

| | |
|--|--|
| Peter Herald Stone 7208 Lamplight Lane Austin, TX 78731 (512) 346-3809 | Associate Professor Dept. of Computer Science University of Texas at Austin Austin, TX 78712 (512) 471-9796 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 2007 – present.
Associate Professor in the Department of Computer Sciences 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.
- **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.
- **Jet Propulsion Laboratory**, May – August 1995.
Summer intern. Worked on automatic planning and scheduling for the New Millennium Project.

AWARDS AND RESEARCH DISTINCTIONS

- Elected **Board Member**, International Machine Learning Society (IMLS), March 2011.
- Leader of **1st-place** teams in the **Trading Agent Competition** (TAC), June 2010; July 2009; July 2008; May 2006; August 2005; August 2003; October 2001; July 2000.
- Adviser 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.
- RoboCup **US Open Champion** team leader, standard platform league, May 2010, May 2009.
- **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), March 2008.
- 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.
- **World Champion** team member in 6 **RoboCup** events: simulator coach competition, July 2005, July 2003; simulator competition, August 1999; simulator and small-size robot competitions, July 1998; small-size robot competition, August 1997.
- 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

- **Associate Professor** at **The University of Texas at Austin**: September 2007 – present.
 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 AAI-07, May – July 2007.
- **Tutorials** on *robot soccer* at AAI-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)
 - 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.
 - 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.
 - Juhyun Lee, current (proposal Autumn 2009).
Color Constancy for Mobile Robots using Computer Graphics.
 - Shivaram Kalyanakrishnan, current (proposal Autumn 2009).
Learning Methods for Sequential Decision Making in Practice.
 - W. Bradley Knox, current (proposal Summer 2010).
Interactively Shaping Agents via Human Reinforcement.
 - Todd Hester, current (proposal Autumn 2010).
Temporal Difference Reinforcement Learning in Time-Constrained Domains.
 - Doran Chakraborty, current (proposal Autumn 2010).
Learning in the Presence of Other Agents.
- **Doctoral Committee Member:** (The University of Texas at Austin)
 - Vinod Valsalam, Computer Science, Supervisor: Risto Miikkulainen.
Utilizing Symmetry In Evolutionary Design, defended August 2010.
 - Jonathan Mugan, Computer Science, Supervisor: Ben 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, current. 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, current. Supervisor: Kathryn McKinley.
Diagnosing And Tolerating Bugs In Deployed Systems, defended September 2008.
- Patrick Beeson, Computer Sciences. Supervisor: Ben 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, current. Supervisor: Don Fussell.
Parametric Kernels for Structured Data Analysis, defended December 2007.
- Karen Fullam, Electrical and Computer Engineering (ECE). Supervisor: K. Suzanne Barber.
Adaptive Trust Modeling in Multi-Agent Systems: Utilizing Experience and Reputation, defended November 2007.
- Rohit Kate, Computer Sciences. Supervisor: Ray Mooney.
Learning For Semantic Parsing With Kernels Under Various Forms Of Supervision, defended August 2007.
- Jefferson Provost, Computer Sciences. Supervisors: Ben 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 (ECE), Supervisor: K. Suzanne Barber.
Trusting Information and Sources in Open Multi-Agent Systems, defended November 2003.
- David Han, ECE, current. Supervisor: K. Suzanne Barber.
- Changhai Xu, Computer Science, current. Supervisor: Kristen Grauman.
- Jacob Schrum, Computer Science, current. Supervisor: Risto Miikkulainen.
- David Chen, Computer Science, current. Supervisor: Ray Mooney.

- Dmitry Kit, Computer Science, current. Supervisor: Dana Ballard.
- Bryan Silverthorn, Computer Science, current. Supervisor: Risto Miikkulainen.
- Andrew Hunter, ECE, current. Supervisor: Jeff Andrews.
- Chinmayi Krishnappa, Computer Science, current. Supervisor: Greg Plaxton.
- Jeremy Stober, Computer Science, current. Supervisors: Ben Kuipers, Risto Miikkulainen.
- Wenke Li, Neuroscience, current. Supervisor: Mike Mauk.
- **Doctoral Committee Member:** (External)
 - 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, Department of 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, current. 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, University of Newcastle, Australia. Supervisor: Stephan Chalup. *Machine Learning on AIBO Robots*. June 2006.
 - Jeff Riley, RMIT, 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)
 - 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)
 - Aravind Gowrisankar, Computer Sciences, Autumn 2008. Supervisor: Risto Miikkulainen. *Evolving 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*.
- **Undergraduate Honors Thesis Supervisor** (The University of Texas at Austin)
 - 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)
 - 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: Ben Kuipers. *Rapid, High Precision Control in Tightly Constrained Environments.*

OTHER ADVISING

- **Postdoctoral Fellows:** Noa Agmon (2010–present), Tsz-Chiu Au (2008–present), Michael Quinlan (2007–2011), Tobias Jung (2008–2010), Ian Fasel (2007–2008), Yaxin Liu (2004–2007), Bikramjit Banerjee (2006).
- **Other Current UT Austin Ph.D. students:** Samuel Barrett, Katie Genter, Matthew Hausknecht, Piyush Khandelwal, Patrick MacAlpine, Daniel Urieli.
- **Other UT Austin undergraduate research:** Nick Collins (2011), Art Richards (2011), Nicu Sturca (2011), Adrian Lopez-Mobilia (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:**
 - **Video Track Chair**, IJCAI, July 2011.
 - **General co-chair**, Autonomous Agents and Multi-Agent Systems (AAMAS), May 2011.
 - **Tutorial co-chair**, AAI, July 2008.
 - **Program co-chair**, Autonomous Agents and Multi-Agent Systems (AAMAS), May 2006.
 - **Workshop co-chair**, AAI, July 2005.
 - **Chair**, RoboCup US Open simulation league committee, May 2005, April 2004.
 - **Tutorial chair**, IJCAI, August 2003.
 - **Entry coordinator**, Trading Agent Competition, October 2001.
 - **Associate chair** in charge of simulation events for RoboCup-2001, August 2001.
 - **Co-chair**, RoboCup simulator competition organizing committee, August 1997 – August 1999.
- **Workshop/Symposium coordination:**
 - **Program Co-chair**, Humanoids Workshop on *Humanoid Soccer Robots*, December 2010.
 - **Co-chair**, AAI workshop on *Multiagent Learning*, July 2005.
 - **Co-chair**, ICML Workshop on *Physiological Data Mining Contest*, July 2004.
 - **Chair**, Information Science and Technology (ISAT) study on *Distributed Cognitive Systems Focused on Team/Multiagent Learning*, May 2005, June 2004.
 - **Co-chair**, IJCAI workshop on *Trading Agent Design and Analysis*, August 2003.
 - **Co-chair**, AAI Fall Symposium on *Personalized Agents*, November 2002.
 - **Co-chair**, AAI Spring Symposium on *Collaborative Learning Agents*, March 2002.
 - **Co-chair**, Agents Workshop on *Learning Agents*, May 2001.
 - **Co-chair**, RoboCup Workshop, August 2000.
 - **Co-chair**, Agents Workshop on *Learning Agents*, May 2000.
- **Editor-in-chief**, *J. of Autonomous Agents and Multi-Agent Systems* (JAAMAS), 2010–present.
- **Associate editor:**
 - *International Conference on Robotics and Automation* (ICRA), May 2011.
 - *Artificial Intelligence Journal* (AIJ), January 2007 – December 2014.
 - *J. of Autonomous Agents and Multi-Agent Systems* (JAAMAS), April 2003 – December 2009
 - *ACM Transactions on Internet Technology* (TOIT), January 2003 – December 2005.
 - *International Journal of Image and Graphics* (IJIG), November 2002 – March 2006.

- **Assistant editor:**
 - *ACM SIGecom Exchanges*, January 2004 – December 2005.
- **Editorial board:**
 - Springer Verlag's *Encyclopedia of Machine Learning*, October 2005 – 2010.
 - *Machine Learning Journal* (MLJ), May 2003 – December 2011.
 - *Journal of Artificial Intelligence Research* (JAIR), February 2002 – February 2005.
- **Organizing committee member:**
 - IJCAI Workshop on *General Game Playing*, July 2009.
 - 2nd *Reinforcement Learning Competition*, March 2008.
 - ICAPS Workshop on *AI Planning and Learning*, September 2007.
 - NIPS Workshop on *The Inaugural Reinforcement Learning Competition*, December 2006.
 - AAAI Fall Symposium on *Real Life Reinforcement Learning*, November 2004.
 - AAMAS Workshop on *Learning and Evolution in Agent Based Systems*, July 2004.
- **Trustee**, RoboCup Federation, July 2003 – present.
- **Executive committee member**, RoboCup Federation, August 1999 – present.
- **Panelist**, SxSW Interactive, Austin, TX on *AI 2010: Wall-e Or Rise Of The Machines?* March 2010.
- **Awards Committee**, IJCAI, August 2010–2015.
- **Advisory Committee**, IJCAI, July 2009.
- **Senior Steering Committee**, 2008 AAAI workshop on *Transfer Learning for Complex Tasks*.
- **Steering committee:**
 - Adaptive and Learning Agents Workshop (ALA), 2008–present.
 - 2007 Pacific Rim Trading Agent Competition.
 - IPTO Cognitive Systems Conference, 2005–2006.
- **Councilor**, Association for the Advancement of Artificial Intelligence (AAAI), July 2005–2008.
- **Board of directors**, Association for Trading Agent Research, August 2003–2009.
- **Guest editor**, ACM SIGecom Exchanges special issue on Trading Agent Design and Analysis, Winter 2004.
- **Consultant**, Information Science and Technology (ISAT) Summer Study on *Automated Intent Recognition on Distributed Organizations* (AIRDO), August 2003.
- **Co-editor**, IEEE Intelligent Systems special issue on “Agents and Markets,” November 2003.
- **Advisory board**, Springer-Verlag book on *Balancing Reactivity and Social Deliberation in MAS*, September 2000 – January 2001.
- **Advisor**, National Academy of Engineering DARPA Prize Authority Workshop, December 2000.
- **Book reviewer:**
 - 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), 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.
 - *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.
- *INFORMS Journal on Computing*, 2006.
- *Journal of Artificial Intelligence Research* (JAIR), 2000 – 2005.
- *Journal of Behavioral Robotics*, 2009.
- *Journal of Intelligent Traffic Systems* (JITS), 2010.
- *Journal of Machine Learning Research* (JMLR), 2009, 2005, 2003.
- *Knowledge and Information Systems* (KAIS), 2002, 2000.
- *Knowledge Engineering Review*, 2003.
- *Machine Learning Journal* (MLJ), 2005–2010, 2003.
- *Neural Networks* (NN), 2008, 2007.
- *Robotics and Autonomous Systems* (RAS), 2007, 2003.
- *Systems, Man and Cybernetics* (SMC), 2005.
- *Transportation Research Part C* (TRC), 2010.
- *Wiley Interdisciplinary Reviews: Cognitive Science*, 2009.
- **Area chair:**
 - International Joint Conference on Artificial Intelligence (IJCAI), July 2009.
 - European Conference on Machine Learning (ECML), October 2005.
 - International Conference on Machine Learning (ICML), August 2003.
- **Senior program committee member:**
 - Autonomous Agents and Multiagent Systems (AAMAS), May 2007, July 2004, July 2003.
 - International Joint Conference on Artificial Intelligence (IJCAI), January 2007.
 - International Conference on Machine Learning (ICML), June 2006.
 - National Conference on Artificial Intelligence (AAAI), July 2004, July 2002.
- **Conference program committee member:**
 - Nat'l Conf. on Artificial Intelligence (AAAI), July 2010, July 2007 (II Track), July 2000.
 - International Conference on Machine Learning (ICML), June 2010, June 2008, June 2000.
 - Autonomous Agents and Multiagent Systems (AAMAS), May 2008
 - Int'l Conf. on Automated Planning and Scheduling. (ICAPS), September 2007, June 2003.
 - Robotics: Science and Systems (RSS), August 2006.
 - International Joint Conference on Artificial Intelligence (IJCAI), August 2005, 2003, 2001.
 - International Conference on Autonomic Computing (ICAC), June 2005, May 2004.
 - ACM Conference on Electronic Commerce (EC), May 2005.
 - Neural Information Processing Systems (NIPS), December 2003, December 2002.
 - European Conference on Machine Learning (ECML), Sept. 2003, August 2002, Sept. 2001.
 - Autonomous Intelligent Networks and Systems Conference (AINS), June 2003.
 - Distributed Autonomous Robotic Systems (DARS), June 2002, October 2000.
 - Intelligent Autonomous Systems (IAS), March 2002.
 - International Conference on Artificial Intelligence (IC-AI), June 2001
 - Autonomous Agents (AA), May 2001, May 2000.
 - International Conference on Multi-Agent Systems (ICMAS), July 2000.
 - International Conference on Enterprise Information Systems (ICEIS), July 2000.
- **Conference paper reviewer**
 - Neural Information Processing Systems (NIPS), December 2011.
 - Autonomous Agents and Multiagent Systems (AAMAS), May 2010.
 - IEEE International Conference on Robotics and Automation (ICRA), May 2010, May 2009, May 2008, April 2005.
 - IEEE Conference on Intelligent Transportation Systems (ITSC), October 2009.
- **Workshop/symposium program committee member:**
 - AAAI *Doctoral Consortium*, August 2011.
 - IEEE International Symposium on Robot and Human Interaction Communication (Ro-Man) — associate editor, July 2011.
 - IJCAI workshop on *Trading Agent Design and Analysis* (TADA), July 2011.

- ECAI workshop on *Benchmarking Intelligent (Multi-)Robot Systems* (BIMRS), August 2010.
- AAAI *AI video competition*, July 2010.
- AAMAS 2010 workshop on *Agents in Real-Time and Dynamic Environments*, May 2010.
- IJCAI *AI video competition*, July 2009.
- AAAI workshop on *Trading Agent Design and Analysis* (TADA), July 2008.
- ECAI workshop on *Cognitive Robotics*, July 2008.
- AAMAS workshop on *Formal Models and Methods for Multi-Robot Systems*, May 2008.
- ICAPS workshop on *AI Planning and Learning* (AIPL), September 2007.
- ICAC workshop on *Adaptive Methods in Autonomic Computing* (AMACS), June 2007.
- Second workshop on *Tackling Computer Systems Problems with Machine Learning Techniques* (SysML), January 2007.
- ICML workshop on *Structural Knowledge Transfer for Machine Learning*, June 2006.
- 3rd International IEEE Latin American Robotic Symposium (LARS), October 2006.
- AAMAS workshop on *Agent Mediated Electronic Commerce* (AMEC/TADA), May 2006.
- IJCAI workshop on *Planning and Learning in A Priori Unknown or Dynamic Domains*
- IJCAI workshop on *Trading Agent Design and Analysis* (TADA), August 2005.
- IJCAI workshop on *Agents in Real-Time and Dynamic Environments*, August 2005.
- AAMAS workshop on *Learning and Adaptation in MAS* (LAMAS), July 2005.
- AAMAS workshop on *Teaching Multiagent Systems* (TeachMAS), July 2005.
- AAMAS workshop on *Trading Agent Design and Analysis* (TADA), July 2004.
- ICML workshop on *Predictive Representations of World State*, July 2004.
- RoboCup Symposium, June 2010, July 2005, July 2004, July 2003, June 2002, August 2001.
- IJCAI Workshop on *Agents in Dynamic Real-Time Environments*, August 2003.
- AAMAS Workshop on *Resource, Role, and Task Allocation in MAS*, July 2003.
- Pacific Rim MultiAgent Workshop (PRIMA), August 2002, July 2001, August 2000.
- AAMAS Workshop on *Coalitions and Team Formation*, July 2002.
- AAMAS Workshop on *MAS Problem Spaces and Their Implications to Achieving Globally Coherent Behavior Coalitions and Team Formation*, July 2002.
- *Agents, Theories, Architectures, and Languages* (ATAL), August 2001.
- Agents Workshop on *Infrastructure for Agents, Multi-Agent Systems, and Scalable Multi-Agent Systems*, May 2001.
- ECAI Workshop on *Balancing Reactivity and Social Deliberation in Multi-Agent Systems*, August 2000.
- ICMAS *Collective Robotics* Workshop, July 1998.
- IROS Workshop on *RoboCup*, November 1996.
- **Project reviewer:**
 - Science Foundation of Ireland (SFI) “Integrated Analysis of System of Systems,” 2009.
 - European Commission (EC) “Ubiquitous Networking Robotics in Urban Settings,” 2008, 2009.
 - UK EPSRC “Market-Based Control of Complex Computational Systems,” 2008.
- **Proposal reviewer:**
 - Army Research Office (ARO) 2011.
 - Instituto de Telecomunicações (IT), Portugal, 2011.
 - National Science Foundation (NSF), 2011, 2010, 2009, 2008, 2004, 2003, 2002, 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:**
 - 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.
 - Doctoral Admissions committee, 2009–10, 2007–08.

- 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.
- Doctoral Admissions committee, 2005–06.
- Ad hoc GSC committee on diversity proposals, 2005.
- Chair, Special Events committee, 2004–05.
- Evaluation of Graduate Programs committee, 2004–05.
- Faculty evaluation committee, 2003–04.
- Departmental Best Dissertation committee, 2003.

INVITED DISTINGUISHED LECTURES

- “Learning and Multiagent Reasoning for Autonomous Agents.”
UC Irvine Computer Science Department Distinguished Lecturer
in Irvine, California. October 2009.
- “Teaching Teammates in Ad Hoc Teams.”
Keynote talk at 10th Bar-Ilan Symposium on the Foundations of Artificial Intelligence (BISFAI)
in Ramat-Gan, Israel. June 2009.
- “Learning and Multiagent Reasoning for Autonomous Agents.”
Keynote talk at 4th IEEE Latin America Robotic Symposium (LARS) in Monterrey Mexico.
November 2007.
- “Learning and Multiagent Reasoning for Autonomous Agents.”
UT Austin Visions of Computing Lecture in Austin, Texas. November 2007.
- “Learning and Multiagent Reasoning for Autonomous Agents.”
IJCAI Computers and Thought Award Lecture in Hyderabad, India. January 2007.
- “Robot Learning.”
National Academy of Sciences spring symposium in Washington, DC. April 2006.
- “Machine Learning on Physical Robots.”
Keynote talk at International Conference on Computing (CIC) in Mexico City. October 2004.
- “The Trading Agent Competition: Two Champion Adaptive Bidding Agents.”
Computer Games Conference in Edmonton, Alberta. July, 2002.

INVITED TALKS

- “Intersections of the Future: Using Fully Autonomous Vehicles”
AAMAS 2011 Workshop on Agents and Data Mining Interaction in Taipei, Taiwan. May 2011.
- “Human-Assisted Reinforcement Learning”
ONR Science of Autonomy Meeting in Arlington, VA. April 2011.
- “Machine Learning and Multiagent Reasoning: from robot soccer to autonomous traffic”
Texas State Computer Science Department Seminar in San Marcos, TX. March 2011.
- “Machine Learning and Multiagent Reasoning: from robot soccer to autonomous traffic”
Virginia Tech Computer Science Department Seminar in Blacksburg, VA. March 2011.
- “Machine Learning and Multiagent Reasoning: from robot soccer to autonomous traffic”
Baylor Physics Department Seminar in Waco, TX. December 2010.
- “Machine Learning and Multiagent Reasoning: From Robot Soccer to Autonomous Traffic”
Johns Hopkins Computer Science Seminar in Baltimore, MD. November 2010.
- “Autonomous Robots Playing Soccer and Traversing Intersections”
Hot Science - Cool Talks Outreach Lecture Series in Austin, Texas. October 2010.
- “Learning and Multiagent Reasoning for Autonomous Agents”
Blue Knot Austin in Austin, Texas. July 2010.
- “Intersections of the Future: Using Fully Autonomous Vehicles”
Taiwan Agent Summer School in Hsinchu, Taiwan. June 2010.

- “Learning and Multiagent Reasoning for Autonomous Agents”
Taiwan Agent Summer School in Hsinchu, Taiwan. June 2010.
- “Machine Learning on Physical Robots”
Seminar in **Mechanical Engineering Department**, UT Austin. April 2010.
- “Autonomous Robots Playing Soccer and Traversing Intersections”
TEDxUT in Austin, Texas. April 2010.
- “Learning and Multiagent Reasoning for Autonomous Agents”
National Instruments in Austin, Texas. January 2010.
- “Intersections of the Future: Using Fully Autonomous Vehicles”
IEEE Latin-American Summer School on Robotics in Santiago, Chile. December 2009.
- “Learning and Multiagent Reasoning for Autonomous Agents”
IEEE Latin-American Summer School on Robotics in Santiago, Chile. December 2009.
- “Intersections of the Future: Using Fully Autonomous Vehicles”
Robotics Seminar at **University Polytechnica de Catalunya** in Barcelona, Spain. November 2009.
- “Teaching Teammates in Ad Hoc Teams”
UT Department of Computer Sciences Forum for AI and UTCS Colloquium in Austin, Texas. September 2009.
- “Learning and Multiagent Reasoning for Autonomous Agents.”
Technion CS Department Pixel Club lecture in Haifa, Israel. June 2009.
- “How Machines Learn: From Robot Soccer to Autonomous Traffic”
HEMDA Center for Science Education in Tel Aviv, Israel. June 2009.
- “Teaching Teammates in Ad Hoc Teams.”
Game theory seminar at **Hebrew University Center for Rationality** in Jerusalem, Israel. May 2009.
- “Teaching Teammates in Ad Hoc Teams.”
AAMAS 2009 workshop on Adaptive Learning Agents in Budapest, Hungary. May 2009.
- “Generalization in Reinforcement Learning.”
Hebrew University Machine Learning Club Talk in Jerusalem, Israel. April 2009. 4/23/09
- “Machine Learning on Physical Robots.”
Hebrew University CS Coloquium in Jerusalem, Israel. April 2009.
- “Generalization in Reinforcement Learning.”
Technion EE Guest Lecture in Haifa, Israel, March 2009.
- “Machine Learning on Physical Robots.”
Haifa Mini-Workshop on Machine Learning: Theory and Practice in Haifa, Israel. March 2009.
- “Learning and Multiagent Reasoning for Autonomous Agents.”
Ben Gurion University in Beér Sheva, Israel. February 2009.
- “Learning and Multiagent Reasoning for Autonomous Agents.”
IBM Haifa Research Lab in Haifa, Israel. December 2009.
- “Learning and Multiagent Reasoning for Autonomous Agents.”
The Israel Association for Artificial Intelligence Symposium in Ashkelon, Israel. November 2008.
- “Learning and Multiagent Reasoning for Autonomous Agents.”
International Workshop on Market-Based Control of Complex Computational Systems in Liverpool, UK. September 2008.
- “Learning and Multiagent Reasoning for Autonomous Agents.”
University of Alberta AI Seminar in Edmonton, Alberta. March 2007.
- “Learning and Multiagent Reasoning for Autonomous Agents.”
University of Southern Alabama in Mobile, Alabama. February 2007.
- “Embracing Mobility.”
DARPA kickoff meeting on Information Theory for Mobile Ad-Hoc Networks in Chicago, IL. November 2006.
- “Layered Learning on Physical Robots.”
University of Amsterdam in Amsterdam, The Netherlands. November 2006.
- “Robust Autonomous Color Learning on a Mobile Robot.”
Robotics Institute Seminar Series, **Carnegie Mellon University** in Pittsburgh, PA. October 2006.
- “Robust Autonomous Color Learning on a Mobile Robot.”
Center for Perceptual Systems Seminar Series, UT Austin. 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. September 2006.
- “State Abstraction Discovery, and Layered Learning on Physical Robots.”
AAMAS workshop on Hierarchical Autonomous Agents and Multiagent Systems in Hakodate, Japan. May 2006.
- “Machine Learning on Physical Robots.”
GRASP seminar series at **University of Pennsylvania**. March 2006.
- “Robot Learning.”
National Academy of Sciences Frontiers of Science Symposium in 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. September 2005.
- “Machine Learning and Multiagent Systems: From robot soccer to autonomic computing.”
IBM Technical Vitality Council in Austin, TX. September 2005.
- “The Trading Agent Competition: Two Champion Adaptive Bidding Agents.”
Intelligent Systems Seminar Series, **McCombs School of Business**, UT Austin. September 2005.
- “Practical Vision-Based Monte Carlo Localization on a Legged Robot.”
IJCAI Workshop on Reasoning with Uncertainty in Robotics in Edinburgh, Scotland. July 2005.
- “RoboCup as an Introduction to Multiagent Systems and Research.”
AAMAS Workshop on Teaching Multiagent Systems in Utrecht, Netherlands. July 2005.
- “Multi-Robot Learning for Continuous Area Sweeping.”
AAMAS Workshop on Learning and Adaptation in Multiagent Systems in Utrecht, Netherlands. July 2005.
- “Speeding up Reinforcement Learning with Behavior Transfer.”
RoboCup US Open in Atlanta, GA. May 2005.
- “Scaling Up Reinforcement Learning via Task Transfer.”
DARPA bidder’s conference on Transfer Learning in Washington, DC. March 2005.
- “Machine Learning on Physical Robots.”
Bar Ilan University in Ramat Gan, Israel. December 2004.
- “Adversarial Agents and Other Agent Topics.”
Air Force Research Lab in Rome, New York. July 2004.
- “Embodied Agents.”
Americas’ School on Agents and Multiagent Systems in New York, New York. July 2004.
- “Coaching, Advising, and Task Transfer for Multiagent Learning.”
Information Science and Technology (ISAT) study on Multiagent/Distributed Learning in Boston, Massachusetts. June 2004.
- “Robot soccer: competitions and research.”
National Instruments Scholarship for Excellence reception in Austin, Texas. June 2004.
- “The RoboCup Challenge: Progress and Research Results in Robot Soccer.”
UT Mechanical Engineering Dept. Robotics Research Group in Austin, Texas. May 2004.
- “The RoboCup Challenge: Progress and Research Results in Robot Soccer.”
UT Department of Computer Sciences Forum for AI in Austin, Texas. November 2003.
- “Policy Gradient Reinforcement Learning for Fast Quadrupedal Locomotion.”
DARPA PI Meeting on Navigation, Locomotion, and Articulation in Washington, DC. Nov. 2003.
- “Autonomous Learning Agents in Dynamic, Multiagent Environments.”
IROS-2003 Workshop on Learning and Evolution in MAS in Las Vegas, Nevada. October, 2003.
- “The RoboCup Challenge: Progress and Research Results in Robot Soccer.”
University of Science and Technology China (USTC) in Hefei, China. October, 2003.
- “Layered Learning towards Autonomic Computing.”
IJCAI-2003 Workshop on AI and Autonomic Computing in Acapulco, Mexico. August 2003.

- “Autonomous Bidding Agents and the Power of Threats.” (with Michael Littman)
CMU Machine Learning Lunch Seminar in Pittsburgh, Pennsylvania. April 2003.
- “Machine Learning Research in the RoboCup Simulation League.”
First RoboCup American Open Workshop in Pittsburgh, Pennsylvania. April 2003.
- “Autonomous Learning Agents in Dynamic, Multiagent Environments.”
UT College of Natural Sciences Advisory Council Meeting in Austin, Texas. April 2003.
- “Scaling Reinforcement Learning toward RoboCup Soccer.”
NIPS 2002 workshop on Multi-Agent Learning in Whistler, British Columbia. December 2002.
- “Autonomous Learning Agents in Dynamic, Multiagent Environments.”
University of Texas School of Library and Information Science in Austin, Texas. September 2002.
- “The Trading Agent Competition: Two Champion Adaptive Bidding Agents.”
UT Department of Computer Sciences Forum for AI in Austin, Texas. September 2002.
- “Multiagent Competitions and Research: Lessons from RoboCup and TAC.”
Trading Agent Competition Workshop in Edmonton, Alberta. July, 2002.
- “Autonomous Learning Agents in Dynamic, Multiagent Environments: Auctions and Soccer.”
Santa Fe Institute Collective Cognition Workshop in Santa Fe, New Mexico. January 2002.
- “Autonomous Learning Agents in Dynamic, Multiagent Environments: Auctions and Soccer.”
University of Alberta AI Seminar in Edmonton, Alberta. November 2001.
- “Autonomous Bidding Agents.”
Brookings Institution Workshop on Multi-Agent Computation in Natural and Artificial Economies in Washington, DC. October 2001.
- “Layered Learning in Multi-Agent Systems: A Winning Approach to Robotic Soccer.”
IEEE Computer Society and DigiPen Institute of Technology in Seattle, Washington. July 2001.
- “Layered Learning in Multi-Agent Systems: A Winning Approach to Robotic Soccer.”
The Boeing Company in Seattle, Washington. July 2001.
- “Layered Learning in Multi-Agent Systems.”
Multi-Strategy Learning Workshop in Guimarães, Portugal. June 2000.
- “The RoboCup Challenge.”
NASA Goddard Space Flight Center in Greenbelt, Maryland. March 2000.
- “The RoboCup Challenge.”
Ohio University in Athens, Ohio. February 2000.
- “The RoboCup Challenge.”
IEEE Real-Time Systems Symposium in Phoenix, Arizona. December, 1999.
- “Layered Learning in Multi-Agent Systems.”
Machines That Learn Workshop in Snowbird, Utah. April 1998.
- “Layered Learning in Multi-Agent Systems.”
SRI International in Palo Alto, California. November 1997.
- “Layered Learning in Multi-Agent Systems.”
University of Washington in Seattle, Washington. November 1997.
- “Task Decomposition and Dynamic Role Assignment for Real-Time Strategic Teamwork.”
Electrotechnical Laboratory (ETL) in Tsukuba, Japan. August 1997.
- “Machine Learning for Agent Control in Real-time Multi-Agent Domains.”
Workshop on Intelligent Robotic Agents in Porto Alegre, Brazil. March 1997.
- “Layered Learning in the Soccer Server.”
Electrotechnical Laboratory (ETL) in Tsukuba, Japan. November 1996.
- “Layered Learning in the RoboCup Soccer Server.”
Osaka University in Osaka, Japan. November 1996.
- “Building a Dedicated Robotic Soccer System.”
Korean Advanced Institute of Science and Technology (KAIST) in Taejon, Korea. August 1996.
- “Towards Collaborative and Adversarial Learning: A Case Study in Robotic Soccer.”
Naval Research Labs (NRL) in Washington DC. July 1996.
- “FLECS: Planning with a Flexible Commitment Strategy.”
Jet Propulsion Laboratory (JPL) in Pasadena, California. July 1995.
- “FLECS: Planning with a Flexible Commitment Strategy.”
USC Intelligent Software Institute (ISI) in 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. 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.
2. Michael P. Wellman, Amy Greenwald, and **Peter Stone**. *Autonomous Bidding Agents: Strategies and Lessons from the Trading Agent Competition*. MIT Press, 2007. (monograph)
3. **Peter Stone**. *Intelligent Autonomous Robotics: A Robot Soccer Case Study*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers, 2007. (monograph)
4. **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.
5. **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.
6. **Peter Stone**. *Layered Learning in Multiagent Systems: A Winning Approach to Robotic Soccer*. MIT Press, 2000. (monograph)

Journal Articles

7. Shivaram Kalyanakrishnan and **Peter Stone**. Characterizing reinforcement learning methods through parameterized learning problems. *Machine Learning (MLJ)*, 2011. To appear.
8. 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)*, 2011. To appear.
9. Tobias Jung, Daniel Polani, and **Peter Stone**. Empowerment for continuous agent-environment systems. *Adaptive Behavior*, 19(1):16–39, 2011.
10. 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.
11. Jonathan Wildstrom, **Peter Stone**, and Emmett Witchel. Autonomous return on investment analysis of additional processing resources. *International Journal on Autonomic Computing*, 1(3), 2010.
12. Matthew E. Taylor and **Peter Stone**. Transfer learning for reinforcement learning domains: A survey. *Journal of Machine Learning Research*, 10(1):1633–1685, 2009.
13. 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*, 21(1):1–27, 2010.
14. 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.
15. Juhyun Lee, W. Bradley Knox, and **Peter Stone**. Inter-classifier feedback for human-robot interaction in a domestic setting. *Journal of Physical Agents*, 2(2):41–50, July 2008. Special Issue on Human Interaction with Domestic Robots.

16. 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*, 2(1):15–30, March 2008. Special issue on Multi-Robot Systems.
17. Kurt Dresner and **Peter Stone**. A multiagent approach to autonomous intersection management. *Journal of Artificial Intelligence Research*, 31:591–656, March 2008.
18. 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**.
19. Matthew E. Taylor, **Peter Stone**, and Yaxin Liu. Transfer learning via inter-task mappings for temporal difference learning. *Journal of Machine Learning Research*, 8(1):2125–2167, 2007.
20. 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.
21. Mohan Sridharan and **Peter Stone**. Structure-Based Color Learning on a Mobile Robot under Changing Illumination. *Autonomous Robots*, 23(3):161–182, 2007.
22. **Peter Stone**. Multiagent learning is not the answer. it is the question. *Artificial Intelligence*, 171:402–405, 2007.
23. Shimon Whiteson, Matthew E. Taylor, and **Peter Stone**. Empirical studies in action selection for reinforcement learning. *Adaptive Behavior*, 15(1):33–50, 2007.
24. **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.
25. 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.
26. Shimon Whiteson and **Peter Stone**. Evolutionary Function Approximation for Reinforcement Learning. *Journal of Machine Learning Research (JMLR)*, 7:877–917, May 2006.
27. 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.
28. **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.
29. 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.

30. 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.
31. **Peter Stone** and Amy Greenwald. The first international trading agent competition: Autonomous bidding agents. *Electronic Commerce Research (EC)*, 5(2):229–65, April 2005.
32. 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.
33. 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.
34. **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.
35. 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.
36. 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.
37. **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.
38. **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.
39. Michael Bowling, Manuela Veloso, and **Peter Stone**. The CMUnited-98 champion small-robot team. *Advanced Robotics*, 13(8):753–766, 2000.
40. 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.
41. **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.
42. **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.
43. 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.

44. **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.
45. 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

46. 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. To appear.
47. 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.
48. 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.
49. **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.
50. 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.
51. **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.
52. 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 Analysis*, volume 44 of *Lecture Notes in Business Information Processing (LNBIP)*, pages 175–89. 2010.
53. 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.**
54. 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. To appear.
55. 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.
56. 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.
57. 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.

58. 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.
59. 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. To appear. **Winner of Best Paper Award**.
60. 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. To appear.
61. 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.
62. 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.
63. 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**.
64. 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.
65. 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.
66. 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**.
67. 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.
68. 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.
69. **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.
70. 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.

71. 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.
72. 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.
73. 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.
74. 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.
75. 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.
76. 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.
77. **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.
78. **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.
79. **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.
80. 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.
81. **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.
82. **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.
83. 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.
84. **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).

85. **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.
86. **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.
87. **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.
88. **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.
89. **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.
90. **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

91. Noa Agmon, Daniel Urieli, and **Peter Stone**. Multiagent patrol generalized to complex environmental conditions. In *Proceedings of the Twenty-Fifth Conference on Artificial Intelligence*, August 2011. To appear.
92. Tsz-Chiu Au, Neda Shahidi, and **Peter Stone**. Enforcing Liveness in Autonomous Traffic Management. In *Proceedings of the Twenty-Fifth Conference on Artificial Intelligence*, August 2011. To appear.
93. 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*, August 2011. To appear.
94. 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’11)*, June 2011.
95. 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 2011)*, May 2011. To appear.
96. 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 2011)*, May 2011.
97. 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 2011)*, May 2011.
98. Michael Quinlan, Tsz-Chiu Au, Jesse Zhu, Nicolae Sturca, and **Peter Stone**. Bringing simulation to life: A mixed reality autonomous intersection. In *Proceedings of IROS 2010-IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2010)*, October 2010.

99. Tobias Jung and **Peter Stone**. Gaussian processes for sample efficient reinforcement learning with rmax-like exploration. In *The European Conference on Machine Learning*, September 2010.
100. Todd Hester and **Peter Stone**. Real time targeted exploration in large domains. In *The Ninth International Conference on Development and Learning (ICDL)*, August 2010.
101. 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 2010)*, June 2010.
102. 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)*, 2010.
103. David Pardoe and **Peter Stone**. Boosting for regression transfer. In *Proceedings of the 27th International Conference on Machine Learning (ICML 2010)*, June 2010.
104. **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*, July 2010.
105. 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 2010)*, May 2010. **Winner of Best Student Paper Award.**
106. 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 2010)*, May 2010.
107. **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.
108. 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.
109. Peter Djeu, Michael Quinlan, and **Peter Stone**. Improving particle filter performance using sse instructions. In *Proceedings of IROS 2009: 2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, October 2009.
110. 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.
111. 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*, September 2009.
112. 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*, September 2009.
113. 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*, pages 749–56, May 2009.
114. 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*, May 2009.

115. W. Bradley Knox and **Peter Stone**. TAMER: Training an Agent Manually via Evaluative Reinforcement. In *IEEE 7th International Conference on Development and Learning*, August 2008.
116. Doran Chakraborty and **Peter Stone**. Online multiagent learning against memory bounded adversaries. In *Machine Learning and Knowledge Discovery in Databases*, volume 5212 of *Lecture Notes in Artificial Intelligence*, pages 211–26, September 2008.
117. Matthew E. Taylor, Nicholas K. Jong, and **Peter Stone**. Transferring instances for model-based reinforcement learning. In *Machine Learning and Knowledge Discovery in Databases*, volume 5212 of *Lecture Notes in Artificial Intelligence*, pages 488–505, September 2008.
118. W. Bradley Knox and **Peter Stone**. Tamer: Training an agent manually via evaluative reinforcement. In *IEEE 7th International Conference on Development and Learning*, August 2008.
119. 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*, July 2008.
120. 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*, July 2008.
121. Jonathan Wildstrom, **Peter Stone**, and Emmett Witchel. CARVE: A Cognitive Agent for Resource Value Estimation. In *International Conference on Autonomic Computing*, June 2008.
122. 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.
123. 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*, May 2008.
124. 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*, May 2008.
125. Juhyun Lee and **Peter Stone**. Person tracking on a mobile robot with heterogeneous inter-characteristic feedback. In *IEEE International Conference on Robotics and Automation*, May 2008.
126. 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*, May 2008.
127. Todd Hester and **Peter Stone**. Negative information and line observations for monte carlo localization. In *IEEE International Conference on Robotics and Automation*, May 2008.
128. 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.
129. 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.
130. 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*, September 2007.
131. 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*, July 2007. Nectar Track.

132. Matthew E. Taylor and **Peter Stone**. Cross-domain transfer for reinforcement learning. In *Proceedings of the Twenty-Fourth International Conference on Machine Learning*, June 2007.
133. 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*, April 2007.
134. 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*, May 2007.
135. 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*, May 2007. **Nominee for Best Student Paper Award.**
136. Shivaram Kalyanakrishnan and **Peter Stone**. Batch reinforcement learning in a complex domain. In *The Sixth International Joint Conference on Autonomous Agents and Multiagent Systems*, May 2007. **Nominee for Best Student Paper Award.**
137. 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*, May 2007.
138. **Peter Stone**. Learning and multiagent reasoning for autonomous agents. In *The 20th International Joint Conference on Artificial Intelligence*, pages 13–30, January 2007.
139. 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*, pages 2212–2217, January 2007.
140. Kurt Dresner and **Peter Stone**. Sharing the road: Autonomous vehicles meet human drivers. In *The 20th International Joint Conference on Artificial Intelligence*, pages 1263–68, January 2007.
141. Bikramjit Banerjee and **Peter Stone**. General game learning using knowledge transfer. In *The 20th International Joint Conference on Artificial Intelligence*, pages 672–677, January 2007.
142. 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*, pages 1113–1118, January 2007.
143. 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.
144. 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.
145. 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.
146. 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.
147. 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.
148. 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.

149. 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.
150. 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.
151. 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.
152. Harish Subramanian, Subramanian Ramamoorthy, **Peter Stone**, and Ben Kuipers. Designing safe, profitable automated stock trading agents using evolutionary algorithms. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, July 2006.
153. 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.
154. David Pardoe and **Peter Stone**. Predictive planning for supply chain management. In *International Conference on Automated Planning and Scheduling (ICAPS)*, June 2006.
155. 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.
156. 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.
157. 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.
158. 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.
159. 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.
160. 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.
161. 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.
162. 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.
163. 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.
164. 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.

165. 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.
166. 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.
167. 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.
168. 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.
169. 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.
170. Nate Kohl and **Peter Stone**. Policy Gradient Reinforcement Learning for Fast Quadrupedal Locomotion. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, New Orleans, LA, May 2004.
171. 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.
172. 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.
173. 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.
174. 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.
175. 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.
176. **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.
177. 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.**
178. **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.
179. **Peter Stone**. TPOT-RL applied to network routing. In *Proceedings of the Seventeenth International Conference on Machine Learning (ICML)*, pages 935–942, 2000.
180. 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.

181. 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.
182. 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.
183. 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.
184. 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.
185. 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.
186. **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.
187. **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

188. Katie Genter, Noa Agmon, and **Peter Stone**. Role-Based Ad Hoc Teamwork. In *AAAI workshop on Plans and Activity Recognition (PAIR)*, August 2011. To appear.
189. 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. To appear.
190. W. Bradley Knox and **Peter Stone**. Augmenting Reinforcement Learning with Human Feedback. In *IJCAI Workshop on Agents Learning Interactively from Human Teachers (ALIHT)*, July 2011. To appear.
191. 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.
192. 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.
193. 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.
194. 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.

195. 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.
 196. 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.
 197. 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.
 198. 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.
 199. 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.
 200. Tobias Jung, Mazda Ahmadi, and **Peter Stone**. Connectivity-based localization in robot networks. In *International Workshop on Robotic Wireless Sensor Networks (WSN)*, June 2009.
 201. 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.
 202. Shimon Whiteson, Brian Tanner, Matthew E. Taylor, and **Peter Stone**. Generalized domains for empirical evaluations in reinforcement learning. In *ICML Workshop on Evaluation Methods for Machine Learning*, June 2009.
 203. 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.
 204. 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.
 205. David Pardoe and **Peter Stone**. The 2007 TAC SCM Prediction Challenge. In *AAAI 2008 Workshop on Trading Agent Design and Analysis (TADA)*, July 2008.
- ***NOTE: REPLACE WITH BOOK CHAPTER
206. 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.
 207. Kurt Dresner and **Peter Stone**. Learning policy selection for autonomous intersection management. In *AAMAS workshop on adaptive and learning agents*, May 2007.
 208. 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.
 209. Matthew E. Taylor, Gregory Kuhlmann, and **Peter Stone**. Accelerating search with transferred heuristics. In *ICAPS workshop on AI Planning and Learning*, September 2007.
 210. Nicholas K. Jong and **Peter Stone**. Model-based exploration in continuous state spaces. In *The Seventh Symposium on Abstraction, Reformulation, and Approximation*, July 2007.
 211. Jonathan Wildstrom, **Peter Stone**, and Emmett Witchel. Autonomous return on investment analysis of additional processing resources. In *2007 Workshop on Adaptive Methods in Autonomic Computing Systems*, June 2007.

212. 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.
213. 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.
214. 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.
215. Daniel Stronger and **Peter Stone**. Expectation-based vision for precise self-localization on a mobile robot. In *AAAI workshop on Cognitive Robotics*, July 2006.
216. Mazda Ahmadi and **Peter Stone**. Instance-based action models for fast action planning. In *AAAI workshop on Cognitive Robotics*, July 2006.
217. **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.
218. 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.
219. 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.
220. Subramanian Ramamoorthy, Harish Subramanian, **Peter Stone**, and Ben 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.
221. 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.
222. 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.
223. 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.
224. 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.
225. 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.
226. 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.
227. 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.
228. **Peter Stone**. Multiagent learning for autonomous spacecraft constellations. In *Proceedings of the NASA Workshop on Planning and Scheduling for Space*, 1997.

229. 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.
230. 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.
231. **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.
232. **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.
233. **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.
234. **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

235. David Pardoe and **Peter Stone**. Designing Adaptive Trading Agents. *ACM SIGecom Exchanges*, 10(2), June 2011.
236. Matthew E. Taylor and **Peter Stone**. An Introduction to Inter-task Transfer for Reinforcement Learning. *AI Magazine*, 32(1):15–20, March 2011.
237. David Pardoe and **Peter Stone**. Developing Adaptive Auction Mechanisms. *ACM SIGecom Exchanges*, 5(3):1–10, April 2005.
238. David Pardoe and **Peter Stone**. TacTex-03: A Supply Chain Management Agent. *ACM SIGecom Exchanges*, 4(3):19–28, Winter 2004.
239. Amy Greenwald, Nicholas R. Jennings, and **Peter Stone**. Guest Editors' Introduction: Agents and Markets. *IEEE Intelligent Systems*, 18(6):12–14, November/December 2003.
240. 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. Hajiaghahi, and E. Chiniforooshan. RoboCup-2001: The fifth robotic soccer world championships. *AI Magazine*, 23(1):55–68, 2002.
241. **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.
242. Amy Greenwald and **Peter Stone**. Autonomous bidding agents in the trading agent competition. *IEEE Internet Computing*, 5(2), 2001.
243. Silvia Coradeschi, Lars Karlsson, **Peter Stone**, Tucker Balch, Gerhard Kraetzschmar, and Minoru Asada. Overview of RoboCup-99. *AI Magazine*, 21(3), 2000.
244. 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

245. 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 <http://www.cs.utexas.edu/~AustinVilla/?p=athome>

Unrefereed Publications

246. Todd Hester, Michael Quinlan, and **Peter Stone**. A real-time model-based reinforcement learning architecture for robot control. Technical Report arXiv e-Prints 1105.1749, arXiv, May 2011.
247. 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.
248. 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.
249. 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.
250. 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.
251. Shivaram Kalyanakrishnan and **Peter Stone**. UT Austin Villa 3D simulation soccer team 2008. In *RoboCup Symposium CD Proceedings*, 2008.
252. Peter Stone Todd Hester, Michael Quinlan. UT Austin Villa 2008 team description paper for the standard platform league. In *RoboCup Symposium CD Proceedings*, 2008.
253. **Peter Stone** and Neda Shahidi. The UT Austin Villa 2008 mixed reality team. In *RoboCup Symposium CD Proceedings*, 2008.
254. **Peter Stone**, Patrick Beeson, Tekin Mericli, and Ryan Madigan. DARPA urban challenge technical report: Austin Robot Technology, June 2007.
255. 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.
256. 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.
257. **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.
258. **Peter Stone**, Todd Hester, Tekin Mericli, Neda Shahidi, and Mohan Sridharan. The ut austin villa 2007 four-legged team description paper. In *RoboCup Symposium CD Proceedings*, 2007.
259. **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.

260. 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.
261. **Peter Stone**, Kurt Dresner, Peggy Fiedelman, 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.
262. 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.
263. **Peter Stone**, Kurt Dresner, Peggy Fiedelman, 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.
264. 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.
265. **Peter Stone**, Kurt Dresner, Selim T. Erdogan, Peggy Fiedelman, 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.
266. **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.
267. 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.
268. **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.
269. 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.
270. **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.
271. **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.
272. **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.

273. 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.
274. 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.
275. 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

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

PRESS

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