

Satinder Singh Baveja

(Note: academic papers written as Satinder Singh)

Associate Professor
Dept. of Electrical Engineering & Computer Science
University of Michigan
Ann Arbor, MI 48109

phone: (734) 936-2831
fax: (734) 763-1260
email: baveja@umich.edu
<http://www.eecs.umich.edu/~baveja>

RESEARCH STATEMENT

My research goals are to develop a theory of and algorithms for building autonomous agents that can *learn* to be broadly competent in complex, partially known, stochastic, and real-world environments. Such agents must be able to sense their environments, build a representation of their *state* from sensory observations, project the consequences of their *actions*, and continually decide on a course of action that maximizes expected *reward*. The field of *reinforcement learning* has focused on these questions and thus has been the locus of my research contributions and effort. More recently, I have become interested in dealing explicitly with the fact that an agent's environment may have other agents in them and so multi-agent decision making and learning or computational game theory has become another research focus. Finally the fact that a single agent's intelligence may itself be organized in a distributed manner has led to forays into a new research area — computational mechanism design. Periodically I test my algorithms and ideas on applications and in the past these have included shopping agents/bots in electronic commerce, spoken natural language dialogue systems that use speech recognition, artificial-intelligence-based chat agents in online virtual communities, and adaptive network routing in telecommunication networks.

EDUCATION

Postdoctoral Fellowship (1993 – 1995), Massachusetts Institute of Technology
Brain and Cognitive Science Department
Advisor: Prof. Michael Jordan

Ph.D. in Computer Science (1993), University of Massachusetts, Amherst.
Dissertation Title: “Learning to Solve Markovian Decision Processes”.
Thesis Advisor: Prof. Andrew G. Barto.

B. Tech. in Electrical Engineering (1987), Indian Institute of Technology (IIT), New Delhi, India.

EXPERIENCE

June 2002 to present: Associate Professor, Department of Electrical Engineering & Computer Science, University of Michigan.

April 2001 to May 2002: Chief Scientist, Syntek Capital (a *sabbatical* period at a multinational venture capital company)
September 1998 to March 2001: Principal Technical Staff Member in the Artificial Intelligence Department at AT&T Labs-Research.
August 1996 to August 1998: Assistant Professor in the Department of Computer Science at the University of Colorado, Boulder.
September 1995 to July 1996: Senior Research Scientist at Harlequin Inc., Cambridge, MA.
September 1993 to August 1995: Postdoctoral Fellow, Center for Biological and Computational Learning, Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology.
May 1988 to August 1993.: Research Assistant, Department of Computer Science, University of Massachusetts.
May 1986 to September 1986.: Software Developer, R&D Center, ICIM, New Delhi, India.

PUBLICATIONS

Journal Articles

1. **S. Singh**, D. Litman, M. Kearns and M. Walker. Optimizing Dialogue Management with Reinforcement Learning: Experiments with the NJFun System. In *Journal of Artificial Intelligence Research (JAIR)*, Volume 16, pages 105-133, 2002.
2. M. Kearns and **S. Singh**. Near-Optimal Performance for Reinforcement Learning in Polynomial Time. *Machine Learning*, Volume 49, Issue 2, pages 209-232, Nov 2002.
3. P. Stone, M. Littman, **S. Singh**, and M. Kearns. ATTac-2000: An Adaptive Autonomous Bidding Agent. In *Journal of Artificial Intelligence Research (JAIR)*, Volume 15, pages 189-206, 2001.
4. **S. Singh**, T. Jaakkola, M. Littman, and C. Szepesvari. Convergence Results for Single-Step On-Policy Reinforcement-Learning Algorithms. *Machine Learning*, Volume 38, Issue 3, pages 287-308, 2000.
5. R. Sutton, D. Precup, and **S. Singh**. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence* Volume 112, pages 181-211, 1999.
6. **S. Singh**, and P. Dayan. Analytical Mean Squared Error Curves for Temporal Difference Learning. *Machine Learning*, Volume 32, Issue 1, pages 5-40 1998.
7. **S. Singh**, and R. S. Sutton. Reinforcement Learning with Replacing Eligibility Traces. *Machine Learning*, Volume 22, Issue 1, pages 123-158, 1996.
8. **S. Singh**, and R. C. Yee. An Upper Bound on the Loss from Approximate Optimal-Value Functions. *Machine Learning*, Volume 16, Issue 3, pages 227-233, 1994.

9. **S. Singh**. Transfer of Learning by Composing Solutions for Elemental Sequential Tasks. *Machine Learning*, Volume 8, Issue 3, page 323–339, 1992.
10. N. E. Berthier, **S. Singh**, A. G. Barto & J. C. Houk. Distributed Representation of Limb Motor Programs in Arrays of Adjustable Pattern Generators. *Journal of Cognitive Neuroscience*, Volume 5:1, pages, 56-78, 1993
11. A.G. Barto, S. Bradtke & **S. Singh**. Real-time Learning and Control Using Asynchronous Dynamic Programming. *Artificial Intelligence*, Volume 72, pages 81-138, 1995.
12. T. Jaakkola, M. I. Jordan, and **S. Singh**. Stochastic Convergence of Iterative DP Algorithms. *Neural Computation*, Volume 6:6, pages 1185-1201, 1994

Refereed Selective-Conference Papers

1. **S. Singh**, A.G. Barto and N. Chentanez. Intrinsically Motivated Reinforcement Learning. To appear in Proceedings of *Advances in Neural Information Processing Systems 17* (NIPS), 2005.
2. D. Parkes, **S. Singh**, and D. Yanovsky. Approximately Efficient Online Mechanism Design. To appear in Proceedings of *Advances in Neural Information Processing Systems 17* (NIPS), 2005.
3. **S. Singh**, M.R. James and M. Rudary. Predictive State Representations: A New Theory for Modeling Dynamical Systems. In *Uncertainty in Artificial Intelligence: Proceedings of the Twentieth Conference* (UAI), pages 512-519, 2004.
4. M.R. James and **S. Singh**. Learning and Discovery of Predictive State Representations in Dynamical Systems with Reset. In *Proceedings of the Twenty-First International Conference on Machine Learning* (ICML), pages 417-424, 2004.
5. M. Rudary, **S. Singh** and M. Pollack. Adaptive Cognitive Orthotics: Combining Reinforcement Learning and Constraint-Based Temporal Reasoning. In *Proceedings of the Twenty-First International Conference on Machine Learning* (ICML), pages 719-726, 2004.
6. **S. Singh**, V. Soni and M. Wellman. Computing Approximate Bayes Nash Equilibria in Tree-Games of Incomplete Information. In *Proceedings of the 5th ACM Conference on Electronic Commerce* (EC), pages 81-90, 2004.
7. C. Kiekentveld, M. Wellman, **S. Singh**, J. Estelle, Y. Vorobeychik, V. Soni and M. Rudary. Distributed Feedback Control for Decision Making on Supply Chains. In *Proceedings of the Fourteenth International Conference on Automated Planning and Scheduling* (ICAPS), pages 384-392, 2004.

8. M. Rudary and **S. Singh**. A Nonlinear Predictive State Representation. In *Advances in Neural Information Processing Systems 16* (NIPS), pages 855-862, 2004.
9. D. Parkes and **S. Singh**. An MDP-Based Approach to Online Mechanism Design. In *Advances in Neural Information Processing Systems 16* (NIPS), pages 791-798, 2004.
10. **S. Singh**, M. Littman, N. Jong, D. Pardoe, and P. Stone. Learning Predictive State Representations. In *Proceedings of the Twentieth International Conference on Machine Learning* (ICML), pages 712-719, 2003.
11. M. Kearns, C. Isbell, **S. Singh**, D. Litman and J. Howe. CobotDS: A Spoken Dialogue System for Chat. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence* (AAAI), pages 425-430, 2002.
12. M. Kearns, M. Littman, and **S. Singh**. Graphical Models for Game Theory. In *Proceedings of 17th Annual Conference on Uncertainty in Artificial Intelligence* (UAI), pages 253-260, 2001.
13. M. Littman, M. Kearns and **S. Singh**. An Efficient Exact Algorithm for Single Connected Graphical Games. In *Advances in Neural Information Processing Systems 14* (NIPS), pages 817-823, 2002.
14. C.L. Isbell, C.R. Shelton, M. Kearns, **S. Singh** and P. Stone. Cobot: A Social Reinforcement Learning Agent. In *Advances in Neural Information Processing Systems 14* (NIPS), pages 1393-1400, 2002.
15. P. Stone, M. Littman, **S. Singh** and M. Kearns. ATTac-2000: An Adaptive Autonomous Bidding Agent. In *Proceedings of the Fifth International Conference on Autonomous Agents* (AGENTS), pages 238-245, 2001.
16. C. Isbell, C. Shelton, M. Kearns, **S. Singh** and P. Stone. A Social Reinforcement Learning Agent. In *Proceedings of the Fifth International Conference on Autonomous Agents* (AGENTS), pages 377-384, 2001. **Winner of Best Paper Award**
17. M.L. Littman, R.S. Sutton, and **S. Singh**. Predictive Representations of State. In *Advances in Neural Information Processing Systems 14* (NIPS), pages 1555-1561, 2001.
18. **S. Singh**, M. Kearns, D. Litman, and M. Walker. Empirical Evaluation of a Reinforcement Learning Spoken Dialogue System. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence* (AAAI), pages 645-651, 2000.
19. C. Isbell, M. Kearns, D. Korman, **S. Singh**, and P. Stone. Cobot in LambdaMOO: A Social Statistics Agent. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence* (AAAI), pages 36-41, 2000.

20. K. Myers, M. Kearns, **S. Singh**, and M. Walker. A Boosting Approach to Topic Spotting on Subdialogues. In *Proceedings of the Seventeenth International Conference on Machine Learning (ICML)*, pages 655-662, 2000.
21. D. Precup, R. Sutton and **S. Singh**. Eligibility Traces for Off-Policy Policy Evaluation. In *Proceedings of the Seventeenth International Conference on Machine Learning (ICML)*, pages 759-766, 2000.
22. M. Kearns, Y. Mansour, and **S. Singh**. Fast Planning in Stochastic Games. In *Uncertainty in Artificial Intelligence: Proceedings of the Sixteenth Conference (UAI)*, pages 309-316, 2000.
23. **S. Singh**, M. Kearns, and Y. Mansour. Nash Convergence of Gradient Dynamics in General-Sum Games. In *Uncertainty in Artificial Intelligence: Proceedings of the Sixteenth Conference (UAI)*, pages 541-548, 2000.
24. M. Kearns and **S. Singh**. "Bias-Variance" Error Bounds for Temporal Difference Updates. In *Proceedings of the Thirteenth Annual Conference on Computational Learning Theory (COLT)*, pages 142-147, 2000.
25. **S. Singh**, M. Kearns, D. Litman, and M. Walker. Reinforcement Learning for Spoken Dialogue Systems. In *Advances in Neural Information Processing Systems 12 (NIPS)*, 2000.
26. R. Sutton, D. McAllester, **S. Singh**, and Y. Mansour. Policy Gradient Methods for Reinforcement Learning with Function Approximation. In *Advances in Neural Information Processing Systems 12 (NIPS)*, 2000.
27. Y. Mansour and **S. Singh**. On the Complexity of Policy Iteration. In *Uncertainty in Artificial Intelligence: Proceedings of the Fifteenth Conference (UAI)*, pages 401-408, 1999.
28. D. McAllester and **S. Singh**. Approximate Planning for Factored POMDPs using Belief State Simplification. In *Uncertainty in Artificial Intelligence: Proceedings of the Fifteenth Conference (UAI)*, pages 409-416, 1999.
29. M. Kearns and **S. Singh**. Finite-Sample Convergence Rates for Q-learning and Indirect Algorithms. In *Advances in Neural Information Processing Systems 11 (NIPS)*, pages 996-1002, 1999.
30. T.X. Brown, H. Tong, and **S. Singh**. Optimizing Admission Control while Ensuring Quality of Service in Multimedia Networks via Reinforcement Learning. In *Advances in Neural Information Processing Systems 11 (NIPS)*, pages 982-988, 1999.
31. R.S. Sutton, **S. Singh**, D. Precup, and B. Ravindran. Improved Switching among Temporally Abstract Actions. In *Advances in Neural Information Processing Systems 11 (NIPS)*, pages 1066-1072, 1999.

32. J.K. Williams and **S. Singh**. Experimental Results on Learning Stochastic Memoryless Policies for Partially Observable Markov Decision Processes. In *Advances in Neural Information Processing Systems 11* (NIPS), pages 1073-1079, 1999.
33. M. Kearns and **S. Singh**. Near-Optimal Performance for Reinforcement Learning in Polynomial Time. In *Machine Learning: Proceedings of the Fifteenth International Conference* (ICML), pages 260-268, 1998.
34. J. Loch and **S. Singh**. Using Eligibility Traces to Find the Best Memoryless Policy in Partially Observable Markov Decision Processes. In *Machine Learning: Proceedings of the Fifteenth International Conference* (ICML), pages 323-331, 1998.
35. R. Sutton, D. Precup and **S. Singh**. Intra-Option Learning about Temporally Abstract Actions. In *Machine Learning: Proceedings of the Fifteenth International Conference* (ICML), pages 556-564, 1998.
36. **S. Singh**, and D. Cohn. How to Dynamically Merge Markov Decision Processes. In *Advances in Neural Information Processing Systems 10* (NIPS), pages 1057-1063, 1998.
37. **S. Singh**, and D. Bertsekas. Reinforcement Learning for Dynamic Channel Assignment. In *Advances in Neural Information Processing Systems 9* (NIPS), pages 974-980, 1997.
38. **S. Singh**, and P. Dayan. Analytical Mean Squared Error Curves in Temporal Difference Learning. In *Advances in Neural Information Processing Systems 9* (NIPS), pages 1054-1060, 1997.
39. D. Cohn, and **S. Singh**. Predicting Lifetimes in Dynamically Allocated Memory. In *Advances in Neural Information Processing Systems 9* (NIPS), pages 939-945, 1997.
40. L. Saul, and **S. Singh**. Learning Curve Bounds for Markov Decision Processes with Undiscounted Rewards. *Proceedings of the 9th Annual Conference on Computational Learning Theory* (COLT), pages 147-156, 1996.
41. P. Dayan, and **S. Singh**. Improving Policies without Measuring Merits, In *Advances in Neural Information Processing Systems 8* (NIPS), pages 1059-1065, 1996.
42. **S. Singh**, T. Jaakkola, and M. I. Jordan. Reinforcement Learning with Soft State Aggregation. *Advances in Neural Information Processing Systems 7* (NIPS), pages 361-368, 1995.
43. L. K. Saul, and **S. Singh**. Markov Decision Processes in Large State Spaces. In *Proceedings of the 8th Annual Conference on Computational Learning Theory* (COLT), pages 281-288, 1995.

44. T. Jaakkola, **S. Singh**, and M. I. Jordan. Reinforcement Learning Algorithm for Partially Observable Markov Decision Problems. In *Advances in Neural Information Processing Systems 7* (NIPS), pages 345-352, 1995.
45. **S. Singh**. Reinforcement Learning Algorithms for Average-Payoff Markovian Decision Processes. In *Proceedings of the 12th National Conference on Artificial Intelligence* (AAAI), pages 700-705, 1994.
46. **S. Singh**, T. Jaakkola, and M. I. Jordan. Learning Without State-Estimation in Partially Observable Markovian Decision Processes. In *Machine Learning: Proceedings of the Eleventh International Conference* (ICML), pages 284-292, 1994.
47. **S. Singh**, A. G. Barto, R. Grupen, and C. Connolly. Robust Reinforcement Learning in Motion Planning. In *Advances in Neural Information Processing Systems 6* (NIPS), pages 655-662, 1994.
48. T. Jaakkola, M. I. Jordan, and **S. Singh**. Stochastic Convergence of Iterative DP Algorithms. *Advances in Neural Information Processing Systems 6* (NIPS), pages 703-710, 1994.
49. **S. Singh**. The Efficient Learning of Multiple Task Sequences. In *Advances in Neural Information Processing Systems 4* (NIPS), pages 251-258, 1992.
50. **S. Singh**. Reinforcement learning with a hierarchy of abstract models. In *Proceedings of the Tenth International Conference on Artificial Intelligence* (AAAI), pages 202-207, 1992.
51. **S. Singh**. Scaling Reinforcement Learning Algorithms by Learning Variable Temporal Resolution Models. In *Machine Learning: Proceedings of the Ninth International Conference* (ICML), pages 406-415, 1992.
52. N. E. Berthier, **S. Singh**, A. G. Barto & J. C. Houk. A Cortico-Cerebellar model that learns to generate distributed motor commands to control a kinematic arm. In *Advances in Neural Information Processing Systems 4* (NIPS), pages 611-618, 1992.
53. **S. Singh**. Transfer of Learning Across Compositions of Sequential Tasks. In L. Birnbaum & G. Collins, editors, *Proceedings of the Eighth International Workshop on Machine Learning*. pages 348-352, 1991.

Other Refereed Conference Papers

1. A.G. Barto, **S. Singh**, and N. Chentanez. Intrinsically Motivated Learning of Hierarchical Collections of Skills. To appear in Proceedings of *International Conference on Developmental Learning (ICDL)*, 2004.
2. J. Estelle, Y. Vorobeychik, M.P. Wellman, **S. Singh**, C. Kiekintveld, and V. Soni. Strategic Interactions in the TAC 2003 Supply Chain Tournament. To appear in *Computer and Games (CG)*, 2004.
3. D. Precup, R. Sutton and **S. Singh**. Theoretical Results on Reinforcement Learning with Temporally Abstract Behaviors. In *Proceedings of the 10th European Conference on Machine Learning (ECML)*, pages 382-393, 1998.
4. J. C. Houk, A. G. Barto, L. N. Eisenman, J. Keiffer, **S. Singh**, T. Sinkjaer & D. Vyas. Motor Programs and Sensorimotor Integration. In *Proceedings of the International Joint Conference on Neural Networks*, pages I-147,148, 1990.

Refereed Workshop Papers

1. J. Estelle, Y. Vorobeychik, M.P. Wellman, **S. Singh**, C. Kiekintveld, and V. Soni. Strategic Procurement in TAC/SCM: An Empirical Game-Theoretic Analysis. In *Workshop on Trading Agent Design and Analysis (TADA)*, 2004.
2. V. Soni, **S. Singh** and M.P. Wellman. Computing Approximate Equilibria in Graphical Games on Arbitrary Graphs. In *Sixth Workshop on Game Theoretic and Decision Theoretic Agents (GTDT)*, 2004.
3. Y. Vorobeychik, M.P. Wellman, and **S. Singh**. Learning Payoff Functions in Infinite Games. In *AAAI Fall Symposium on Artificial Multi-Agent Learning*, 2004.
4. J. Csirik, M Littman, **S. Singh** and P. Stone. FAucs: An FCC Spectrum Auction Simulator for Autonomous Bidding Agents. In *Electronic Commerce: Proceedings of the Second International Workshop*. Springer Verlag, Heidelberg, 2001.
5. C. Isbell, M. Kearns, D. Korman, **S. Singh** and P. Stone. Cobot in LambdaMOO: A Social Statistics Agent. In *Workshop on Interactive Robotics and Entertainment (WIRE)*, 2000.
6. A. McGovern, D. Precup, B. Ravindran, **S. Singh** and R. Sutton. Hierarchical Optimal Control of MDPs. In *Tenth Yale Workshop on Adaptive and Learning Systems*, 1998.
7. D. Precup, R.S. Sutton, and **S. Singh**. Planning with Closed-Loop Macro Actions. In *Proceedings of AAAI Fall Symposium on Model-directed Autonomous Systems*, 1997.

8. R. S. Sutton & **S. Singh**. On Step-Size and Bias in Temporal-Difference Learning *In Eighth Yale Workshop on Adaptive and Learning Systems*, 1994.
9. A. G. Barto & **S. Singh**. Reinforcement Learning and Dynamic Programming. *In Sixth Yale Workshop on Adaptive and Learning Systems*, pages 83-88, 1990.

Refereed Book Chapters and Magazine Articles

1. C. Kiekintveld, M. Wellman, **S. Singh**, and V. Soni. Value-Driven Procurement in the TAC Supply Chain Game. *SIGecom Exchanges*, Volume 4.3, pages 9-19, 2004.
2. P. Stone, R. Sutton and **S. Singh**. Reinforcement Learning for 3 vs. 2 Keepaway. In P. Stone, T. Balch and G. Kraetschmar, editors, *RoboCup-2000: Robot Soccer World Cup IV*, Springer Verlag, 2001.
3. **S. Singh**, P. Norvig and D. Cohn. How to Make Software Agents Do the Right Thing: An Introduction to Reinforcement Learning. *In Dr. Dobbs journal*, pages 28-33, March 1997.
4. J.C. Houk, **S. Singh**, C. Fisher and A. Barto. An Adaptive Sensorimotor Network Inspired by the Physiology of the Cerebellum. In W.T. Miller, R.S. Sutton and P.J. Werbos, editors, *Neural Networks for Control*, pages 301-348, MA: MIT Press, 1990.
5. A. G. Barto & **S. Singh**. On the Computational Economics of Reinforcement Learning. *In Proceedings of the 1990 Connectionist Models Summer School*. 1990.

Unrefereed Papers

- S. Singh**. Introduction (to Machine Learning journal special issue on Reinforcement Learning), *Machine Learning*, Volume 49, Issue 2, pages 107-109, 2002.
- S. Singh**. Learning to Solve Markovian Decision Processes. CMPSCI Technical Report 93-77, (Ph.D Thesis). University of Massachusetts
- N. E. Berthier, **S. Singh**, A. G. Barto & J. C. Houk. Distributed Representation of Limb Motor Programs in Arrays of Adjustable Pattern Generators. Technical Report NPB 3, Institute of Neuroscience, 1991.
- A. G. Barto, N. E. Berthier, **S. Singh** & J. C. Houk. Network Model of the Cerebellum and Motor Cortex that Learns to Control Planar Limb Movements. *Soc. Neurosci. Abstr.* Vol. 16, Part 2:1223, 1990.

PROFESSIONAL ACTIVITIES

Conference Program Committees/Chairs:

- **Program Committee:** The Twentieth National Conference on Artificial Intelligence (AAAI), 2005.
- **Program Committee:** The Seventeenth Annual Conference on Learning Theory (COLT), 2004.
- **Workshops Co-Chair:** Neural Information Processing Systems (NIPS), 2004.
- **Workshops Co-Chair:** Neural Information Processing Systems (NIPS), 2003.
- **Program Committee:** Second International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS), 2003.
- **Area Chair:** (Equivalent to Senior Program Committee) International Conference on Machine Learning (ICML), 2003.
- **Program Committee:** Seventeenth National Conference on Artificial Intelligence, (AAAI) 2002.
- **Program Committee:** Eighteenth Conference on Uncertainty in Artificial Intelligence (UAI), 2002.
- **Program Committee:** Sixteenth Conference on Uncertainty in Artificial Intelligence (UAI), 2000.
- **Program Committee:** Fifteenth National Conference on Artificial Intelligence, (AAAI) 1998.
- **Program Committee:** Fifteenth International Conference on Machine Learning, (ICML) 1998.
- **Program Committee:** Fourteenth International Conference on Machine Learning, (ICML) 1997.
- **Program Committee:** Tenth Annual Conference on Computational Learning Theory (COLT), 1997.
- **Tutorials Chair:** Neural Information Processing Conference (NIPS) 1997.
- **Program Committee:** Thirteenth National Conference on Artificial Intelligence, (AAAI) 1996.
- **Area Chair:** (Equivalent to Senior Program Committee) Control and Navigation, Advances in Neural Information Processing Systems (NIPS), 1995.

Workshop/Symposia Committees and Organizing:

- **Co-Chair:** AAAI Fall Symposium on *Real-Life Reinforcement Learning*, 2004.
- **Co-Chair:** Workshop on *Predictive Representations of World Knowledge* at International Conference on Machine Learning (ICML), 2004.
- **Workshop Co-Chair:** Workshop on *Beyond MDPs: Representations and Algorithms* at Uncertainty in Artificial Intelligence (UAI), 2000.
- **Program Committee:** Workshop on *Value Function Approximation* at International Conference on Machine Learning (ICML), 1995.
- **Program Committee:** Workshop on *Reinforcement Learning: What we Know, What we Need* at International Conference on Machine Learning (ICML), 1993.

Editorial:

- Editor: Special Issue on Reinforcement Learning for the Machine Learning Journal, 2000.
- Editorial Board: *Journal of Machine Learning Research*, 2003-present.
- Editorial Board: *Machine Learning journal*, 2000-2002.
- Editorial Board: *Journal of Artificial Intelligence Research (JAIR)*, 2001-2003

Other Offices and Service:

- ISAT: Participant in ISAT meeting on *Network as Economy*, 2004.
- DARPA: Participated on DARPA workshop on *Cognitive Networking*, 2003.
- NSF: Served on a few proposal review panels.
- Best Student Paper Committee: International Conference on Machine Learning (ICML), 2003.
- Member of Research Advisory Board: Idilia Corporation, Montreal, CA, 2002-2003.
- Member of Research Advisory Board: Athene Software, Boulder CO, 1999-2001.
- Faculty on Summer Professional Seminar 1, New Computing Architectures: Theory and Applications of Neural Networks, CBCL, MIT. 1994.

Invited Talks:

- Invited colloquia and seminar talks at the University of Alberta, Georgia Tech, Massachusetts Institute of Technology, Carnegie Mellon University, University of Massachusetts at Amherst, University of Pennsylvania, Michigan State University, University of Colorado at Boulder, Brown University, UCSD, Salk Institute, University College (London), University of Washington, GTE Laboratories, AT&T Labs-Research, Microsoft Research, IBM Research at Hawthorne.
- At NIPS workshop on Machine Learning meets the User Interface, 2003.
- At NIPS workshop on Planning for the Real World: The promises and challenges of dealing with uncertainty.
- Invited plenary talk at the John Vincent Atanasoff Symposium at Iowa State University, November 1, 2003.
- At Workshop on Advances in Machine Learning, at Centre de Recherches Mathematiques, Montreal, June 8-11, 2003.
- At UM-SFI Workshop on Emergence and Engineering in Complex Systems, November 2003.
- At SARDINE (State-Action-Reward Day in New England), May 31, 2002.
- At The Learning Workshop at Snowbird, 2001 and 2000.
- At Planning with Partially Observable Markov Decision Processes, AAAI Fall Symposium, 1998.
- At NSF Workshop on Reinforcement Learning, 1996.
- At Robot Learning Workshop, *Neural Information Processing Systems*, 1992.
- At Active Learning Workshop, *Neural Information Processing Systems*, 1991.

Reviewing:

- Reviewed for various conferences including AAAI, IJCAI, ICML, NIPS, UAI, and AAMAS.
- Reviewed for various journals including Neural Networks, Neural Computation, Machine Learning, Artificial Intelligence, JAIR, and JMLR.