

**Michael L. Littman, Ph.D.**

52 Church St.  
Barrington, RI 02806  
(908) 229-8804

**Employment**

- 2012–: Professor of Computer Science, Brown University, Providence, RI. (Associate Chair, 2018–2019 Academic Year.)
- 2013–: Adjunct Professor of Interactive Computing, Georgia Institute of Technology, Atlanta, GA.
- 2009–2012: Chair of the Department of Computer Science, Rutgers University, Piscataway, NJ.
- 2003–2012: Professor of Computer Science, Rutgers University, Piscataway, NJ. Tenure granted (2004). Member of the RUCCS (Cognitive Science) Faculty (2004). Member of the Winlab Faculty (2007). Promoted to Full Professor (2008). Member of the ECE Graduate Faculty (2010).
- 2002–2003: Associate Research Professor of Computer Science, Rutgers University, Piscataway, NJ.
- 2002–2004: Consultant, Bernardsville, NJ. Clients: Knowledge Analysis Technologies, Boulder, CO. University of Pennsylvania, Philadelphia, PA. Ornarose, Inc., Madison, NJ.
- 2001–2002: Lecturer in Computer Science, Princeton University, Princeton, NJ.
- 2000–2002: Principal Technical Staff Member, Artificial Intelligence Principles Research Department, AT&T Labs–Research, Florham Park, NJ.
- 1999–: Officer (Co-Founder and CEO), OneAcross.com, Harrisonburg, VA.
- 1996–2000: Assistant Professor of Computer Science, Duke University, Durham, NC.
- 1992–1996: Member of Technical Staff, Computer Graphics and Interactive Media, Bellcore, Morristown, NJ.
- 1988–1992: Member of Technical Staff, Cognitive Science Research Group, Bellcore, Morristown, NJ.

## Education

- Ph.D. in Computer Science, Brown University, 1996. Advisor: Leslie Pack Kaelbling. Thesis title: Algorithms for sequential decision making.
- Graduate study in Computer Science, Carnegie Mellon University, 1993. Advisor: Avrim Blum.
- M.S. in Computer Science, Yale University, 1988. Advisor: Marina Chen. Thesis title: An exploration of asynchronous data-parallelism.
- B.S. in Computer Science, Yale University, 1988. Graduated summa cum laude with exceptional distinction in Computer Science.

## Awards and Honors

- Inducted as an ACM Fellow for contributions to the design and analysis of sequential decision making algorithms in artificial intelligence, 2018.
- The Philip J. Bray Award for Excellence in Teaching in the Physical Sciences, Brown University, 2018.
- TEDxBoston, “A Cooperative Path to Artificial Intelligence”, 2017.
- TEDxProvidence, “How I Learned to Stop Worrying and Be Realistic About AI”, 2017.
- Invited Keynote Speaker, 24th International Joint Conference on Artificial Intelligence (IJCAI-15), Buenos Aires, Argentina, 2015.
- IFAAMAS Influential Paper Award, AAMAS-2015: “Markov games as a framework for multi-agent reinforcement learning” (published in ICML-1994).
- Popular Youtube channel (<http://www.youtube.com/mlittman>) with computer science research videos and educational music videos (over 440,000 views as of August, 2018).
- “Dancing with the Profs” winner, Brown University, 2015.
- The “Shakey” Award for Most Entertaining Video, AAAI 2014 (“Overfitting: Machine Learning Music Video”) with Charles Isbell and Aaron Gross.
- Classic Paper Award, AAAI-13: “Acting Optimally in Partially Observable Stochastic Domains” (published in AAAI-94 by Cassandra, Kaelbling, and Littman).
- Invited Keynote Speaker, International Conference on Grammatical Induction, Baltimore, MD, 2012.
- Warren I. Susman Award for Excellence in Teaching, Rutgers University’s highest honor for outstanding and innovative performance in the classroom by a tenured faculty member, 2011.

- Winner of Lemonade Stand Game 2011 (with Wunder, Kaisers, Yaros).
- Inducted as a Fellow of the Association for the Advancement of Artificial Intelligence, 2010.
- The “Shakey” Award for Best AI Video Narration, IJCAI 2009 (“Real Live Robot Learning”) with Kaushik Subramanian.
- Speaker, Brown University, Distinguished Lecture Series, Providence, RI, 2009.
- Speaker, University of British Columbia, ICICS/CS Distinguished Lecture Series, Vancouver, BC, Canada, 2009.
- Invited Keynote Speaker, Quantitative Evaluation of Systems (QEST) Conference, St. Malo, France, 2008.
- Best Student Paper Award (with Munoz de Cote), Uncertainty in Artificial Intelligence, 2008 (“A Polynomial-time Nash Equilibrium Algorithm for Repeated Stochastic Games”).
- Best Student Paper Award (with Li, Walsh), International Conference on Machine Learning, 2008 (“Knows What It Knows: A Framework for Self-Aware Learning”).
- Rutgers University 2008 SAS Award for Distinguished Contributions to Undergraduate Education.
- The First “Shakey” Award for Best AI Video Short, AAAI 2007 (“Artificial Intelligence: An Instance of Aibo Ingenuity”).
- The First Annual Reinforcement Learning Competition, First Prize, Pentathlon (with Li, Nouri, Walsh), 2006.
- Connecticut–Trinity–Wesleyan Computer Science Consortium Joint Distinguished Lecture Series, 2006.
- Eclipse Innovation Grant 2005, IBM.
- Best Paper Award (with Strehl), International Conference on Tools for Artificial Intelligence, 2004 (“An Empirical Evaluation of Interval Estimation for Markov Decision Processes”).
- Teacher of the Year, 2003–2004, Rutgers Computer Science Graduate Student Society.
- Invited Alumnus Speaker, Brown University Department of Computer Science, 25th Anniversary Symposium and Celebration Banquet, 2004.
- Invited Keynote Speaker, The Seventeenth Canadian Conference on Artificial Intelligence, Ontario, Canada, 2004.

- Honorable Mention (with Goldsmith, Mundhenk), IJCAII-JAIR Best Paper Prize, 2003 (“The Computational Complexity of Probabilistic Planning”).
- Winner of ICMAS-00 Trading Agent Competition (for ATTac, with Stone, Singh, Kearns) and the 2003 Trading Agent Competition (for ATTac-01, with Stone, Schapire, McAllester, Csirik).
- Distinguished Graduate Award in Academics, Plymouth Whitemarsh High School, Plymouth Meeting, PA, 2002.
- Invited Keynote Speaker, Second International Conference on Computers and Games, Hamamatsu, Japan, 2000.
- Robert B. Cox Trinity College Distinguished Teaching Award, Duke University, 1999.
- Best Paper Award (with Keim, Shazeer, Agarwal, Cheves, Fitzgerald, Grosland, Jiang, Pollard, Weinmeister), American Association of Artificial Intelligence, 1999 National Conference (“Proverb: The Probabilistic Cruciverbalist”).
- National Science Foundation CAREER Award, 1997 (“Planning Under Uncertainty in Large Domains”).

## Journal Publications

- [1] Bei Peng, James MacGlashan, Robert Loftin, Michael L Littman, David L Roberts, and Matthew E Taylor. Curriculum design for machine learners in sequential decision tasks. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 2018.
- [2] Mark K. Ho, James MacGlashan, Michael L. Littman, and Fiery Cushman. Social is special: A normative framework for teaching with and learning from evaluative feedback. *Cognition*, 167:91–106, 2017.
- [3] Adam Morris, James MacGlashan, Michael L. Littman, and Fiery Cushman. Evolution of flexibility and rigidity in retaliatory punishment. *Proceedings of the National Academy of Sciences*, 2017.
- [4] Douglas H. Fisher, Jr. Charles L. Isbell, and Michael L. Littman. Ask me anything about MOOCs. *AI Magazine*, 38(2), 2017.
- [5] C. W. Kahler, W. V. Lechner, J. MacGlashan, T. B. Wray, and M. L. Littman. Initial progress towards development of a voice-based computer delivered motivational intervention for heavy drinking college students: An experimental study. *JMIR Mental Health*, 2017.
- [6] Mark K Ho, Michael L. Littman, James MacGlashan, Fiery Cushman, and Joseph L. Austerweil. Showing versus doing: Teaching by demonstration. In *Advances In Neural Information Processing Systems*, pages 3027–3035, 2016.

- [7] Mark K. Ho, James MacGlashan, Amy Greenwald, Elizabeth M. Hilliard, Carl Trim-bach, Stephen Brawner, Joshua B. Tenenbaum, Max Kleiman-Weiner, and Joseph L. Austerweil. Feature-based joint planning and norm learning in collaborative games. In *Proceedings of the 38th Annual Meeting of the Cognitive Science Society*, 2016.
- [8] Max Kleiman-Weiner, Mark K. Ho, Joseph L. Austerweil, Michael L. Littman, and Joshua B. Tenenbaum. Coordinate to cooperate or compete: Abstract goals and joint intentions in social interaction. In *Proceedings of the 38th Annual Meeting of the Cognitive Science Society*, 2016.
- [9] Robert T. Loftin, Bei Peng, James MacGlashan, Michael L. Littman, Matthew E. Taylor, Jeff Huang, and David L. Roberts. Learning behaviors via human-delivered discrete feedback: Modeling implicit feedback strategies to speed up learning. *Autonomous Agents and Multi-Agent Systems*, 30(1):30–59, 2016.
- [10] Michael L. Littman. Reinforcement learning improves behaviour from evaluative feedback. *Nature*, 521(7553):394–556, 2015.
- [11] Nikos Vlassis, Michael L. Littman, and David Barber. On the computational complexity of stochastic controller optimization in POMDPs. *ACM Transactions on Computation Theory*, 4(4), 2012.
- [12] Michael L. Littman. Technical perspective: A new way to search game trees. *Communications of the ACM*, 55(3):105, 2012.
- [13] Changhe Yuan, Heejin Lim, and Michael L. Littman. Most Relevant Explanation: Computational complexity and approximation methods. *Annals of Mathematics and Artificial Intelligence*, 61(3), 2011.
- [14] Shimon Whiteson and Michael L. Littman. Introduction to the special issue on empirical evaluations in reinforcement learning. *Machine Learning*, 84(1–2), 2011.
- [15] Lihong Li, Michael L. Littman, Thomas J. Walsh, and Alexander L. Strehl. Knows what it knows: A framework for self-aware learning. *Machine Learning*, 82(3):399–443, 2011.
- [16] Fusun Yaman, Thomas J. Walsh, Michael L. Littman, and Marie desJardins. Democratic approximation of lexicographic preference models. *Artificial Intelligence*, 175(7–8):1290–1307, 2011.
- [17] Merlise Clyde, Joyee Ghosh, and Michael Littman. Bayesian adaptive sampling for variable selection and model averaging. *Journal of Computational and Graphical Statistics*, 20(1):80–101, 2011.
- [18] Michael L. Littman and Daniel Reeves. Puzzle: Baffling raffling. *SIGecom Exchanges*, 10(1), 2011.

- [19] Brian Russell, Michael L. Littman, and Wade Trappe. Integrating machine learning in ad hoc routing: A wireless adaptive routing protocol. *International Journal of Communication Systems*, 24(7):950–966, 2011.
- [20] Lihong Li and Michael L. Littman. Reducing reinforcement learning to KWIK online regression. *Annals of Mathematics and Artificial Intelligence*, 58(3–4):217–237, 2010.
- [21] Ali Nouri and Michael L. Littman. Dimension reduction and its application to model-based exploration in continuous spaces. *Machine Learning*, 81(1):85–98, 2010.
- [22] Michael L. Littman. A tutorial on partially observable Markov decision processes. *Journal of Mathematical Psychology*, 53(3):119–125, 2009.
- [23] Alexander L. Strehl, Lihong Li, and Michael L. Littman. Reinforcement learning in finite MDPs: PAC analysis. *Journal of Machine Learning Research*, 10:2413–2444, 2009.
- [24] Emma Brunskill, Bethany R. Leffler, Lihong Li, Michael L. Littman, and Nicholas Roy. Provably efficient learning with typed parametric models. *Journal of Machine Learning Research*, 10:1955–1988, 2009.
- [25] Thomas J. Walsh, Ali Nouri, Lihong Li, and Michael L. Littman. Learning and planning in environments with delayed feedback. *Autonomous Agents and Multi-Agent Systems*, 18(1):83–105, 2009.
- [26] Alexander L. Strehl and Michael L. Littman. An analysis of model-based interval estimation for Markov decision processes. *Journal of Computer and System Sciences*, 74:1309–1331, 2008. Special issue on Learning Theory.
- [27] David L. Roberts, Charles L. Isbell, Jr., and Michael L. Littman. Optimization problems involving collections of dependent objects. *Annals of Operations Research, Special Issue on Decision Theory and Computer Science*, 163(1):255–270, 2008.
- [28] Andreas Buja, Deborah F. Swayne, Michael L. Littman, Nathaniel Dean, Heike Hofmann, and Lisha Chen. Data visualization with multidimensional scaling. *Journal of Computational and Graphical Statistics*, 444–472:29, 2008.
- [29] Amy Greenwald and Michael L. Littman. Introduction to the special issue on learning and computational game theory. *Machine Learning*, 67(1–2), 2007.
- [30] Martin Zinkevich, Amy Greenwald, and Michael L. Littman. A hierarchy of prescriptive goals for multiagent learning. *Artificial Intelligence*, 171:440–447, 2007.
- [31] Håkan L. S. Younes, Michael L. Littman, David Weissman, and John Asmuth. The First Probabilistic Track of the International Planning Competition. *Journal of Artificial Intelligence Research*, 24:851–887, 2005.

- [32] Peter D. Turney and Michael L. Littman. Corpus-based learning of analogies and semantic relations. *Machine Learning*, 60(1–3):251–278, 2005.
- [33] Michael L. Littman and Peter Stone. A polynomial-time Nash equilibrium algorithm for repeated games. *Decision Support Systems*, 39(1):55–66, 2005.
- [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*, 19:209–242, 2003.
- [35] Peter D. Turney and Michael L. Littman. Measuring praise and criticism: Inference of semantic orientation from association. *ACM Transactions on Information Systems (TOIS)*, 21(4):315–346, 2003.
- [36] Stephen M. Majercik and Michael L. Littman. Contingent planning under uncertainty via stochastic satisfiability. *Artificial Intelligence*, 147(1–2):119–162, 2003.
- [37] Michael L. Littman, Greg A. Keim, and Noam Shazeer. A probabilistic approach to solving crossword puzzles. *Artificial Intelligence*, 134(1–2):23–55, 2002.
- [38] Peter Stone, Michael L. Littman, Satinder Singh, and Michael Kearns. ATTac-2000: An adaptive autonomous bidding agent. *Journal of Artificial Intelligence Research*, 16:189–206, 2001.
- [39] Michael L. Littman, Stephen M. Majercik, and Toniann Pitassi. Stochastic Boolean satisfiability. *Journal of Automated Reasoning*, 27(3):251–296, 2001.
- [40] Michael L. Littman. Value-function reinforcement learning in Markov games. *Cognitive Systems Research*, 2(1):55–66, 2001.
- [41] Satinder Singh, Tommi Jaakkola, Michael L. Littman, and Csaba Szepesvári. Convergence results for single-step on-policy reinforcement-learning algorithms. *Machine Learning*, 39:287–308, 2000.
- [42] Michael L. Littman. Computers and language games. *IEEE Intelligent Systems*, 14(6):17–18, November/December 1999. In *Trends and Controversies, Playing with AI* (Haym Hirsch, editor).
- [43] Csaba Szepesvári and Michael L. Littman. A unified analysis of value-function-based reinforcement-learning algorithms. *Neural Computation*, 11(8):2017–2059, 1999.
- [44] Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1–2):99–134, 1998.
- [45] Michael L. Littman, Judy Goldsmith, and Martin Mundhenk. The computational complexity of probabilistic planning. *Journal of Artificial Intelligence Research*, 9:1–36, 1998.

- [46] Eugene Charniak, Glenn Carroll, John Adcock, Anthony Cassandra, Yoshihiko Gotoh, Jeremy Katz, Michael Littman, and John McCann. Taggers for parsers. *Artificial Intelligence*, 85(1–2):45–57, August 1996.
- [47] Leslie Pack Kaelbling, Michael L. Littman, and Andrew W. Moore. Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4:237–285, 1996.
- [48] Kiran Chilakamarri, Nathaniel Dean, and Michael Littman. Three-dimensional Tutte embedding. *Congressus Numerantium*, 107:129–140, 1995.
- [49] Richard J. Gerrig and Michael L. Littman. Disambiguation by community membership. *Memory and Cognition*, 18(4):331–338, 1990.

### **Selected Refereed Conference Publications**

- [1] David Abel, Dilip Arumugam, Kavosh Asadi, Yuu Jinnai, Michael Littman, and Lawson Wong. State abstraction as compression in apprenticeship learning. In *Proceedings of the Thirty-Third Association for the Advancement of Artificial Intelligence Conference*, 2019.
- [2] Michael Shum, Max Kleiman-Weiner, Michael L. Littman, and Joshua Tenenbaum. Theory of minds: Understanding behavior in groups through inverse planning. In *Proceedings of the Thirty-Third Association for the Advancement of Artificial Intelligence Conference*, 2019.
- [3] Kavosh Asadi, Dipendra Misra, and Michael L. Littman. Lipschitz continuity in model-based reinforcement learning. In *Proceedings of the 35th International Conference on Machine Learning*, 2018.
- [4] David Abel, Dilip Arumugam, Lucas Lehnert, and Michael L. Littman. State abstractions for lifelong reinforcement learning. In *Proceedings of the 35th International Conference on Machine Learning*, 2018.
- [5] David Abel, Yuu Jinnai, Yue Guo, George Konidaris, and Michael L. Littman. Policy and value transfer in lifelong reinforcement learning. In *Proceedings of the 35th International Conference on Machine Learning*, 2018.
- [6] Mark K Ho, Michael L. Littman, Fiery Cushman, and Joseph L. Austerweil. Effectively learning from pedagogical demonstrations. In C. Kalish, M. Rau, T. Rogers, and J. Zhu, editors, *Proceedings of the 40th Annual Meeting of the Cognitive Science Society*, 2018.
- [7] Nakul Gopalan, Marie desJardins, Michael L. Littman, James MacGlashan, Shawn Squire, Stefanie Tellex, John Winder, and Lawson L. S. Wong. Planning with abstract markov decision processes. In *ICAPS*, 2017.



- [8] James MacGlashan, Mark K Ho, Robert Loftin, Bei Peng, Guan Wang, David L. Roberts, Matthew E. Taylor, and Michael L. Littman. Interactive learning from policy-dependent human feedback. In *Proceedings of the Thirty-Fourth International Conference on Machine Learning*, 2017.
- [9] Kavosh Asadi and Michael L. Littman. An alternative softmax operator for reinforcement learning. In *ICML*, pages 243–252, 2017.
- [10] David Abel, David Hershkowitz, and Michael Littman. Near optimal behavior via approximate state abstraction. In *Proceedings of The 33rd International Conference on Machine Learning*, pages 2915–2923, 2016.
- [11] Bei Peng, James MacGlashan, Robert Loftin, Michael L. Littman, David L. Roberts, and Matthew E. Taylor. A need for speed: Adapting agent action speed to improve task learning from non-expert humans. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, pages 957–965, 2016.
- [12] Blase Ur, Melwyn Pak Yong Ho, Stephen Brawner, Jiyun Lee, Sarah Mennicken, Noah Picard, Diane Schulze, and Michael L. Littman. Trigger-action programming in the wild: An analysis of 200,000 IFTTT recipes. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI)*, 2016.
- [13] J. MacGlashan, M. Babes-Vroman, M. desJardins, M. Littman, S. Muresan, S. Squire, S. Tellex, D. Arumugam, and L. Yang. Grounding English commands to reward functions. In *Proceedings of Robotics: Science and Systems*, 2015.
- [14] James MacGlashan and Michael L. Littman. Between imitation and intention learning. In *Proceedings of IJCAI*, 2015.
- [15] Mark K Ho, Michael L. Littman, Fiery Cushman, and Joseph L. Austerweil. Teaching with rewards and punishments: Reinforcement or communication? In *Proceedings of the 37th Annual Meeting of the Cognitive Science Society*, 2015.
- [16] Blase Ur, Elyse McManus, Melwyn Pak Yong Ho, and Michael L. Littman. Practical trigger-action programming in the smart home. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI)*, 2014.
- [17] Robert Loftin, James MacGlashan, Michael L. Littman, Matthew E. Taylor, David L. Roberts, and Jeff Huang. A strategy-aware technique for learning behaviors from discrete human feedback. In *Proceedings of the Twenty-Eighth Association for the Advancement of Artificial Intelligence Conference*, 2014.
- [18] Robert Loftin, Bei Peng, James MacGlashan, Michael L. Littman, Matthew E. Taylor, Jeff Huang, and David L. Roberts. Learning something from nothing: Leveraging implicit human feedback strategies. In *Proceedings of the Twenty-Third IEEE International Symposium on Robot and Human Communication (ROMAN 14)*, 2014.

- [19] Eric Sodomka, Elizabeth Hilliard, Michael Littman, and Amy Greenwald. Coco-Q: Learning in stochastic games with side payments. *JMLR Workshop and Conference Proceedings: Proceedings of The 30th International Conference on Machine Learning*, 28(3):1471–1479, 2013.
- [20] Sergiu Goschin, Ari Weinstein, and Michael Littman. The cross-entropy method optimizes for quantiles. *JMLR Workshop and Conference Proceedings: Proceedings of The 30th International Conference on Machine Learning*, 28(3):1193–1201, 2013.
- [21] Ari Weinstein and Michael L. Littman. Open-loop planning in large-scale stochastic domains. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence*, 2013.
- [22] Thomas J. Walsh, Michael L. Littman, and Alexander Borgida. Learning web-service task descriptions from traces. *Web Intelligence and Agent Systems*, 10(4):397–421, 2012.
- [23] Sergiu Goschin, Ari Weinstein, Michael L. Littman, and Erick Chastain. Planning in reward-rich domains via PAC bandits. *JMLR: Workshop and Conference Proceedings (10th European Workshop on Reinforcement Learning)*, 24:25–42, 2012.
- [24] Ari Weinstein, Michael L. Littman, and Sergiu Goschin. Rollout-based game-tree search outprunes traditional alpha-beta. *JMLR: Workshop and Conference Proceedings (10th European Workshop on Reinforcement Learning)*, 24:155–166, 2012.
- [25] Zongzhang Zhang, XiaoPing Chen, and Michael Littman. Covering number as a complexity measure for POMDP planning and learning. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*, pages 1853–1859, 2012.
- [26] A. Weinstein and M. Littman. Bandit-based planning and learning in continuous-action Markov decision processes. In *ICAPS 2012*, 2012.
- [27] Michael Wunder, John Robert Yaros, Michael Kaisers, and Michael Littman. A framework for modeling population strategies by depth of reasoning. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, 2012.
- [28] Monica Babes, Vukosi N. Marivate, Michael L. Littman, and Kaushik Subramanian. Apprenticeship learning about multiple intentions. In *Proceedings of the International Conference on Machine Learning*, pages 897–904, 2011.
- [29] Sergiu Goschin, Michael L. Littman, and David H. Ackley. The effects of selection on noisy fitness optimization. In *13th Annual Genetic and Evolutionary Computation Conference, GECCO 2011*, pages 2059–2066, 2011.

- [30] Michael Wunder, Michael Kaisers, Michael Littman, and John Robert Yaros. Using iterated reasoning to predict opponent strategies. In *The Tenth International Conference on Autonomous Agents and Multiagent Systems*, 2011.
- [31] Ari Weinstein, Chris Mansley, and Michael Littman. Sample-based planning for continuous action Markov decision processes. In *21st International Conference on Automated Planning and Scheduling*, 2011.
- [32] Jordan Ash, Monica Babes, Gal Cohen, Sameen Jalal, Sam Lichtenberg, Michael L. Littman, Vukosi N. Marivate, Phillip Quiza, Blase Ur, and Emily Zhang. Scratchable devices: User-friendly programming for household appliances. In *Human-Computer Interaction. Towards Mobile and Intelligent Interaction Environments—14th International Conference*, pages 137–146, 2011.
- [33] John Asmuth and Michael L. Littman. Learning is planning: Near Bayes-optimal reinforcement learning via Monte-Carlo tree search. In *UAI’11*, 2011.
- [34] Thomas J. Walsh, Kaushik Subramanian, Michael L. Littman, and Carlos Diuk. Generalizing apprenticeship learning across hypothesis classes. In *Proceedings of the Twenty-Seventh International Conference on Machine Learning*, pages 1119–1126, 2010.
- [35] Thomas J. Walsh, Sergiu Goschin, and Michael L. Littman. Integrating sample-based planning and model-based reinforcement learning. In *Proceedings of AAAI-10*, 2010.
- [36] Michael Wunder, Michael Littman, and Monica Babes. Classes of multiagent Q-learning dynamics with epsilon-greedy exploration. In *Proceedings of the Twenty-Seventh International Conference on Machine Learning (ICML-10)*, pages 1167–1174, 2010.
- [37] Marie desJardins and Michael L. Littman. Broadening student enthusiasm for computer science with a Great Insights course. In *SIGCSE*, pages 157–161, 2010.
- [38] John Asmuth, Lihong Li, Michael L. Littman, Ali Nouri, and David Wingate. A Bayesian sampling approach to exploration in reinforcement learning. In *25th Conference on Uncertainty in Artificial Intelligence (UAI’09)*, 2009.
- [39] Thomas Walsh, István Szita, Carlos Diuk, and Michael Littman. Exploring compact reinforcement-learning representations with linear regression. In *25th Conference on Uncertainty in Artificial Intelligence (UAI’09)*, 2009.
- [40] Lihong Li, Michael L. Littman, and Christopher R. Mansley. Online exploration in least-squares policy iteration. In *AAMAS*, volume 2, pages 733–739, 2009.
- [41] Ali Nouri and Michael L. Littman. Multi-resolution exploration in continuous spaces. In *Advances in Neural Information Processing Systems 21*, pages 1209–1216, 2008.

- [42] John Asmuth, Michael L. Littman, and Robert Zinkov. Potential-based shaping in model-based reinforcement learning. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence*, pages 604–609, 2008.
- [43] Thomas J. Walsh and Michael L. Littman. Efficient learning of action schemas and web-service descriptions. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence*, pages 714–719, 2008.
- [44] Enrique Munoz de Cote and Michael L. Littman. A polynomial-time Nash equilibrium algorithm for repeated stochastic games. In *24th Conference on Uncertainty in Artificial Intelligence (UAI'08)*, 2008.
- [45] Ronald Parr, Lihong Li, Gavin Taylor, Christopher Painter-Wakefield, and Michael L. Littman. An analysis of linear models, linear value-function approximation, and feature selection for reinforcement learning. In *International Conference on Machine Learning (ICML-2008)*, 2008. See <http://www.cs.duke.edu/~parr/icml08-addendum.html>.
- [46] Carlos Diuk, Andre Cohen, and Michael Littman. An object-oriented representation for efficient reinforcement learning. In *Proceedings of the Twenty-Fifth International Conference on Machine Learning (ICML-08)*, 2008.
- [47] Emma Brunskill, Bethany R. Leffler, Lihong Li, Michael L. Littman, and Nicholas Roy. CORL: A continuous-state off-set-dynamics reinforcement learner. In *24th Conference on Uncertainty in Artificial Intelligence*, pages 53–61, 2008.
- [48] Lihong Li and Michael Littman. Efficient value-function approximation via online linear regression. In *International Symposium on Artificial Intelligence and Mathematics*, 2008.
- [49] Alexander L. Strehl and Michael L. Littman. Online linear regression and its application to model-based reinforcement learning. In *Advances in Neural Information Processing Systems 20*, 2008.
- [50] Monica Babes, Enrique Munoz de Cote, and Michael L. Littman. Social reward shaping in the prisoner’s dilemma. In *7th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 1389–1392, 2008.
- [51] Alexander L. Strehl, Carlos Diuk, and Michael L. Littman. Efficient structure learning in factored-state MDPs. In *Proceedings of the Twenty-Second National Conference on Artificial Intelligence (AAAI-07)*, 2007.
- [52] Thomas J. Walsh and Michael L. Littman. A multiple representation approach to learning dynamical systems. In *Computational Approaches to Representation Change During Learning and Development: AAAI Fall Symposium*, 2007.

- [53] Thomas J. Walsh and Michael L. Littman. Planning with conceptual models mined from user behavior. In *Proceedings of the AAAI-07 Workshop on Acquiring Planning Knowledge via Demonstration*, 2007.
- [54] Fancong Zeng and Michael L. Littman. Just-in-time failure detection. In *ICAC-2007 Workshop on Adaptive Methods in Autonomic Computing Systems (AMACS)*, 2007.
- [55] Bethany R. Leffler, Michael L. Littman, and Timothy Edmunds. Efficient reinforcement learning with relocatable action models. In *Proceedings of the Twenty-Second Conference on Artificial Intelligence (AAAI-07)*, 2007.
- [56] Ronald Parr, Christopher Painter-Wakefield, Lihong Li, and Michael Littman. Analyzing feature generation for value-function approximation. In *International Conference on Machine Learning (ICML-2007)*, 2007.
- [57] Alexander L. Strehl, Lihong Li, and Michael L. Littman. Incremental model-based learners with formal learning-time guarantees. In *Proceedings of the 22nd Conference on Uncertainty in Artificial Intelligence (UAI 2006)*, 2006.
- [58] Michael L. Littman, Nishkam Ravi, Arjun Talwar, and Martin Zinkevich. An efficient optimal-equilibrium algorithm for two-player game trees. In *Twenty-Second Conference on Uncertainty in Artificial Intelligence (UAI-06)*, 2006.
- [59] Martin Zinkevich, Amy R. Greenwald, and Michael L. Littman. Cyclic equilibria in Markov games. In *Advances in Neural Information Processing Systems 18*, 2005.
- [60] R Chadha, A Poylisher, B Deb, M Littman, and B Sabata. Adaptive dynamic server placement in MANETs. In *IEEE Military Communications Conference (MILCOM 2005)*, pages 1083–1089, 2005.
- [61] David L. Roberts, Mark J. Nelson, Charles Lee Isbell, Jr., Michael Mateas, and Michael L. Littman. Targeting specific distributions of trajectories in MDPs. In *Proceedings of The Twenty-First National Conference on Artificial Intelligence*, 2006.
- [62] Alexander L. Strehl, Lihong Li, Eric Wiewiora, John Langford, and Michael L. Littman. PAC model-free reinforcement learning. In *Proceedings of the Twenty-third International Conference on Machine Learning (ICML-06)*, 2006.
- [63] Alexander L. Strehl, Chris Mesterharm, Michael L. Littman, and Haym Hirsh. Experience-efficient learning in associative bandit problems. In *Proceedings of the Twenty-third International Conference on Machine Learning (ICML-06)*, 2006.
- [64] Lihong Li, Thomas J. Walsh, and Michael L. Littman. Towards a unified theory of state abstraction for MDPs. In *Ninth International Symposium on Artificial Intelligence and Mathematics*, 2006.

- [65] Carlos Diuk, Michael Littman, and Alexander Strehl. A hierarchical approach to efficient reinforcement learning in deterministic domains. In *Fifth International Conference on Autonomous Agents and Multiagent Systems (AAMAS-06)*, 2006.
- [66] Lihong Li and Michael L. Littman. Lazy approximation for solving continuous finite-horizon MDPs. In *The Twentieth National Conference on Artificial Intelligence*, pages 1175–1180, 2005.
- [67] Nishkam Ravi, Nikhil Dandekar, Preetham Mysore, and Michael L. Littman. Activity recognition from accelerometer data. In *Seventeenth Innovative Applications of Artificial Intelligence Conference*, pages 1541–1546, 2005.
- [68] Bethany R. Leffler, Michael L. Littman, Alexander L. Strehl, and Thomas J. Walsh. Efficient exploration with latent structure. In *Proceedings of Robotics: Science and Systems*, 2005.
- [69] Alexander L. Strehl and Michael L. Littman. An empirical evaluation of interval estimation for Markov decision processes. In *The 16th IEEE International Conference on Tools with Artificial Intelligence (ICTAI-2004)*, pages 128–135, 2004.
- [70] Michael L. Littman, Nishkam Ravi, Eitan Fenson, and Rich Howard. An instance-based state representation for network repair. In *Proceedings of the Nineteenth National Conference on Artificial Intelligence (AAAI)*, pages 287–292, 2004.
- [71] Satinder Singh, Michael L. Littman, Nicholas K. Jong, David Pardoe, and Peter Stone. Learning predictive state representations. In *The Twentieth International Conference on Machine Learning (ICML-2003)*, pages 712–719, 2003.
- [72] Peter D. Turney, Michael L. Littman, Jeffrey Bigham, and Victor Shnayder. Combining independent modules in lexical multiple-choice problems. In *Recent Advances in Natural Language Processing III: Selected Papers from RANLP 2003*, volume 260, pages 101–110. John Benjamins, 2004.
- [73] 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*, volume 2531 of *Lecture Notes in Artificial Intelligence*, pages 289–306. Springer Verlag, 2002.
- [74] Michail G. Lagoudakis, Ronald E. Parr, and Michael L. Littman. Least-squares methods in reinforcement learning for control. In *Proceedings of the 2nd Hellenic Conference on Artificial Intelligence (SETN-02), Lecture Notes on Artificial Intelligence, Vol. 2308*, pages 249–260. Springer, 2002.
- [75] Michael L. Littman, Richard S. Sutton, and Satinder Singh. Predictive representations of state. In *Advances in Neural Information Processing Systems 14*, pages 1555–1561, 2002.

- [76] Michael L. Littman, Michael Kearns, and Satinder Singh. An efficient, exact algorithm for solving tree-structured graphical games. In *Advances in Neural Information Processing Systems 14*, pages 817–823, 2002.
- [77] Sanjoy Dasgupta, Michael Littman, and David McAllester. PAC generalization bounds for co-training. In *Advances in Neural Information Processing Systems 14*, pages 375–382, 2002.
- [78] Michael Kearns, Michael L. Littman, and Satinder Singh. Graphical models for game theory. In *Proceedings of the 17th Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 253–260, 2001.
- [79] Michael L. Littman. Friend-or-foe Q-learning in general-sum games. In *Proceedings of the Eighteenth International Conference on Machine Learning*, pages 322–328. Morgan Kaufmann, 2001.
- [80] Michael L. Littman and Peter Stone. Implicit negotiation in repeated games. In *Eighth International Workshop on Agent Theories, Architectures, and Languages (ATAL-2001)*, pages 393–404, 2001.
- [81] Justin A. Boyan and Michael L. Littman. Exact solutions to time-dependent MDPs. In Todd K. Leen, Thomas G. Dietterich, and Volker Tresp, editors, *Advances in Neural Information Processing Systems 13*, pages 1026–1032, 2001.
- [82] Michail G. Lagoudakis and Michael L. Littman. Learning to select branching rules in the DPLL procedure for satisfiability. In Henry Kautz and Bart Selman, editors, *Electronic Notes in Discrete Mathematics (ENDM)*, volume 9. Elsevier Science Publishers, 2001.
- [83] Jiefu Shi and Michael L. Littman. Abstraction methods for game theoretic poker. In *Computers and Games*, pages 333–345. Lecture Notes in Computer Science, Springer-Verlag, 2001.
- [84] Michael L. Littman. Computer language games. In *Computers and Games*, pages 396–404. Lecture Notes in Computer Science, Springer-Verlag, 2001.
- [85] Michail G. Lagoudakis and Michael L. Littman. Algorithm selection using reinforcement learning. In Pat Langley, editor, *Proceedings of the Seventeenth International Conference on Machine Learning*, pages 511–518, San Francisco, 2000. Morgan Kaufmann.
- [86] Fan Jiang and Michael L. Littman. Approximate dimension equalization in vector-based information retrieval. In Pat Langley, editor, *Proceedings of the Seventeenth International Conference on Machine Learning*, pages 423–430, San Francisco, 2000. Morgan Kaufmann.

- [87] Michael L. Littman. Initial experiments in stochastic satisfiability. In *Proceedings of the Sixteenth National Conference on Artificial Intelligence*, pages 667–672, 1999.
- [88] Michael L. Littman, Fan Jiang, and Greg A. Keim. Learning a language-independent representation for terms from a partially aligned corpus. In Jude Shavlik, editor, *Proceedings of the Fifteenth International Conference on Machine Learning*, pages 314–322. Morgan Kaufmann, 1998.
- [89] Stephen M. Majercik and Michael L. Littman. MAXPLAN: A new approach to probabilistic planning. In Reid Simmons, Manuela Veloso, and Stephen Smith, editors, *Proceedings of the Fourth International Conference on Artificial Intelligence Planning*, pages 86–93. AAAI Press, 1998.
- [90] Stephen M. Majercik and Michael L. Littman. Using caching to solve larger probabilistic planning problems. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence*, pages 954–959. The AAAI Press/The MIT Press, 1998.
- [91] Anthony Cassandra, Michael L. Littman, and Nevin L. Zhang. Incremental Pruning: A simple, fast, exact method for partially observable Markov decision processes. In *Proceedings of the Thirteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-97)*, pages 54–61, San Francisco, CA, 1997. Morgan Kaufmann Publishers.
- [92] Michael L. Littman. Probabilistic propositional planning: Representations and complexity. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence*, pages 748–754. AAAI Press/The MIT Press, 1997.
- [93] Michael L. Littman and Csaba Szepesvári. A generalized reinforcement-learning model: Convergence and applications. In Lorenza Saatta, editor, *Proceedings of the Thirteenth International Conference on Machine Learning*, pages 310–318, 1996.
- [94] Michael L. Littman, Anthony R. Cassandra, and Leslie Pack Kaelbling. Learning policies for partially observable environments: Scaling up. In Armand Prieditis and Stuart Russell, editors, *Proceedings of the Twelfth International Conference on Machine Learning*, pages 362–370. Morgan Kaufmann, San Francisco, CA, 1995. Reprinted in *Readings in Agents*, M. H. Huhns and M. P. Singh, eds., Morgan Kaufmann, 1998.
- [95] Michael L. Littman, Thomas L. Dean, and Leslie Pack Kaelbling. On the complexity of solving Markov decision problems. In *Proceedings of the Eleventh Annual Conference on Uncertainty in Artificial Intelligence (UAI-95)*, pages 394–402, Montreal, Québec, Canada, 1995.
- [96] David H. Ackley and Michael L. Littman. Altruism in the evolution of communication. In Rodney A. Brooks and Pattie Maes, editors, *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*, pages 40–49, Cambridge, MA, 1994. Bradford Books/MIT Press.



- [97] Michael L. Littman. Markov games as a framework for multi-agent reinforcement learning. In *Proceedings of the Eleventh International Conference on Machine Learning*, pages 157–163, 1994.
- [98] Michael L. Littman. Memoryless policies: Theoretical limitations and practical results. In Dave Cliff, Philip Husbands, Jean-Arcady Meyer, and Stewart W. Wilson, editors, *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, pages 238–245, Cambridge, MA, 1994. The MIT Press.
- [99] Justin A. Boyan and Michael L. Littman. Packet routing in dynamically changing networks: A reinforcement learning approach. In J. D. Cowan, G. Tesauro, and J. Al-spector, editors, *Advances in Neural Information Processing Systems 6*, pages 671–678. Morgan Kaufmann, San Mateo, CA, 1994.
- [100] Robert Allen, Pascal Obry, and Michael Littman. An interface for navigating clustered document sets returned by queries. In *Conference on Organizational Computing Systems (COOCS)*, pages 166–171. SIGOIS, Milpitas, November 1993.
- [101] Michael L. Littman and Justin A. Boyan. A distributed reinforcement learning scheme for network routing. In Joshua Al-spector, Rodney Goodman, and Timothy X. Brown, editors, *Proceedings of the 1993 International Workshop on Applications of Neural Networks to Telecommunications*, pages 45–51. Lawrence Erlbaum Associates, Hillsdale NJ, 1993.
- [102] David H. Ackley and Michael L. Littman. A case for distributed Lamarckian evolution. In C. Langton, C. Taylor, J. D. Farmer, and S. Ramussen, editors, *Artificial Life III: Santa Fe Institute Studies in the Sciences of Complexity*, pages 487–509. Addison-Wesley, Redwood City, CA, 1993.
- [103] Michael L. Littman, Deborah F. Swayne, Nathaniel Dean, and Andreas Buja. Visualizing the embedding of objects in Euclidean space. In H. Joseph Newton, editor, *Computing Science and Statistics, Proceedings of the 24th symposium on the Interface*, volume 24, pages 208–217. Interface Foundation of North America, 1992.
- [104] Laurence Brothers, James Hollan, W. Scott Stornetta, Jakob Nielsen, Steven Abney, George W. Furnas, and Michael L. Littman. Supporting informal communication via ephemeral interest groups. In *Proceedings of the Computer Supported Cooperative Work (CSCW) '92 conference*. The Association For Computing Machinery, Toronto, November 1992.
- [105] Michael L. Littman. An optimization-based categorization of reinforcement learning environments. In I. H. Meyer, H. Roithlat, and S. Wilson, editors, *From Animals to Animats: Proceedings of the Second International Conference on Simulation and Adaptive Behavior*. The MIT Press, 1992.

- [106] Dennis E. Egan, Michael E. Lesk, R. Daniel Ketchum, Carol C. Lochbaum, Joel R. Remde, Michael L. Littman, and Thomas K. Landauer. Hypertext for the electronic library? CORE sample results. In *Proceedings of Hypertext '91*. Association of Computing Machinery, 1991.
- [107] Michael L. Littman and David H. Ackley. Adaptation in constant utility non-stationary environments. In Rik K. Belew and Lashon Booker, editors, *Proceedings of the Fourth International Conference on Genetic Algorithms*, pages 136–142, San Mateo, CA, 1991. Morgan Kaufmann.
- [108] David H. Ackley and Michael L. Littman. Interactions between learning and evolution. In C. Langton, C. Taylor, J. D. Farmer, and S. Ramussen, editors, *Artificial Life II: Santa Fe Institute Studies in the Sciences of Complexity*, volume 10, pages 487–509. Addison-Wesley, Redwood City, CA, 1991.
- [109] Thomas K. Landauer and Michael L. Littman. Fully automatic cross-language document retrieval using latent semantic indexing. In *Proceedings of the Sixth Annual Conference of the UW Centre for the New Oxford English Dictionary and Text Research*, pages 31–38, Waterloo Ontario, October 1990.
- [110] David H. Ackley and Michael L. Littman. Generalization and scaling in reinforcement learning. In D. S. Touretzky, editor, *Advances in Neural Information Processing Systems 2*, pages 550–557, San Mateo, CA, 1990. Morgan Kaufmann.
- [111] David H. Ackley and Michael S. Littman. Learning from natural selection in an artificial environment. In *Proceedings of the International Joint Conference on Neural Networks*, volume 1, pages 189–193. Lawrence Erlbaum Associates, Washington DC, January 1990.

## Other Selected Publications

- [1] Marie desJardins and Michael L. Littman. Evolutionary Huffman encoding. *Inroads*, 9(2):80, 2018.
- [2] Lucas Lehnert and Michael L. Littman. Transfer with model features in reinforcement learning. arXiv preprint arXiv:1807.01736, 2018.
- [3] Kavosh Asadi, Dipendra Misra, and Michael L. Littman. Lipschitz continuity in model-based reinforcement learning. arXiv preprint arXiv:1804.07193, 2018.
- [4] Kavosh Asadi, Evan Cater, Dipendra Misra, and Michael L. Littman. Equivalence between Wasserstein and value-aware model-based reinforcement learning. arXiv preprint arXiv:1806.01265, 2018.
- [5] Yuhang Song, Christopher Grimm, Xianming Wang, and Michael L. Littman. Learning approximate stochastic transition models. arXiv preprint arXiv:1710.09718, 2017.

- [6] Lucas Lehnert, Stefanie Tellex, and Michael L. Littman. Advantages and limitations of using successor features for transfer in reinforcement learning. arXiv preprint arXiv:1708.00102, 2017.
- [7] Cameron Allen, Kavosh Asadi, Melrose Roderick, Abdel rahman Mohamed, George Konidaris, and Michael Littman. Mean actor critic. arXiv preprint arXiv:1709.00503, 2017.
- [8] Christopher Grimm, Dilip Arumugam, Siddharth Karamcheti, David Abel, Lawson LS Wong, and Michael L Littman. Modeling latent attention within neural networks. arXiv preprint arXiv:1706.00536, 2017.
- [9] J. L. Austerweil, S. Brawner, A. Greenwald, E. Hilliard, M. Ho, M. L. Littman, J. MacGlashan, and C. Trimbach. The impact of outcome preferences in a collection of non-zero-sum grid games. AAAI Spring Symposium 2016 on Challenges and Opportunities in Multiagent Learning for the Real World, 2016.
- [10] Michael Littman. Puzzle: Hey, hi! *AI Matters*, 2(2), 2015.
- [11] Pushkar Kolhe, Michael L. Littman, and Charles L. Isbell. Peer reviewing short answers using comparative judgement. In *Proceedings of the Third ACM Conference on Learning @ Scale (L@S '16)*, 2016.
- [12] Guan Wang and Michael L. Littman. Making the intelligent home smart through touch-control trigger-action programming. In *Applications for Future Internet*, 2016.
- [13] Eric Eaton, Tom Dietterich, Maria Gini, Barbara J. Grosz, Charles L. Isbell, Subbarao Kambhampati, Michael Littman, Francesca Rossi, Stuart Russell, Peter Stone, Toby Walsh, and Michael Wooldridge. Who speaks for ai? *AI Matters*, 2(2), 2015.
- [14] Michael Littman. Puzzle: Notable names. *AI Matters*, 2(1), 2015.
- [15] M.L. Littman. Markov decision processes. In James D. Wright, editor, *International Encyclopedia of the Social & Behavioral Sciences*, volume 14, pages 573–575. Oxford: Elsevier, 2nd edition, 2015.
- [16] Michael Littman. ‘rise of the machines’ is not a likely future. *Live Science*, January 2015.
- [17] Michael L. Littman, Charles Isbell, and Aaron Gross. Overfitting: Machine learning music video. In *AAAI Video Competition*, 2014.
- [18] Michael Littman and Samuel Kortchmar. The path to a programmable world. *Footnote1.com*, 2014.
- [19] James MacGlashan, M. Littman, R. Loftin, B. Peng, D. Roberts, and M. E. Taylor. Training an agent to ground commands with reward and punishment. In *AAAI-14 workshop on Machine Learning for Interactive Systems*, 2014.

- [20] Marie desJardins and Michael L. Littman. AAI-13 preface. In *AAAI 2013*, 2013.
- [21] Michael L. Littman. Heuristic functions: A new musical lecture. In *AAAI Video Competition*, 2013.
- [22] James MacGlashan, Monica Babes-Vroman, Kevin Winner, Ruoyuan Gao, Richard Adjogah, Marie desJardins, Michael Littman, and Smaranda Muresan. Learning to interpret natural language instructions. In *Proceedings of AAI-2012 Workshop on Grounding Language for Physical Systems, AAI Technical Report WS-12-07*, 2012.
- [23] Vukosi N. Marivate, Jessica Chemali, Michael Littman, and Emma Brunskill. Discovering multi-modal characteristics in observational clinical data. In *Machine Learning for Clinical Data Analysis and Healthcare NIPS Workshop*, 2013.
- [24] Vukosi Marivate and Michael Littman. An ensemble of linearly combined reinforcement-learning agents. In *AAAI Late Breaking Paper Track*, 2013.
- [25] Michael L. Littman. Inducing partially observable Markov decision processes. *JMLR: Workshop and Conference Proceedings (The 11th ICGI)*, 21:145–148, 2012.
- [26] Matthew Adams, Robert Loftin, Matthew E. Taylor, Michael Littman, and David Roberts. An empirical analysis of RL’s drift from its behaviorism roots. In *ALA 2012 - Workshop at AAMAS*, 2012.
- [27] Michael Wunder, Michael L. Littman, Michael Kaisers, and John Robert Yaros. A cognitive hierarchy model applied to the lemonade game (AAAI workshop). In *Interactive Decision Theory and Game Theory*, 2010.
- [28] Jordan Ash, Monica Babes, Gal Cohen, Sameen Jalal, Michael Littman, Luis Piloto, Phillip Quiza, and Blase Ur. Scratchable devices: We’d like to teach the world to code. Presentation at Scratch@MIT, 2010.
- [29] Ali Nouri, Michael L. Littman, Lihong Li, Ronald Parr, Christopher Painter-Wakefield, and Gavin Taylor. A novel benchmark methodology and data repository for real-life reinforcement learning. In *Multidisciplinary Symposium on Reinforcement Learning, Poster Abstracts*, 2009.
- [30] Michael Wunder, Michael Littman, and Matthew Stone. Communication, credibility and negotiation using a cognitive hierarchy model. In *Multi-agent Sequential Decision-Making in Uncertain Domains (AAMAS Workshop)*, 2009.
- [31] Andrea Pohorecky Danyluk, Léon Bottou, and Michael L. Littman. *Proceedings of the 26th Annual International Conference on Machine Learning (ICML)*. ACM, 2009.
- [32] Michael Littman and Kaushik Subramanian. Real live robot learning. In *AAAI Video Competition*, 2009.

- [33] Carlos Diuk and Michael L. Littman. Hierarchical reinforcement learning. In Juan Ramón Rabuñal Dopico, Julian Dorado, and Alejandro Pazos, editors, *Encyclopedia of Artificial Intelligence*. IGI Global, 2008.
- [34] M. Littman and M. Zinkevich. The 2006 AAAI computer-poker competition. *International Computer Games Association Journal*, 29, September 2007.
- [35] Michael Littman. Artificial intelligence: An instance of Aibo ingenuity. In *AAAI Video Competition*, 2007.
- [36] Thomas J. Walsh, Lihong Li, and Michael L. Littman. Transferring state abstractions between MDPs. In *Proceeding of the ICML-06 Workshop on Structural Knowledge Transfer for Machine Learning*, 2006.
- [37] Michael L. Littman. Letter to the editor on “Science’s Gatekeepers”. *New York Times (Science Time Online)*, May 2006.
- [38] Ronald Brachman, David Gunning, Rodney Brooks, Laurie Damianos, Randall Fish, Michael Littman, Lauri Riek, and Manuela Veloso. Search engine for the real world: The grand treasure hunt. In L. Damianos, editor, *Selected Grand Challenges in Cognitive Science*, pages 4–20, 2005.
- [39] Nicholas L. Cassimatis, Sean Luke, Simon D. Levy, Ross Gayler, Pentti Kanerva, Chris Eliasmith, Timothy W. Bickmore, Alan C. Schultz, Randall Davis, James A. Landay, Robert Miller, Eric Saund, Thomas F. Stahovich, Michael L. Littman, Satinder P. Singh, Shlomo Argamon, and Shlomo Dubnov. Reports on the 2004 AAAI Fall Symposia. *AI Magazine*, 26(1):98–102, 2005.
- [40] Håkan L. S. Younes and Michael L. Littman. PPDDL1.0: An extension to PDDL for expressing planning domains with probabilistic effects. Technical Report CMU-CS-04-167, Carnegie Mellon University, 2004.
- [41] Yukio Ohsawa, Peter McBurney, Simon Parsons, Christopher A. Miller, Alan Schultz, Jean Scholtz, Michael A. Goodrich, Eugene Santos, Jr., Benjamin Bell, Charles Lee Isbell, Jr., and Michael L. Littman. AAAI-2002 Fall Symposium Series. *AI Magazine*, 24(1):95–98, 2003.
- [42] Sofus A. Macskassy, Foster J. Provost, and Michael L. Littman. Confidence bands for ROC curves. bib2html CeDER Working Paper IS-03-04, Stern School of Business, New York University, NY, NY 10012, 2003.
- [43] Michael Littman, Simon Parsons, and Elizabeth Sklar. Metrobots team description. In *Proceedings of RoboCup-2003: Robot Soccer World Cup VII*, pages 1–5, 2003.
- [44] Michael L. Littman and Håkan L. S. Younes. IPC 2004 probabilistic planning track: FAQ 0.1. In *Proceedings of the ICAPS-03 Workshop on the Competition: Impact, Organization, Evaluation, Benchmarks*, pages 7–12, 2003.

- [45] Michael L. Littman, Thu Nguyen, Haym Hirsh, Eitan M. Fenson, and Richard Howard. Cost-sensitive fault remediation for autonomic computing. In *Workshop on AI and Autonomic Computing: Developing a Research Agenda for Self-Managing Computer Systems*, 2003.
- [46] Michael Littman. What’s the scoop? *Games*, 27(3):78, April 2003.
- [47] Michael Littman. Computer eyes. *Archives of Ophthalmology*, 121:419, March 2003.
- [48] Michael Littman. Apply yourself. *Games*, 27(1):78, February 2003.
- [49] Michael L. Littman, Toniann Pitassi, and Russell Impagliazzo. On the complexity of counting satisfying assignments. Working notes of the Satisfiability Testing: Theory and Applications workshop, LICS-2001 (Symposium on Logic In Computer Science), 2001.
- [50] János A. Csirik, Michael L. Littman, Satinder Singh, and Peter Stone. FAucS: An FCC spectrum auction simulator for autonomous bidding agents. In *Second International Workshop on Electronic Commerce (WELCOM-2001)*, pages 193–151, 2001.
- [51] Fan Jiang and Michael L. Littman. Approximate dimension reduction at NTCIR. In *Proceedings of the Second NTCIR Workshop Meeting on Evaluation of Chinese & Japanese Text Retrieval and Text Summarization*, pages 286–290, 2001.
- [52] Harry Q. Bovik, Judy Q. Goldsmith, Andrew Q. Klapper, and Michael Q. Littman. Markov indecision processes: A formal model of decision-making under extreme confusion. Awarded *Toughest Submission to Decide Award* at the Eighteenth International Conference on Machine Learning, 2001.
- [53] Sebastian Thrun and Michael L. Littman. A review of *Reinforcement Learning*. *AI Magazine*, 21(1):103–105, 2000.
- [54] Michael L. Littman. Markov decision processes. In Neil J. Smelser and Paul B. Baltes, editors, *International Encyclopedia of the Social & Behavioral Sciences*. Elsevier, 2001.
- [55] Giuseppe De Giacomo, Marie desJardins, Dolores Cañamero, Glenn S. Wasson, Michael L. Littman, Gerard Allwein, Kim Marriott, Bernd Meyer, Barbara Webb, and Tom Con. The AAAI Fall Symposia. *AI Magazine*, 20(3):87–89, 1999.
- [56] Fan Jiang, Ravindran Kannan, Michael L. Littman, and Santosh Vempala. Improved document sampling for efficient singular value decomposition. bib2html CS-99-05, Duke University, Durham, NC, 1999.
- [57] Michael L. Littman, Greg A. Keim, and Noam M. Shazeer. Solving crosswords with Proverb. In *Proceedings of the Sixteenth National Conference on Artificial Intelligence*, pages 914–915, 1999.

- [58] Greg A. Keim and Michael L. Littman. Design of optimal classification dialogues. Manuscript, 1999.
- [59] Michael L. Littman and Fan Jiang. A comparison of two corpus-based methods for translingual information retrieval. bib2html CS-1998-11, Duke University, Durham, NC, 1998.
- [60] Michael L. Littman, Susan T. Dumais, and Thomas K. Landauer. Automatic cross-language information retrieval using latent semantic indexing. In Gregory Grefenstette, editor, *Cross Language Information Retrieval*. Kluwer, 1998.
- [61] Bob Rehder, Michael L. Littman, Susan Dumais, and Thomas K. Landauer. Automatic 3-language cross-language information retrieval with latent semantic indexing. In *The Sixth Text Retrieval Conference*, pages 103–110. National Institute of Standards and Technology Special Publication, 1997.
- [62] Ming-Yang Kao and Michael L. Littman. Algorithms for informed cows. In Sven Koenig, editor, *AAAI-97 Working Notes, On-line Search*, pages 55–61, 1997.
- [63] Stephen M. Majercik and Michael L. Littman. Reinforcement learning for selfish load balancing in a distributed memory environment. In Paul P. Wang, editor, *Proceedings of the International Conference of Information Sciences*, volume 2, pages 262–265, 1997.
- [64] Michael S. Fulkerson, Michael L. Littman, and Greg A. Keim. Speeding safely: Multi-criteria optimization in probabilistic planning. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence*, page 831. AAAI Press/The MIT Press, 1997.
- [65] Michael L. Littman and Stephen M. Majercik. Large-scale planning under uncertainty: A survey. In *NASA Workshop on Planning and Scheduling in Space*, pages 27:1–8, 1997.
- [66] Michael Lederman Littman. *Algorithms for Sequential Decision Making*. bib2html, Department of Computer Science, Brown University, February 1996. Also Technical Report CS-96-09.
- [67] Michael L. Littman, Anthony R. Cassandra, and Leslie Pack Kaelbling. Efficient dynamic-programming updates in partially observable Markov decision processes. bib2html CS-95-19, Brown University, Providence, RI, 1996.
- [68] Michael L. Littman. Simulations combining evolution and learning. In Rik K. Belew and Melanie Mitchell, editors, *Adaptive Individuals in Evolving Populations: Models and Algorithms: Santa Fe Institute Studies in the Sciences of Complexity*, volume XXVI, pages 465–477. Addison-Wesley Publishing Company, Reading, MA, 1996.
- [69] Michael L. Littman and Fil Menczer. Preface: The effect of memory length on individual fitness in a lizard. In Rik K. Belew and Melanie Mitchell, editors, *Adaptive Individuals in Evolving Populations: Models and Algorithms: Santa Fe Institute Studies*

*in the Sciences of Complexity*, volume XXVI, pages 169–171. Addison-Wesley Publishing Company, Reading, MA, 1996.

- [70] Michael L. Littman and David H. Ackley. Comparative studies of algorithms for switchport allocation. bib2html TM ARH-018713, Bellcore, Morristown NJ, February 1991.
- [71] Michael Bianchi, Michael L. Littman, and Deborah F. Swayne. A review of technology transfer tales: Aphorisms for researchers. In *Bellcore Quality Innovations and Applications Conference*. Bellcore, 1992.
- [72] Michael L. Littman. MaxiBook: User-centered text browsing on an ASCII terminal. In *Proceedings of the 1991 Bellcore/BCC Symposium on User-Centered Design: Making It Happen*, number SR OPT-002130 in Bellcore Special Report, pages 88–91, November 1991.
- [73] Susan T. Dumais and Michael L. Littman. An interface for information retrieval using latent semantic indexing. In *Proceedings of the Bellcore Symposium on User Centered Design: Making Contact*, number SR STS-001658 in Bellcore Special Report, page 45, May 1990.
- [74] Michael S. Littman and Christopher D. Metcalf. An exploration of asynchronous data-parallelism. Technical Report YALEU/DCS/TR-684, Yale University Department of Computer Science, New Haven CT, October 1988.

## Grants and Contracts

- RI: Medium: Collaborative Research: Teaching Computers to Follow Verbal Instructions (NSF, with Muresan@Rutgers/SC&I and desJardins@UMBC/CS). 9/1/2011–8/31/2014, \$703,922.
- RI: Small: Understanding Value-based Multiagent Learning and Its Applications (NSF, solo PI). 7/1/2010–6/30/2013, \$450,000.
- Acquiring and Exploiting Rich Causal Models for Robust Decision Making (AFOSR, subcontract through MIT with Tenenbaum@MIT/CogSci and Kaelbling@MIT/EECS). 9/1/2009–2/28/2011, \$130,000.
- Collaborative Research: Pilot Research on Language-Based Strategies for Creative Problem Solving (NSF, co-PI with Gerrig@StonyBrook/Psych). 7/1/2008–6/30/2010, \$99,683.
- RI: Feature Discovery and Benchmarks for Exportable Reinforcement Learning (NSF, PI with Parr@Duke/CS). 10/1/2007–9/30/2011, \$241,000. [REU Supplement.]
- HSD-DRU: The role of communication in the dynamics of effective decision making (NSF, PI with Stone@Rutgers/CS and Sopher@Rutgers/Econ and Lau@Rutgers/PoliSci). 3/1/2007–2/28/2011, \$684,995.



- Transfer Learning (DARPA, subcontractor through Berkeley/MIT), 2006–2009.
- Integrated Learning (DARPA, subcontractor through BBN), 2006–2009.
- Learning for distributed network threat detection (Homeland Security SBIR, subcontract through PnP Networks, Inc.) 2005.
- Decision Theoretic Planning for Planetary Exploration (NASA, PI with Meuleau@NASA). 1/2005.
- Learning to Create Knowledge: Bridging the Representation Gap (DARPA, PI with Singh@Michigan/CS and Barto@UMass/CS). 8/1/2004–8/1/2005, \$455,970.
- ITR: Collaborative Research: Representation and Learning in Computational Game Theory (NSF, co-PI with Schapire@Princeton/CS and Kearns@Penn/CS and Warmuth@UCSB/CS). 9/1/2003–9/1/2008, \$1,600,000.
- Cognitive policy-based network management (DARPA, subcontract through Telcordia with Chadha@Telcordia). 9/1/2003–12/1/2003, \$20,000.
- Evaluating Next Generation Probabilistic Planners (NSF, solo PI). 7/1/2003–7/1/2006, \$244,000.
- Creating the First International Probabilistic Planning Competition (NSF, solo PI). 5/1/2003–5/1/2004, \$10,232.
- Self-healing networks via self-aware software (DARPA, subcontract through PnP Networks, Inc with Howard@PnP). 9/1/2002–8/31/2003, \$175,000.
- Graduate assistance in areas of national need for fellowships in experimental computer science (US Department of Education, co-PI). 8/1/1998–7/31/2001, \$126,110.
- Cross-language text retrieval using latent semantic analysis (SBIR, subcontract through University of Colorado with Landauer@CU/Psych). 9/1/1997–3/1/1999, \$103,212.
- Planning under uncertainty in large domains (NSF CAREER). 3/1/1997–2/28/2001, \$300,000.
- Visually-guided planning for robot control (Duke’s Arts & Sciences Research Council, solo PI). 5/1/1997, \$2,500.
- CURIOUS: (C)enter for (U)ndergraduate Education and (R):esearch: Integration Thr(OU)gh Performance and Vi(S)ualization (NSF, co-PI with Astrachan@Duke/CS). 9/1/1996–8/31/1996, \$405,200.

## Oral Presentations

### • Tutorials:

- Model-based Reinforcement Learning, *Neural Information Processing Systems*, Vancouver, Canada, 2009.
- Reinforcement Learning, *Machine Learning Summer School*, Cambridge University, Cambridge, UK, 2009.
- Probabilistic Planning and Reinforcement Learning, *1st International Seminar on New Issues of Artificial Intelligence*, Universidad Carlos III de Madrid, Madrid, Spain, 2008.
- Reinforcement Learning, Machine Learning Summer School, Canberra Australia, 2006. [Presentation canceled because of extreme weather.]
- Learning and Planning in Markov Environments, *Brazil Agents School*, Porto Alegre, Brazil, 2005.
- Learning and Planning in Markov Environments, *Americas School on Agents and Multiagent Systems*, Pittsburgh, PA, 2005.
- Decision Theoretic Techniques, *Americas School on Agents and Multiagent Systems*, New York, NY, 2004.
- Learning Topics in Game-Theoretic Decision Making, *COLT 2003: The Sixteenth Annual Conference on Learning Theory*, Washington, DC, 2003.
- Multiagent learning from a game theoretic perspective (with Michael Bowling), *International Joint Conference on Artificial Intelligence*, Acapulco, Mexico, 2003. (Also part of the *Second Americas School on Agents and Multiagent Systems*.)
- Introduction to decision-theory for agents, *First Americas School on Agents and Multiagent Systems*, Los Angeles, CA, 2002.
- Markov decision processes and planning under uncertainty (with Leslie Pack Kaelbling), *American Association for Artificial Intelligence*, Orlando, FL, 1999.
- Reinforcement Learning (with Leslie Pack Kaelbling and Andrew Moore), *American Association for Artificial Intelligence*, Seattle, WA, 1994. [Prepared materials, but did not attend due to family emergency.]

### • Invited Seminars:

- Artificial Intelligence and Machine Learning, *Briefing to Naval Operations Strategic Studies Group (SSG)*, Newport, RI, 2016.
- Joint Intention and Cooperation, *Department of Computing Science, University of Alberta*, Edmonton, CA, 2016.
- Hour of Code (part of Computer Science Education Week), *Barrington High School*, Barrington, RI, 2013.

- Efficiently Learning to Behave Efficiently, *Department of Computer Science, Johns Hopkins University*, Baltimore, MD, 2013.
- Bayesian Model Sampling in Reinforcement Learning, *Pattern Theory Seminar, Brown Department of Applied Mathematics*, Providence, RI, 2013.
- Introduction to Machine Learning, *Department of Computer Science, Bard College*, Annandale-on-Hudson, NY, 2012.
- Experiences Teaching Thousands Online, *CSEE Colloquium, University of Maryland, Baltimore County*, Baltimore, MD, 2012.
- Advances in Model-based Reinforcement Learning or Q-learning Considered Harmful, *Rosetta Stone*, Harrisonburg, VA; *Department of Biostatistics, Columbia University*, New York, NY, 2011.
- Bayesian Model Sampling in Reinforcement Learning, *PriML Kickoff Meeting*, Philadelphia, PA, 2010.
- Classes of Multiagent Q-learning Dynamics with epsilon-greedy Exploration, *Santa Fe Institute Workshop on Decentralized Control in Systems of Strategic Actors*, Santa Fe, NM, 2010.
- Bayesian Models in Reinforcement Learning, *Department of Statistics, Columbia University*, New York, NY, 2010.
- Efficiently Learning to Behave Efficiently, *The Gatsby Computational Neuroscience Unit*, London, UK, 2010; *Department of Neuroscience, Yale University*, New Haven, CT, 2010; *GRASP Lab, University of Pennsylvania*, Philadelphia, PA, 2009.
- KWIK Learning: Explicit Uncertainty and Decision Making, *NIPS 2008 Workshop Model Uncertainty and Risk in Reinforcement Learning*, Whistler, BC, 2008.
- Modular Approaches to Computer Programs for Language Games, *ITA Software*, Cambridge, MA, 2008.
- Exploring Approximations to Models of Exploration, *NIPS 2008 Workshop on Machine Learning Meets Human Learning*, Whistler, BC, 2008.
- Introducing Computer Science from Scratch, *Scratch@MIT Conference*, Cambridge, MA, 2008.
- Efficient Model Learning for Reinforcement Learning, *New York University*, New York, NY, 2008; *Georgia Tech*, Atlanta, GA, 2008; *Google Research*, New York, NY, 2008; *Drexel University*, Philadelphia, PA, 2008; *New York Academy of Sciences*, New York, NY, 2007; *Universite Pierre & Marie Curie*, Paris, France, 2007; *Brown University*, Providence, RI, 2007; *MIT*, Cambridge, MA, 2007.
- Robotics and Reinforcement Learning, *NJ Governor's School of Engineering and Technology*, 2007.

- Representation learning in the reinforcement learning setting, *AAAI Fall Symposium on Computational Approaches to Representation Change During Learning and Development*, Arlington, VA, 2007.
- Reinforcement Feedback in Autonomic Computing Applications, *ICAC-2007 Workshop on Adaptive Methods in Autonomic Computing Systems (AMACS)*, Jacksonville, FL, 2007.
- Reward Bonuses for Efficient, Effective Exploration (or, KWIK Learners at Play), *Cognitive Development Society*, Santa Fe, NM, 2007.
- Knowing What You Know, *Air Force Office of Research Support, Robust Decision Making Workshop*, Arlington, VA, 2007.
- Advancing the Theory and Practice of Model-based Reinforcement Learning, *University of Amsterdam*, Amsterdam, The Netherlands, 2006; *University of Washington*, Seattle, WA, 2006.
- Advances in Model-based Reinforcement Learning or “Q-learning considered harmful”, *IBM*, Hawthorne, NY, 2005; *Carnegie Mellon University*, Pittsburgh, PA, 2005; *Naval Research Institute*, Washington, DC, 2005; *University of Toronto*, Toronto, ON, 2005.
- Multiagent Reinforcement Learning: Clues from the brain, *University of Maryland, Baltimore County*, Baltimore, MD, 2005; *Telcordia*, Piscataway, NJ, 2004.
- Towards Real-Life Reinforcement Learning, *IBM*, Yorktown Heights, NY, 2006; *The School of Informatics, The University of Edinburgh*, Edinburgh, Scotland, 2006; *Wesleyan College*, Middletown, CT, 2006; *Moravian College*, Bethlehem, PA, 2006; *Google Research*, New York, NY, 2005; *University of Pennsylvania*, Philadelphia, PA, 2005; *Brigham Young University*, Provo, UT, 2005; *Lockheed-Martin*, Cherry Hill, NJ, 2004; *Vrije Universiteit Brussel*, Brussels, Belgium, 2004; *Carnegie Mellon University*, Pittsburgh, PA, 2004; *Rutgers Computer Science Department Open House*, Piscataway, NJ, 2004; *University of Alberta*, Edmonton Canada, 2004; *Princeton University*, Pervasive Computing class, Princeton, NJ, 2004; *Rutgers Cognitive Science Center*, Piscataway, NJ, 2004.
- An Analysis of Model-based Interval Estimation, *University of Alberta*, Edmonton Canada, 2004.
- Algorithmic Problems in Sequential Decision Making, *DIMACS*, Piscataway, NJ, 2004.
- An Instance-based Learning Approach to Network Repair, *University of Alberta*, Edmonton Canada, 2004; *DIMACS*, Piscataway, NJ, 2004; *UT Austin*, Austin, TX, 2004; *Georgia Tech*, Atlanta, GA, 2004; *Stanford*, Stanford, CA, 2003; *University of Maryland, Baltimore County*, Baltimore, MD, 2003.
- A Polynomial-time Nash Equilibrium Algorithm For Repeated Games, *University of Alberta*, Edmonton Canada, 2004; *Carnegie Mellon University*, Pittsburgh, PA, 2003.

- Learning Predictive State Representations, *MIT*, Cambridge, MA, 2002; *Rutgers*, Piscataway, NJ, 2002; *Brown*, Providence, RI, 2002; *Stanford*, Stanford, CA, 2002; *University of Massachusetts*, Amherst, MA, 2002.
- Language games and other meaningful pursuits, *IBM*, Yorktown Heights, NY, 2002. *Google*, Mountainview, CA, 2002. *Stanford University*, Stanford, CA, 2002. *AT&T Labs*, Florham Park, NJ, 2001.
- Sequential decision making algorithms in filtering, *DIMACS*, Piscataway, NJ, 2003.
- Leading best-response strategies in repeated games, *MIT*, Cambridge, MA, 2001; *Brandeis*, Waltham, MA, 2001.
- Are Crossword Harder than Chess?, *University of Alberta*, Edmonton, Canada, 2004; *New Jersey Institute of Technology*, Newark, NJ, 2003; *Bowdoin College*, Bowdoin, ME, 2003.
- Probabilistic constraint satisfaction and an application, *University of Toronto*, Toronto, Canada, 2001; *University of Illinois*, Urbana, IL, 2001; *Cornell*, Ithaca, NY, 2001.
- Approximate dimension equalization in vector-based information retrieval, *MIT*, Cambridge, MA, 2000; *Telcordia*, Morristown, NJ, 2000.
- Solving crossword puzzles via probabilistic constraint satisfaction, *Institute for Advanced Study*, Princeton, NJ, 2000; *IBM Research*, Hawthorne, NY, 2000; *Telcordia*, Morristown, NJ, 2000; *Brown University*, Providence, RI, 2000; *University of Maryland*, College Park, MD, 2000; *University of Alberta*, Edmonton, CA, 2000; *Microsoft Research*, Redmond, WA, 2000.
- Microrobotics research at Duke, *The Durham FM Association (DFMA)*, Durham, NC, 1999.
- Constraint satisfaction, probability, and solving crossword puzzles, *Rutgers University*, Piscataway, NJ, 1999; *Yale University*, New Haven, CT, 1999; *Swarthmore College*, Swarthmore, PA, 1999; *AT&T Labs*, Florham Park, NJ, 1999; *Columbia University*, New York, NY, 1999; *Princeton University*, Princeton, NJ, 1999.
- Planning under uncertainty: Lessons from complexity theory, *AT&T Labs*, Florham Park, NJ, 1998; *Carnegie Mellon University*, Pittsburgh, PA, 1997.
- Learning a language-independent representation for terms from a partially aligned corpus, *University of Colorado, Boulder*, Boulder, CO, 1998.
- AI robs from the rich and gives to the dumb, *Duke University*, ACM Student Chapter, Durham, NC, 1998.
- An introduction to reinforcement learning, *Duke University, Psychology: Experimental Departmental Seminar*, Durham, NC, 1997; *Triangle Area Neural Network Society*, Research Triangle Park, NC, 1997.

- Computationally efficient sampling, *Statistics Week 1997, Duke University*, Durham, NC, 1997.
- Generalized Markov decision processes, *University of Maryland*, College Park, MD, 1997.
- Decision-making under uncertainty via partially observable Markov decision processes, *Fuqua School of Business, Duke University*, Durham, NC, 1997; *University of Kentucky*, Lexington, KY, 1997; *North Carolina State University*, Raleigh, NC, 1996; *Institute of Statistics & Decision Sciences, Duke University*, Durham, NC, 1996.
- Algorithms for sequential decision making, *University of Pennsylvania*, Philadelphia, PA, 1996; *Duke University*, Durham, NC, 1996; *Carnegie Mellon University*, Pittsburgh, PA, 1996; *University of Pittsburgh*, Pittsburgh, PA, 1996; *University of Delaware*, Newark, DE, 1996.
- Markov models for sequential decision making, *University at Stony Brook*, Stony Brook, NY, 1995; *Brandeis University*, Waltham, MA, 1995; *University of Pennsylvania*, Philadelphia, PA, 1995.
- Partially observable Markov decision processes, *Massachusetts Institute of Technology*, Cambridge, MA, 1994; *Carnegie Mellon University*, Pittsburgh, PA, 1993.
- Solving reinforcement learning environments, *Rowland Institute for Science*, Cambridge, MA, 1994.
- Learning probabilistic policies by reinforcement, *Carnegie Mellon University*, Pittsburgh, PA, 1993.
- State of the art information retrieval, *University of Delaware*, Newark, DE, 1991.
- Combining evolution and learning, *University of San Diego, AI graduate seminar*, La Jolla, CA, 1991.

• **Invited Workshop Presentations:**

- Humanitys Long-Term Future with AI, *Future of AI Symposium*, New York, NY, 2016.
- Reinforcement Learning from users: New algorithms and frameworks, *NIPS: Machine Learning From and For Adaptive User Technologies: From Active Learning & Experimentation to Optimization & Personalization*, Montreal, CA, 2015.
- Coordination and Joint Intention, *NIPS: Learning, Inference and Control of Multi-Agent Systems*, Montreal, CA, 2015.
- Computational Game Theory in Sequential Environments, *First Reinforcement Learning Decision Making Meeting*, Princeton, NJ, 2013.
- Bad Dell!: Teaching Computers With Rewards and Punishments, *Computer Science Open House*, Providence, RI, 2013.

- Reinforcement learning in Scratch, *Artemis Summer Program* (Computer Science introduction for Middle School girls), Providence, RI, 2013.
- The ICML Cycle system, *AAAI-13 Spotlight Track, “What’s Hot”*, Bellevue, WA, 2013.
- Interacting with Reinforcement Learning, *Microsoft Research Faculty Summit*, Redmond, WA, 2013.
- Ethics and Reviewing, *AAAI-13 Panel on Conference Reviewing: Best Practices*, Bellevue, WA, 2013.
- Administration in “If I’d Only Known!” panel *Grace Hopper Celebration of Women In Computing Conference*, Baltimore, MD, 2012.
- Reinforcement Learning and Utility-Based Decisions, *KDD Workshop on Utility-based Datamining*, Philadelphia, PA, 2006.
- Algorithm and complexity issues in discrete multistage games (keynote), *Workshop on Decision Making in Adversarial Domains*, College Park, MD, 2005.
- Evolution and multiagent learning: Clues from the brain, *AAMAS-04 Workshop on Learning and Evolution in Agent Based Systems*, New York, NY, 2004, also *AAAI Fall Symposium*, Washington, DC, 2004.
- Evaluating predictions, *ICML-2004 Predictive Representations of World Knowledge*, Banff, AB, 2004.
- Introduction to Planning Under Uncertainty, *Neural Information Processing Systems Conference: Planning for the Real World*, Whistler, BC, 2003.
- Remembering the past by predicting the future, *Eighteenth International Conference on Machine Learning: Hierarchy and Memory in Reinforcement Learning*, Williams, MA, 2001.
- Stochastic satisfiability: Algorithms and applications, *Sixteenth Annual IEEE Symposium on Logic in Computer Science: Workshop on Theory and Applications of Satisfiability Testing*, Boston, MA, 2001.
- Leading best-response strategies in repeated games, *Eleventh Yale Workshop on Adaptive and Learning Systems*, New Haven, CT, 2001.
- An Exponential Algorithm for  $k$ -MAJSAT, *DIMACS Workshop on Faster Exact Solutions for NP-Hard Problems*, Princeton, NJ, 2000.
- Finding and exploiting relationships in large text collections, *DIMACS Workshop on Mathematical Methods for High Performance Data Mining Applications*, Princeton, NJ, 1998.
- Robotic exploration via partially observable Markov decision processes, *Conference on Automated Learning and Discovery: Robot Exploration and Learning*, Pittsburgh, PA, 1998.

- Planning under uncertainty: Lessons from complexity theory, *DARPA Young AI Researchers Workshop*, Providence, RI, 1997.
- Complexity of planning in stochastic domains, *American Association of Artificial Intelligence: Structural Issues in Planning and Temporal Reasoning Workshop*, Portland, OR, 1996.
- Generalized Markov decision processes, *INFORMS '96*, Washington, DC, 1996; *NSF Workshop on Reinforcement Learning*, Harpers Ferry, WV, 1996.
- Reinforcement learning with hidden state (with Lonnie Chrisman), *International Conference on Machine Learning: Reinforcement Learning Workshop*, Amherst, MA, 1993.
- High-dimensional graph embedding, *DIMACS Workshop on Visualizing Graphs*, New Brunswick, NJ, 1992.
- Large scale browsing interfaces, *Computer Human Interaction Conference: Information Worlds Workshop*, New Orleans, LA, 1992.
- Combining evolution and learning, *International Conference on Genetic Algorithms: Artificial Life Workshop*, San Diego, CA, 1991.
- Perpetual exploration, *Neural Information Processing Systems Conference: Active Learning Workshop*, Denver, CO, 1991.
- Interactions between evolution and learning, *Neural Information Processing Systems Conference: Evolution and Learning Workshop*, Denver, CO, 1990.

## Patent

- Awarded U.S. Patent 5,301,109 for cross-language information retrieval (with Thomas K. Landauer), 1994.

## Professional Activities

- **Professional Committees:** American Association for Artificial Intelligence, Executive Council Member, 2002–2005.
- **Journal Advisory Board Member:** Journal of Artificial Intelligence Research, 2003–.
- **Journal Associate Editor / Action Editor:** Guest Editor: Machine Learning Journal special issue on Empirical Evaluations in Reinforcement Learning (with Shimon Whiteson), 2009–2011; Guest Editor: Machine Learning Journal special issue on Game Theory and Learning (with Amy Greenwald), 2005–2007; Journal of Machine Learning Research, 2003–2010; Journal of Artificial Intelligence Research, 2001–2003; Adaptive Behavior, 1998–2001.



- **Conference Chair and Workshop Organizer:** The Future of AI Symposium at NYU, co-organizer, 2015; DARPA ISAT Workshop on “Centaur’s”, co-organizer, 2015; ICML, General chair, 2013; AAI, Program Co-chair (with Marie desJardins), 2013; ICML, Programme Co-chair (with Leon Bottou), 2009; New York Academy of Sciences Machine Learning Symposium, 2006–2008; ICML, Workshop Co-chair (with Sanjoy Dasgupta), 2008; Epigenetic Robotics, Local Arrangements Chair, 2007; AAI Human–Computer Poker Competition (arbiter), 2007; ICML, Publicity Chair, 2006; AAI Computer Poker Competition (arbiter, with Martin Zinkevich), 2006; NIPS Workshop, Reinforcement Learning: Benchmarks and Bake-offs II, 2004 (with Martin Reidmiller); AAI Fall Symposium, Real-Life Reinforcement Learning, 2004 (with Satinder Singh); NIPS Workshop, Reinforcement Learning: Benchmarks and Bake-offs, 2004 (with Rich Sutton); ICAPS, Competition Chair, 2004; NIPS Workshop, Multi-Agent Learning: Theory and Practice, 2002 (with Gerry Tesauro); AAI, Tutorial Chair, 2002; AAI Fall Symposium, Personalized Agents, 2002 (with Charles Isbell and Peter Stone); AAI, Tutorial Chair, 2000; Uncertainty in Artificial Intelligence, Beyond MDPs: Representations and Algorithms, 2000 (with Michael Kearns and Satinder Singh); AAI Spring Symposium, Planning with Partially Observable Markov Decision Processes, 1998 (with Tony Cassandra); AAI, Student Abstract and Poster Program Chair, 1998.
- **Senior Program Committee Member:** American Association for Artificial Intelligence, 2005–2007; International Conference on Machine Learning, 2004–2007; Uncertainty in Artificial Intelligence, 2004; Electronic Commerce, 2004; International Joint Conference on Artificial Intelligence, 2003; Neural Information Processing Systems (NIPS), 2003; Artificial Life, 1994.
- **Journal Editorial Board Member:** Journal of Machine Learning Research, 2000–2003; Journal of Artificial Intelligence Research, 1997–2001.
- **Officer:** Triangle Area Neural Network Society, President, 1998–2000; Triangle Area Neural Network Society, Membership Chair, 1997.
- **Awards Committee Member:** Awards Committee, IJCAI-JAIR Best Paper Prize, 2005; Nominator for quadrennial Kyoto Prize in Information Science, 2004; Best Paper Awards Committee, 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, 2001.
- **Organizing Committee Member:** AAI 2013 Spring Symposium: Lifelong Machine Learning, 2013; International Conference on Automated Planning and Scheduling Workshop: Monte-Carlo Tree Search: Theory and Applications, 2011; AAI Fall Symposium: Computational Approaches to Representation Change During Learning and Development, 2007; AAI AI Video Competition, 2007; AAI doctoral consortium panelist, 2005–2007; ICML workshop, Predictive Representations of World Knowledge, 2004; ICAPS workshop, Planning under Uncertainty and Incomplete Information,

2003; ICAPS workshop, PDDL, 2003; ICAPS workshop, Workshop on the Competition: Impact, Organization, Evaluation, Benchmarks, 2003; AAAI Spring Symposium, Collaborative Learning Agents, 2002; AAAI Workshop, Probabilistic Approaches in Search, 2001; AAAI Fall Symposium, Using Uncertainty Within Programming, 2001; Johns Hopkins University Center for Language and Speech Processing Summer Workshop, Reading Comprehension, 2000; AIPS-00 Workshop on Decision-Theoretic Planning, 2000.

- **Funding Agency Referee:** Alberta CERC (National Chair), 2012; AFOSR reviewer, 2012; Israeli Science Foundation, 2005; NSERC Proposal Reviewer, 2003; NASA Panel member, 2001; NSF Panel Member, 1997, 1999, 2002–2004, 2007, 2012; NSF SBIR Panel Member, 1999.
- **Book and Journal Referee:** Adaptive Behavior, 1995, 1996; Algorithm Journal, 1999; Artificial Intelligence, 1999, 2002; Automatica, 1999; Computational Intelligence, 2005; Computer Journal, 1998; Decision Support Systems Journal, 2003; Electronic Transactions on Artificial Intelligence, 2002, 2003; Games and Economic Behavior, 1998; IEEE Journal on Selected Areas in Communications, 2001; IEEE Transactions on Systems, Man and Cybernetics, Part B, 2005; Information Retrieval, 2002; International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems, 2002; Journal of Artificial Intelligence Research, 1995, 1997, 1999, 2004; Journal of Discrete Applied Mathematics, 2002; Journal of Intelligent Information Systems, 2001; Journal of Machine Learning Research, 2002; Journal of the ACM, 1998; Journal of the American Society for Information Science and Technology, 2004; Journal of Natural Language Engineering, 2001; Kluwer Press, 1994, 1999; Linear Algebra Applications, 1999; Machine Learning, 1993, 1994, 1996, 1997, 1999–2001; Management Science, 2004; Mathematics of Operations Research, 2005, 2006; Operations Research, 1995, 1997, 1998; Society for Industrial and Applied Mathematics, 1998; Transactions on Information Systems, 1992, 1993.
- **Conference Referee:** AISTats, 2011; American Association for Artificial Intelligence, 1996, 1999, 2000, 2004; American Association for Artificial Intelligence (Student Abstract and Poster Program), 1997, 1999; Artificial Intelligence Planning and Scheduling, 2000; Artificial Life, 1998; Autonomous Agents & Multi Agent Systems, 2002, 2005, 2005 (workshop); Computer Games, 2006; Electronic Commerce, 2005, 2006; European Conference on Machine Learning, 2005 (workshop); International Conference on Machine Learning, 1995, 1997–2001; International Conference on Planning and Scheduling, 2000, 2002; International Joint Conference on Artificial Intelligence, 1995, 1997, 1999, 2001, 2005 (workshop), 2006; International Symposium on Artificial Intelligence and Mathematics, 1999, 2000, 2006; Neural Information Processing Systems, 1991, 1992, 1994, 1996–2001, 2005; RoboCup, 2004; Robotics: Science & Systems, 2005, 2006; Simulations of Adaptive Behavior, 1994, 1996, 2000; Uncertainty in Artificial Intelligence, 2000, 2002, 2007.

### **Brown Departmental and University Service**

- Executive Committee, 2014–2018.
- Curriculum Committee, 2012–2014.

### **Rutgers Departmental and University Service**

- Department Chair, 2009–2012.
- Executive Committee, 2005–2012.
- Undergraduate Education, 2006–2009.
- Peer Evaluation Committee, 2006.
- Publicity Committee, 2004, 2006.
- Distinguished Colloquium Committee, 2004–2005.
- Workload Committee Creation Committee, 2004.
- Tenure Reading Committees, 2004–2006.
- Faculty Meeting Secretary, 2003–2004, 2005–2007.
- Departmental Colloquium Organizer, 2002–2003, 2009–2011.

### **Duke Departmental and University Service**

- Outstanding Project Committee, 1999.
- Department Colloquium Chair, 1998.
- Triangle Lecture Series Coordinator, Department Colloquium Chair, 1998.
- Graduate Compensation Committee, 1998.
- Faculty Search Committee, 2000 (Artificial Intelligence), 1999 (Cognitive Neuroscience), 1998 (Algorithms).
- New Faculty Orientation List Consultant, 1997.

## Teaching Experience

- **Professor:**

- Rutgers, CS 110 (service course): “Introduction to Computers and Applications” (Spring 2011).
- Rutgers, CS 500 (graduate proseminar): “Seminar in Machine Learning” (Spring 2010).
- Rutgers, CS 195 (honors seminar): “Programming for the Masses” (Spring 2010).
- Rutgers, CS 105 (service course): “Great Insights in Computer Science” (Spring 2006, Fall 2006, Spring 2007, Fall 2007, Fall 2008).
- Rutgers, CS 500 (graduate proseminar): “A Machine Learning Approach to Solving Pitfall” (Spring 2009).
- Rutgers, CS 500 (graduate proseminar): “Planning in Learned Environments” (Fall 2007).
- Rutgers, CS 500 (graduate proseminar): “Multiagent Reinforcement Learning” (Spring 2007).
- Rutgers, CS 500 (graduate proseminar): “Reinforcement Learning Theory” (Fall 2006).
- Rutgers, CS 671/672 (graduate): “Learning and Sequential Decision Making” (Spring 2004, Fall 2005, Spring 2009 as CS598).
- Rutgers, CS 500 (graduate proseminar): “Social Reinforcement Learning” (Spring 2005).
- Rutgers, CS 205 (undergraduate): “Introduction to Discrete Structures” (Fall 2004, Spring 2005).
- Rutgers, CS 536 (graduate): “Machine Learning” (Fall 2003, Fall 2004).
- Rutgers, CS 500 (graduate proseminar): “Learned Representations in AI” (Fall 2003).
- Rutgers, CS 500 (graduate proseminar): “Learning Robot Reading Group” (Spring 2003).
- Princeton, COS 302 (undergraduate): “Introduction to Artificial Intelligence” (Fall 2001).
- Duke, CPS 196 (undergraduate): “Programming Under Uncertainty” (Fall 1999).
- Duke, CPS 271 (graduate): “Numeric Artificial Intelligence” (Spring 1997, 1998, 1999).
- Duke, CPS 130 (undergraduate): “Introduction to the Design and Analysis of Algorithms” (Fall 1997, 1998).

- Duke, CPS 370 (mixed seminar): “Using the Web to Solve Crossword Puzzles” (Fall 1998).
- University of North Carolina at Chapel Hill, INLS 379 (graduate seminar): “Research in Information Retrieval” (Spring 1998, co-organized with Newby at UNC).
- Duke, CPS 370 (graduate seminar): “Statistical Natural Language Processing” (Fall 1997).
- Duke, CPS 370 (graduate): “Planning Under Uncertainty” (Fall 1996).

• **Research Advisor:**

- Ph.D.: Sam Saarinen (Brown), Lucas Lehnert (Brown), Jun Ki Lee (Brown), David Abel (Brown), Guan “Royal” Wang (Brown), Kavosh Asadi (Brown), Carl Trimbach (Brown), Stephen Brawner (Brown, 2018), Vukosi Marivate (Rutgers, 2014), Monica Babes (Rutgers, 2014), Sergiu Goschin (Rutgers, 2014, primary advisor), Michael Wunder (Rutgers, 2013, secondary advisor), Ari Weinstein (Rutgers, 2013), John Asmuth (Rutgers, 2013), Ali Nouri (Rutgers, 2010), Tom Walsh (Rutgers, 2010), Carlos Diuk (Rutgers, 2010), Lihong Li (Rutgers, 2009), Bethany Leffler (Rutgers, 2008), Brian Russell (Rutgers, 2008, primary advisor), Alexander L. Strehl (Rutgers, 2007), Fancong Zeng (Rutgers, 2007), Stephen Majercik [NASA Fellow] (Duke, 2000), Fan Jiang (Duke, 2000).
- Past Lab Members: Zongzhang Zhang (visitor, 2010–2011), Ewin Davis (visitor, 2010), Marek Grzes (visitor, 2008), Scott Selikoff (2008), Jose Enrique Munoz de Cote (visitor, 2007), David LeRoux (2005).
- Master’s: Kaushik Subramanian (2010), Rati Sharma (thesis, 2006), Ahmed Chowdhury (2006), Paul Batchis (2005), Jianyu Yang (2004), Peng Song (2004), Robin Carnow (2004), Bing Bai (2003), Zhi Wei (2003), Sushant Agarwal (Duke, 2000), Hao Cheng (Duke, 1999).
- Second-year project: Michail Lagoudakis (2000), Jiefu Shi (2000).
- Undergraduate thesis: Karl Weinmeister (1999), Michael Cafarella (1996).

• **External Committee Member:**

- Ph.D.: Min Wen (Penn), Robert Loftin (NCSU), Evan Russek (NYU, Neuroscience, 2018), Pedro Tsividis (MIT, Brain and Cognitive Sciences, 2018), Kevin Waugh (CMU), Mark Ho (Brown, CLPS, 2017), Nan Jiang (U. of Michigan, 2017), Nicholas Franklin (Brown, CLPS, 2017), Michael Johansson (U. Alberta, 2016), Adrian Boteanu (WPI, 2015), Jonathan Scholz (GA Tech, 2015), Liangyuan Hu (Brown, Biostatistics, 2015), Arthur Guez (Gatsby, 2015), Jeffrey Cockburn (Brown, Cognitive Neuroscience, 2015), James MacGlashan (U. Maryland, Baltimore County, 2013), Michael Kaisers (Maastricht University, 2012), Todd Hester (UT Austin, 2012), Dylan Alex Simon (NYU, Neuroscience, 2012), Peng Zang

(Georgia Tech, 2011), Umar Syed (Princeton, 2010), Emma Brunskill (MIT, 2009), Colin McMillen (CMU, 2009), Michael Smith (U. Maryland, Baltimore County, 2009), Emmanuel Rachelson (University of Toulouse, 2009), Michael Holmes (Georgia Tech, 2008), Özgür Şimşek (University of Massachusetts, 2008), Enrique Munoz de Cote (Politecnico di Milano, 2008), Thomas Degris (Université Pierre & Marie Curie, 2007), Charles Madeira (Université Pierre & Marie Curie, 2007), Shimon Whiteson (UT Austin, 2007), Fuat Balci (Rutgers, Psychology, 2007), Shahar Cohen (Tel-Aviv University, 2007), Darse Billings (U. Alberta, 2007), Charles Gretton (Australian National University, 2006), Matthijs Spaan (Universiteit Van Amsterdam, 2006), Mykel Kochenderfer (University of Edinburgh, 2006), Rosemary Emery-Montemerlo (CMU and Stanford, 2005), Yuhuan Chang (MIT, 2005), Pascal Poupart (U. Toronto, 2005), Sara Cordes (Rutgers, Psychology, 2005), Katja Verbeeck (Vrije Universiteit Brussel, 2004), Susana Eyheramendy (Rutgers, Statistics, 2004), Martin Zinkevich (CMU, 2004), Joelle Pineau (CMU, 2004), Olivier Buffet (de l'Université Henri Poincaré, 2003), Michail Lagoudakis (Duke University, 2003), Miles Efron (University of North Carolina, Library School, 2003), Eric Hansen (University of Massachusetts, 1998).

• **Committee Member:**

- Ph.D.: Do Kook Choe (Brown, 2017), Hannah Quay-de la Vallee (Brown, 2017), Steven Gomez (Brown, 2015), Layla Oesper (Brown, 2015), Mark Buller (Brown, 2015), Eric Sodomka (Brown, 2013), Timothy Edmunds (Rutgers, 2008), Kleisouris Konstantinos (Rutgers, 2008), David DeVault (Rutgers, 2008), Christopher Mesteharm (Rutgers, 2007), Fabio Oliveira (Rutgers, 2007), Bing Bai (Rutgers, 2007), Ramana Isukapalli (Rutgers, 2007), Eiman Elnahrawy (Rutgers, 2006), Sofus Macskassy (Rutgers, 2002), Tong Luo (Duke, 2000), Kasturirangan Varadarajan (Duke, 1998), Deganit Armon (Duke, 1997).
- Master's: Yuriko Ishihara (Duke, 1998), Cecilia Magdalena Procopiuc (Duke, 1998), Steve Myers (Duke, 1997), Hakan Bakircioglu (Duke, 1997, EE).
- Undergraduate thesis: Gretta Bartels (1998).