- **The University of Texas at Austin**, September 2007 – August 2012.
  Associate Professor in the Department of Computer Science and Center for Perceptual Systems.
- **The Hebrew University of Jerusalem** and **Bar Ilan University**, September 2008 – June 2009.
  Visiting Professor in the Computer Science Departments.
- **The University of Texas at Austin**, June 2002 – August 2007.
  Assistant Professor in the Department of Computer Sciences and Center for Perceptual Systems.
- **RobotsLab**, September 2012 – December 2015.
  Consultant. Helping design educational robotics curriculum.
- **SAIC**, then **Leidos** April 2013 – May 2014.
  Consultant. Working 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

- 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.
- **IEEE Senior Member**, 2015.
- Elected to UT Austin **Academy of Distinguished Teachers**, 2014.
- **World Champion** team member in 13 **RoboCup** events: 3D simulator competition, July 2017, July 2016, July 2015, July 2014, June 2012, July 2011; standard platform competition, 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 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: *highest award in Artificial Intelligence*.

- **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 343H *Artificial Intelligence: Honors.* Autumn 2017.
  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
  CS 378 *Autonomous Multiagent Systems.* Spring 2005. Instructor rating: 4.9/5.0
  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
- **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)
  - Patrick MacAlpine, defended July 2017.
    Multilayered Skill Learning and Movement Coordination for Autonomous Robotic Agents.
  - 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.
  - Jacob Menashe, current (proposal Autumn 2016).
    Intrinsically-motivated Hierarchical Reinforcement Learning.
  - Elad Liebman, current (proposal Summery 2017).
    Sequential Decision Making in Musical Intelligence.
  - Sanmit Narvekar, current (proposal Summer 2017).
    Curriculum Learning in Reinforcement Learning.

- **Doctoral Committee Member:** (The University of Texas at Austin)
  - Kwan Suk Kim, Mechanical Engineering. Supervisor: Luis Setnis.
    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, 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 Application 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.
- Jesse Thomason, Computer Science, current. Supervisor: Raymond Mooney.
- Jason Liang, Computer Science, current. Supervisor: Risto Miikkulainen.
- Igor Karpov, Computer Science, current. Supervisor: Risto Miikkulainen.
- Julian Bishop, Computer Science, current. Supervisor: Risto Miikkulainen.

- **Doctoral Committee Member:** (External)
  - 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)
  - 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)
  - 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.
    Supervisor: Jeff Rosenschein.
    Collaboration in Ad Hoc Settings: Novel Approaches to Implementation.
    December 2014.
- **Undergraduate Honors Thesis Supervisor** (The University of Texas at Austin)
  - 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.
  - 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)
  - 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:** Patrick MacAlpine (2017–), Justin Hart (2016–), Guni Sharon (2015–), 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:** Josiah Hanna, Shih-Yun Lo, Foraz Torabi, Ishan Durugkar
- **Other UT Austin undergraduate research:** Rolando Fernandez (2015–17), Maxwell Svetlik (2015–17), Shun Zhang (2012–14), Andrew Sharp (2012–13), Art Richards (2011), Nicu Stiurca (2011), Francisco Barrera (2011), Bartley Gillan (2007), Mickey Ristroph (2007), Srinivas Ashok (2007), David Li (2007), David Reaves (2007), Thomas Nelson (2006–07), Augustine Mathew (2006–07), Ben Bradley (2004), Aashish Parekh (2004), Prashanth Govindarajan (2003), Bharat Kejriwal (2003), Justin Lallinger (2003), Ali Niaz (2003).
- **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**:
  - **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 Artificialy 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*, 2011, 2010.
  - **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, 2005, 2004.
  - **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.
  - *ACM Transactions on Internet Technology* (TOIT), 2003–2005.
  - *International Journal of Image and Graphics* (IJIG), 2002–2006.
- **Assistant editor**:
  - *ACM SIGecom Exchanges*, 2004–2005.
- **Editorial board**:
  - Springer Verlag's *Encyclopedia of Machine Learning*, 2005–2010.
  - *Machine Learning Journal* (MLJ), 2003–present.
  - *Journal of Artificial Intelligence Research* (JAIR), 2002–2005.
- **Organizing committee member**:
  - 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.
- **Vice President**, RoboCup Federation, 2013–present.
- **Trustee**, RoboCup Federation, 2003–present.
- **Executive committee member**, RoboCup Federation, 1999–present.
- **Secretary**, RoboCup, US, 2016–present.
- **Chair**, First Study Panel of the One Hundred Year Study on AI, 2015–2016.
- **Guest editor**:
  - 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.
- **Academic Committee**, University of Science and Technology China (USTC)-GOOCOO Robotics Research Center, 2015 – present.
- **Fellows Committee**, AAAI 2016–2018.
- **Feigenbaum Prize Committee**, AAAI 2014–2017.
- **Conference Committee**, AAAI, 2014–2016.
- **Advisory Committee**, AAAI, 2015.
- **Awards Committee**, IJCAI, 2010–2015.
- **Advisory Committee**, IJCAI, 2009.
- **Senior Steering Committee**, AAAI workshop on *Transfer Learning for Complex Tasks*, 2008.
- **Steering committee**:
  - LPNMR workshop on *Knowledge Representation and Planning in Robotics and Autonomous Systems* (KRPRAS), 2017.
  - AAAI workshop on *Multiagent Interaction without Prior Coordination*, 2016, 2015, 2014.
  - Adaptive and Learning Agents Workshop (ALA), 2008–present.
  - Pacific Rim Trading Agent Competition, 2007.
  - IPTO Cognitive Systems Conference, 2005–06.
- **Councilor**, Association for the Advancement of Artificial Intelligence (AAAI), 2005–2008.

- **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.
- **Co-editor**, IEEE Intelligent Systems special issue on "Agents and Markets," 2003.
- **Industrial Advisory Board**:
  - Esquared, 2015–present.
  - OneSpot, 2013–present.
- **Advisory board**, Springer-Verlag book on *Balancing Reactivity and Social Deliberation in MAS*, 2000–2001.
- **Advisor**, National Academy of Engineering DARPA Prize Authority Workshop, 2000.
- **Book reviewer**:
  - Elsevier, 2016, 2015.
  - Synthesis Lecture Series, 2011.
  - Cambridge University Press, 2010.
  - John Wiley & Sons, 2007, 2006.
  - 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), 2016, 2014, 2013, 2006, 2005, 2002.
  - *Autonomous Agents and Multi-Agent Systems* (JAAMAS), 2002 – 2007, 2000.
  - *Autonomous Robots*, 1999.
  - *Communications of the ACM* (CACM), 2010, 2009.
  - *Computational Intelligence*, 2003.
  - *Data Mining and Knowledge Discovery* (DMKD) 2007.
  - *Decision Support Systems* (DSS), 2007, 2006, 2004, 2003.
  - *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 Transactions on Intelligent Transportation Systems*, 2010.
  - *IEEE Transactions on Knowledge and Data Engineering* (IEEE TKDE), 2002, 1999.
  - *IEEE Transactions on Robotics* (IEEE TRO), 2004 – 2007.
  - *IEEE Transactions on Robotics and Automation* (IEEE TRA), 2002, 2001.
  - *International Journal of Robotics Research* (IJRR), 2013, 2012, 2011.
  - *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), 2013, 2012, 2000 – 2005.
  - *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), 2011, 2009, 2005, 2003.
  - *Knowledge and Information Systems* (KAIS), 2002, 2000.
  - *Knowledge Engineering Review*, 2003.
  - *Machine Learning Journal* (MLJ), 2015, 2005–2013, 2003.
  - *Neural Networks* (NN), 2008, 2007.
  - *Robotics and Autonomous Systems* (RAS), 2007, 2003.
  - *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.
- **Senior Area chair**:
  - Neural Information Processing Systems (NIPS), 2017, 2003, 2002.
- **Area chair**:
  - International Joint Conference on Artificial Intelligence (IJCAI), 2018, 2016, 2013, 2009.
  - AAAI Conference on Artificial Intelligence, 2018.
  - International Conference on Machine Learning (ICML), 2016, 2015, 2012, 2003.
  - European Conference on Machine Learning (ECML), 2005.
- **Senior program committee member**:
  - Autonomous Agents and Multiagent Systems (AAMAS), 2016, 2007, 2004, 2003.
  - International Joint Conference on Artificial Intelligence (IJCAI), 2015, 2007.
  - International Conference on Machine Learning (ICML), 2006.
  - AAAI Conference on Artificial Intelligence, 2004, 2002.
- **Conference program committee member**:
  - Multi-disciplinary Conference on Reinforcement Learning and Decision Making (RLDM), 2017, 2015.
  - Conference on Artificial Intelligence (AAAI), 2012 (Computational Sustainability Track), 2010, 2007 (Integrated Intelligence Track), 2000.
  - International Conference on Machine Learning (ICML), 2010, 2008, 2000.
  - Autonomous Agents and Multiagent Systems (AAMAS), 2008.
  - International Conf. on Automated Planning and Scheduling. (ICAPS), 2007, 2003.
  - Robotics: Science and Systems (RSS), 2006.
  - International Joint Conference on Artificial Intelligence (IJCAI), 2005, 2003, 2001.
  - International Conference on Autonomic Computing (ICAC), 2005, 2004.
  - ACM Conference on Electionic Commerce (EC), 2005.
  - Neural Information Processing Systems (NIPS), 2003, 2002.
  - European Conference on Machine Learning (ECML), 2001–2003.
  - Autonomous Intelligent Networks and Systems Conference (AINS), 2003.
  - Distributed Autonomous Robotic Systems (DARS), 2002, 2000.
  - Intelligent Autonomous Systems (IAS), 2002.
  - International Conference on Artificial Intelligence (IC-AI), 2001.
  - Autonomous Agents (AA), 2001, 2000.
  - International Conference on Multi-Agent Systems (ICMAS), 2000.
  - International Conference on Enterprise Information Systems (ICEIS), 2000.
- **Conference paper reviewer**
  - 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), 2015, 2013.
  - 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), 2008–2010, 2005.
  - IEEE Conference on Intelligent Transportation Systems (ITSC), 2009.
- **Workshop/symposium program committee member**:
  - AAAI Fall Symposium on *Natural Communication for Human-Robot Collaboration*, 2017.
  - RoboCup Symposium, 2016, 2015, 2013, 2010, 2001–2005.
  - RSS workshop on *Combining AI Reasoning and Cognitive Science with Robotics*, 2015.
  - IJCAI *AI video competition*, 2015, 2009.
  - AAMAS workshop on *Multiagent Sequential Decision Making Under Uncertainty* (MSDM), 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 Communciation* (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.
- **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:**
  - 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, Robotics Industrial Affiliates Program, 2017–current.
  - Chair, Provost's Future of Computing Task Force, 2017.
  - Associate Chair of Computer Science Department, 2015–current.
  - Chair, Graduate Portfolio Program in Robotics, 2015–current.
  - Faculty Awards and Honors Committee, 2016–17(chair).
  - GDC Advisory Board, 2015–16.
  - Search Committee for Neuroscience Department Chair, 2015–16.
  - Faculty Recruiting committee, 2015–16, 2014–15, 2012–13.
  - Colloquia committee, 2013–14.
  - College of Natural Sciences Strategic Planning Task Force, 2013.
  - Doctoral Admissions committee, 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.
  - Faculty Recruiting committee, 2006–07.
  - 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.
  - Faculty evaluation committee, 2003–04.
  - Departmental Best Dissertation committee, 2003.

INVITED DISTINGUISHED LECTURES

- "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 Univerity College of Computing & Informatics 6th Annual **Jay Modi Memorial Lecture**, Philadelphia, PA. March 2014.
- "PRISM – Practical RL: Representation, Interactions, Synthesis, and Mortality"
  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 (BISFAI)
  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

- "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 talk at Deutsche Bank Internet Conference.
  Palo Alto, CA. August 2017.
- "Reasoning about Hypothetical Agent Behaviours and their Parameters"
  ONR Science of Autonomy Program Review.
  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 talk at Synopsis SNUG Conference.
  Santa Clara, CA. March 2017.
- "Artificial Intelligence and Life in 2030"
  Keynote talk 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 talk 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"
  IJCAI 2016 Industry Day Keynote talk (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 talk 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 talk 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."
  IJCAI 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.

2. 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. 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.

9. Corey White, Elad Liebman, and **Peter Stone**. Decision mechanisms underlying mood-congruent emotional classification. *Cognition and Emotion*, pages 1–10, 2017.

10. 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*, 2017.

11. Stefano Albrecht, Somchaya Liemhetcharat, and **Peter Stone**. Special issue on multiagent interaction without prior coordination: Guest editorial. *Autonomous Agents and Multi-Agent Systems*, 31:4, pages 765–66, July 2017.

12. Katie Genter, Tim Laue, and **Peter Stone**. Three years of the robocup standard platform league drop-in 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.

13. Matteo Leonetti, Luca Iocchi, and **Peter Stone**. A synthesis of automated planning and reinforcement learning for efficient, robust decision-making. *Artificial Intelligence*, 241:103 – 130, September 2016.

14. Samuel Barrett, Avi Rosenfeld, Sarit Kraus, and **Peter Stone**. Making friends on the fly: Cooperating with new teammates. *Artificial Intelligence*, October 2016.

15. 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.

16. Katie Genter and **Peter Stone**. Ad hoc teamwork behaviors for influencing a flock. *Acta Polytechnica*, 56(1), 2016.

17. Elad Liebman, Benny Chor, and **Peter Stone**. Representative selection in nonmetric datasets. *Applied Artificial Intelligence*, 29:807–838, 2015.

18. Todd Hester and **Peter Stone**. Intrinsically motivated model learning for developing curious robots. *Artificial Intelligence*, May 2015.

19. W. Bradley Knox and **Peter Stone**. Framing reinforcement learning from human reward: Reward positivity, temporal discounting, episodicity, and performance. *Artificial Intelligence*, 225, August 2015.

20. 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.

21. **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*, 203:35–65, October 2013.

22. Doran Chakraborty and **Peter Stone**. Multiagent learning in the presence of memory-bounded agents. *Autonomous Agents and Multiagent Systems (JAAMAS)*, 2013.

23. Todd Hester and **Peter Stone**. TEXPLORE: Real-time sample-efficient reinforcement learning for robots. *Machine Learning*, 90(3), 2013.

24. 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.

25. 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*, 4:409–421, October 2012.

26. 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.

27. Shivaram Kalyanakrishnan and **Peter Stone**. Characterizing reinforcement learning methods through parameterized learning problems. *Machine Learning* (MLJ), 84(1):205–247, June 2011.

28. 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.

29. Tobias Jung, Daniel Polani, and **Peter Stone**. Empowerment for continuous agent-environment systems. *Adaptive Behavior*, 19(1):16–39, 2011.

30. 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.

31. Jonathan Wildstrom, **Peter Stone**, and Emmett Witchel. Autonomous return on investment analysis of additional processing resources. *International Journal on Autonomic Computing*, 1(3), 2010.

32. 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.

33. 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.

34. 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.

35. 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.

36. 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.

37. Kurt Dresner and **Peter Stone**. A multiagent approach to autonomous intersection management. *Journal of Artificial Intelligence Research* (JAIR), 31:591–656, March 2008.

38. 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.**

39. 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.

40. 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.

41. Mohan Sridharan and **Peter Stone**. Structure-Based Color Learning on a Mobile Robot under Changing Illumination. *Autonomous Robots*, 23(3):161–182, 2007.

42. **Peter Stone**. Multiagent learning is not the answer. it is the question. *Artificial Intelligence* (AIJ), 171:402–405, 2007.

43. Shimon Whiteson, Matthew E. Taylor, and **Peter Stone**. Empirical studies in action selection for reinforcement learning. *Adaptive Behavior*, 15(1):33–50, 2007.

44. **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.

45. 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.

46. Shimon Whiteson and **Peter Stone**. Evolutionary Function Approximation for Reinforcement Learning. *Journal of Machine Learning Research* (JMLR), 7:877–917, May 2006.

47. 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.

48. **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.

49. 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.

50. 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.

51. **Peter Stone** and Amy Greenwald. The first international trading agent competition: Autonomous bidding agents. *Electronic Commerce Research* (EC), 5(2):229–65, April 2005.

52. 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.

53. 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.

54. **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.

55. 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.

56. 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.

57. **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.

58. **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.

59. Michael Bowling, Manuela Veloso, and **Peter Stone**. The CMUnited-98 champion small-robot team. *Advanced Robotics*, 13(8):753–766, 2000.

60. 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.

61. **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.

62. **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.

63. 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.

64. **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.

65. 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

66. 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*, Lecture Notes in Artificial Intelligence. Springer Verlag, Berlin, 2017.
Nominee for Best Paper Award.

67. 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*, Lecture Notes in Artificial Intelligence. Springer Verlag, Berlin, 2016.

68. 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*, Lecture Notes in Artificial Intelligence. Springer Verlag, Berlin, 2016.

69. 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.

70. Tsz-Chiu Au, Shun Zhang, and **Peter Stone**. Autonomous intersection management for semi-autonomous vehicles. In *Handbook of Transportation*. Routledge, 2016.

71. 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.

72. 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.

73. 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.

74. 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.

75. 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.

76. 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.

77. 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.

78. 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.

79. 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.

80. 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.

81. 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.

82. 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.

83. 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.

84. 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.

85. 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*, Lecture Notes in Artificial Intelligence. Springer Verlag, pages 254–65, 2011.

86. **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.

87. 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.

88. **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.

89. 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.

90. Shivaram Kalyanakrishnan and **Peter Stone**. Learning complementary multiagent behaviors: A case study. In *Proceedings of the RoboCup International Symposium 2009*. Springer Verlag, 2009. **Winner of Best Student Paper Award.**

91. 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.

92. 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.

93. 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.

94. 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.

95. 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.

96. 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. **Winner of Best Paper Award.**

97. 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.

98. 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.

99. 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*. volume 4434, pages 158–170. Springer Verlag, Berlin, 2007.

100. 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*. volume 4434, pages 72–85. Springer Verlag, Berlin, 2007. **Winner of Best Student Paper Award.**

101. 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*. volume 4434, pages 270–278. Springer Verlag, Berlin, 2007.

102. 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*. volume 4434, pages 59–71. Springer Verlag, Berlin, 2007.

103. 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*. volume 4434, pages 98–109. Springer Verlag, Berlin, 2007. **Nominee for Best Paper Award.**

104. 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.

105. 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*, volume 4020, pages 673–81. Springer Verlag, Berlin, 2006.

106. **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*, volume 4020, pages 93–105. Springer Verlag, Berlin, 2006.

107. 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.

108. 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.

109. 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.

110. 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.

111. 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.

112. 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.

113. 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.

114. **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.

115. **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.

116. **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.

117. 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.

118. **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.

119. **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.

120. 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.

121. **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).

122. **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.

123. **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.

124. **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.

125. **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.

126. **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.

127. **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

128. 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.

129. 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.

130. Piyush Khandelwal and **Peter Stone**. Multi-robot human guidance: Human experiments and multiple concurrent requests. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, May 2017.

131. 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.

132. 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.

133. 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.

134. 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 Sytems (AAMAS)*, May 2017.

135. 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.

136. 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.

137. 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.

138. 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.

139. 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.

140. 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.

141. 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.

142. 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.

143. 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.

144. 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.

145. 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.

146. 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.

147. 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.

148. 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.

149. 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.

150. Matthew Hausknecht and **Peter Stone**. Deep reinforcement learning in parameterized action space. In *Proceedings of the International Conference on Learning Representations (ICLR)*, May 2016.

151. 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.

152. **Peter Stone**. What's hot at RoboCup (extended abstract). In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI)*, February 2016.

153. 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.

154. 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.

155. 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 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, September 2015.
Finalist for RoboCup Best Paper Award.

156. 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.

157. 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.

158. 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.
Winner of Computational Sustainability Outstanding Paper Award.

159. 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.

160. 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.

161. 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.

162. 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.

163. 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.

164. 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.

165. 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.

166. 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.

167. 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.

168. 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.

169. 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.

170. 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.

171. 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.

172. 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.

173. 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.

174. 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. **Winner of Best Paper Award.**

175. 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.

176. 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.

177. 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.

178. 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.

179. 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.

180. 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.

181. 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.

182. 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.

183. W. Bradley Knox and **Peter Stone**. Learning non-myopically from human-generated reward. In *In Proceedings of the International Conference on Intelligent User Interfaces* (IUI), March 2013.

184. 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.

185. 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.**

186. 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. **Winner of Paper of Excellence Award.**

187. 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.

188. 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.

189. 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.

190. 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.

191. 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.

192. 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.

193. 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.

194. 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.

195. 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.

196. 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.

197. 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.

198. 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.

199. 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.

200. 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.

201. 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.

202. 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.

203. 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.

204. 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.

205. 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.

206. 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.

207. 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.

208. 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.

209. Michael Quinlan, Tsz-Chiu Au, Jesse Zhu, Nicolae Stiurca, and **Peter Stone**. Bringing simulation to life: A mixed reality autonomous intersection. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems* (IROS), October 2010.

210. 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.

211. Todd Hester and **Peter Stone**. Real time targeted exploration in large domains. In *The Ninth International Conference on Development and Learning* (ICDL), August 2010.

212. 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.

213. 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.

214. David Pardoe and **Peter Stone**. Boosting for regression transfer. In *Proceedings of the 27th International Conference on Machine Learning* (ICML), June 2010.

215. **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.

216. 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. **Winner of Best Student Paper Award.**

217. 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.

218. **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.

219. 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.

220. Peter Djeu, Michael Quinlan, and **Peter Stone**. Improving particle filter performance using sse instructions. In *Proceedings of IEEE/RSJ International Conference on Intelligent RObots and Systems* (IROS), October 2009.

221. 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.

222. 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.

223. 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.

224. 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.

225. 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* (AAMAS), May 2009.

226. 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.

227. 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.

228. 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.

229. 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.

230. 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.

231. Jonathan Wildstrom, **Peter Stone**, and Emmett Witchel. CARVE: A Cognitive Agent for Resource Value Estimation. In *International Conference on Autonomic Computing* (ICAC), June 2008.

232. 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.

233. 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.

234. 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.

235. 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.

236. 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.

237. 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.

238. 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.

239. M. Sridharan and **Peter Stone**. Color learning on a mobile robot: Towards full autonomy under changing illumination. In *The IEEE International Conference on Intelligent Robots and Systems* (IROS), 2007.

240. 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.

241. 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.

242. 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.

243. 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.

244. 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.

245. 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.**

246. 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.**

247. 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.

248. **Peter Stone**. Learning and multiagent reasoning for autonomous agents. In *The 20th International Joint Conference on Artificial Intelligence* (IJCAI), pages 13–30, January 2007.

249. 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.

250. 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.

251. 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.

252. 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.

253. 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.

254. 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.

255. 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.

256. 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.

257. 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.

258. 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.

259. 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.

260. 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.

261. 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.

262. 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.

263. 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.
Winner of Best Paper Award, GA Track.

264. David Pardoe and **Peter Stone**. Predictive planning for supply chain management. In *International Conference on Automated Planning and Scheduling* (ICAPS), June 2006.

265. 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.

266. Mohan Sridharan and **Peter Stone**. Real-time vision on a mobile robot platform. In *IEEE/RSJ International Conference on Intelligent Robots and Systems* (IROS), August 2005.

267. 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.

268. 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.

269. 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.

270. 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.

271. 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.

272. 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.

273. 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.

274. 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.

275. 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.

276. 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.

277. 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.

278. 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.

279. 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.

280. 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.

281. 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.

282. 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.

283. 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.

284. 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.

285. 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.

286. **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.

287. 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. **Winner of Best Paper Award.**

288. **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.

289. **Peter Stone**. TPOT-RL applied to network routing. In *Proceedings of the Seventeenth International Conference on Machine Learning* (ICML), pages 935–942, 2000.

290. 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.

291. 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.

292. 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.

293. 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.

294. 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.

295. 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.

296. **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.

297. **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

298. 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.

299. 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.

300. Shiqi Zhang and **Peter Stone**. Integrated Commonsense Reasoning and Probabilistic Planning. In *ICAPS workshop on Planning for Robotics*, June 2017.

301. Patrick MacAlpine and **Peter Stone**. Evaluating ad hoc teamwork performance in drop-in player challenges. In *AAMAS Workshop on Multiagent Interaction without Prior Coordination (MIPC)*, May 2017.

302. Guni Sharon and **Peter Stone**. A protocol for mixed autonomous and human-operated vehicles at intersections. In *AAMAS Workshop on Agent-based modelling of urban systems (ABMUS)*, May 2017.

303. Shiqi Zhang, Jivko Sinapov, Suhua Wei, and **Peter Stone**. Robot behavioral exploration and multimodal perception using POMDPs. In *Proceedings of 2017 AAAI Spring Symposium on Interactive Multi-Sensory Perception for Embodied Agents*, March 2017.

304. Matthew Hausknecht and **Peter Stone**. Grounded semantic networks for learning shared communication protocols. In *NIPS workshop on Deep Reinforcement Learning*, December 2016.

305. 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.

306. 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.

307. 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.

308. Shiqi Zhang, Piyush Khandelwal, and **Peter Stone**. Dynamically constructed (PO)MDPs for adaptive robot planning. In *Proceedings of the IJCAI Workshop on Autonomous Mobile Service Robots (WSR)*, July 2016.

309. 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.

310. Shih-Yun Lo, Shiqi Zhang, and **Peter Stone**. Integrated Task and Motion Planning for Mobile Service Robots. In *RSS Workshop on Integrating Task and Motion Planning*, June 2016.

311. 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 Learning Agents (ALA)*, May 2016.

312. Matthew Hausknecht, Yilun Chen, and **Peter Stone**. Deep imitation learning for parameterized action spaces. In *AAMAS Workshop on Adaptive Learning Agents (ALA)*, May 2016.

313. Jonathan Grizou, Samuel Barrett, Manuel Lopes, and **Peter Stone**. Collaboration in ad hhoc teamwork: ambiguous tasks, roles, and communication. In *AAMAS Workshop on Adaptive Learning Agents (ALA)*, May 2016.

314. 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.

315. 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.

316. Shiqi Zhang, Fangkai Yang, Piyush Khandelwal, and **Peter Stone**. Mobile robot planning using action language *BC* with hierarchical domain abstractions. In *The 7th Workshop on Answer Set Programming and Other Computing Paradigms (ASPOCP)* , July 2014.

317. Piyush Khandelwal and **Peter Stone**. Multi-robot human guidance using topological graphs. In *AAAI Spring 2014 Symposium on Qualitative Representations for Robots*, March 2014.

318. 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.

319. 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.

320. 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.

321. Patrick MacAlpine, Elad Liebman, and **Peter Stone**. Simultaneous learning and reshaping of an approximated optimization task. In *AAMAS Workshop on Adaptive Learning Agents* (ALA), May 2013.

322. Doran Chakraborty, Noa Agmon, and **Peter Stone**. Targeted opponent modeling of memory-bounded agents. In *AAMAS Workshop on Adaptive Learning Agents* (ALA), May 2013.

323. Patrick MacAlpine and **Peter Stone**. Using dynamic rewards to learn a fully holonomic bipedal walk. In *AAMAS Workshop on Adaptive Learning Agents* (ALA), June 2012.

324. Samuel Barrett, **Peter Stone**, Sarit Kraus, and Avi Rosenfeld. Learning teammate models for ad hoc teamwork. In *AAMAS Workshop on Adaptive Learning Agents* (ALA), June 2012.

325. 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.

326. W. Bradley Knox and **Peter Stone**. Augmenting Reinforcement Learning with Human Feedback. In *ICML Workshop on Imitation Learning*, July 2011.

327. 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.

328. 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.

329. 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.

330. 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.

331. 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.

332. 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.

333. 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.

334. 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.

335. Samuel Barret, Matthew E. Taylor, and **Peter Stone**. Transfer learning for reinforcement learning on a physical robot In *AAMAS 2010 Workshop on Adaptive Learning Agents (ALA)*, May 2010.

336. Tobias Jung, Mazda Ahmadi, and **Peter Stone**. Connectivity-based localization in robot networks. In *International Workshop on Robotic Wireless Sensor Networks (WSN)*, June 2009.

337. 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.

338. 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.

339. 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.

340. 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.

341. Kurt Dresner and **Peter Stone**. Learning policy selection for autonomous intersection management. In *AAMAS workshop on adaptive and learning agents*, May 2007.

342. 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.

343. Matthew E. Taylor, Gregory Kuhlmann, and **Peter Stone**. Accelerating search with transferred heuristics. In *ICAPS workshop on AI Planning and Learning*, September 2007.

344. Nicholas K. Jong and **Peter Stone**. Model-based exploration in continuous state spaces. In *The Seventh Symposium on Abstraction, Reformulation, and Approximation*, July 2007.

345. 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.

346. 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.

347. 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.

348. Daniel Stronger and **Peter Stone**. Expectation-based vision for precise self-localization on a mobile robot. In *AAAI workshop on Cognitive Robotics*, July 2006.

349. Mazda Ahmadi and **Peter Stone**. Instance-based action models for fast action planning. In *AAAI workshop on Cognitive Robotics*, July 2006.

350. **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.

351. 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.

352. 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.

353. 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.

354. 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.

355. 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.

356. 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.

357. 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.

358. 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.

359. 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.

360. 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.

361. **Peter Stone**. Multiagent learning for autonomous spacecraft constellations. In *Proceedings of the NASA Workshop on Planning and Scheduling for Space*, 1997.

362. 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.

363. 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.

364. **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.

365. **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.

366. **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.

367. **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

368. 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.

369. Daniele Nardi, Itsuki Noda, Fernando Ribeiro, **Peter Stone**, Oskar von Stryk, and Manuela Veloso. RoboCup soccer leagues. *AI Magazine*, 35(3):77–85, 2014.

370. David Pardoe and **Peter Stone**. Designing Adaptive Trading Agents. *ACM SIGecom Exchanges*, 10(2), June 2011.

371. Matthew E. Taylor and **Peter Stone**. An Introduction to Inter-task Transfer for Reinforcement Learning. *AI Magazine*, 32(1):15–20, March 2011.

372. David Pardoe and **Peter Stone**. Developing Adaptive Auction Mechanisms. *ACM SIGecom Exchanges*, 5(3):1–10, April 2005.

373. David Pardoe and **Peter Stone**. TacTex-03: A Supply Chain Management Agent. *ACM SIGecom Exchanges*, 4(3):19–28, Winter 2004.

374. Amy Greenwald, Nicholas R. Jennings, and **Peter Stone**. Guest Editors' Introduction: Agents and Markets. *IEEE Intelligent Systems*, 18(6):12–14, November/December 2003.

375. 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.

376. **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.

377. Amy Greenwald and **Peter Stone**. Autonomous bidding agents in the trading agent competition. *IEEE Internet Computing*, 5(2), 2001.

378. Silvia Coradeschi, Lars Karlsson, **Peter Stone**, Tucker Balch, Gerhard Kraetzschmar, and Minoru Asada. Overview of RoboCup-99. *AI Magazine*, 21(3), 2000.

379. 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

380. Manuela Veloso and **Peter Stone**. Video: Robocup robot soccer history 1997 – 2011, October 2012. *Proceedings of the 2015 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.

381. 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

382. Decebal Constantin Mocanu, Elena Mocanu, **Peter Stone**, Phuong H. Nguyen, Madeleine Gibescu, and Antonio Liotta. Evolutionary training of sparse artificial neural networks: A network science perspective. *arXiv e-Prints 1707.04780*, August 2017.

383. Kara Kockelman, Stephen Boyles, **Peter Stone**, Dan Fagnant, Rahul Patel, Michael W. Levin, Guni Sharon, Michele 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.

384. 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.

385. 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.

386. **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.

387. 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.

388. 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.

389. 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.

390. 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.

391. 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.

392. Shivaram Kalyanakrishnan and **Peter Stone**. UT Austin Villa 3D simulation soccer team 2008. In *RoboCup Symposium CD Proceedings*, 2008.

393. Peter Stone Todd Hester, Michael Quinlan. UT Austin Villa 2008 team description paper for the standard platform league. In *RoboCup Symposium CD Proceedings*, 2008.

394. **Peter Stone** and Neda Shahidi. The UT Austin Villa 2008 mixed reality team. In *RoboCup Symposium CD Proceedings*, 2008.

395. **Peter Stone**, Patrick Beeson, Tekin Mericli, and Ryan Madigan. DARPA urban challenge technical report: Austin Robot Technology, June 2007.

396. 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.

397. 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.

398. **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.

399. **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.

400. **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.

401. 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.

402. **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.

403. 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.

404. **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.

405. 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.

406. **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.

407. **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.

408. 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.

409. **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.

410. 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.

411. **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.

412. **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.

413. **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.

414. 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.

415. 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.

416. 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, Senior Member (IEEE)
- American Association for the Advancement of Science (AAAS)
- Association for Computing Machinery (ACM)

PRESS

Interviewed and quoted regarding research several times on television, on radio, and in magazines and newspapers including CNN, NPR, 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.

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).