

Justin A. Boyan

Massachusetts Institute of Technology
Artificial Intelligence Lab
545 Technology Square NE43-753
Cambridge, MA 02139

Email: jboyan@arc.nasa.gov
Web: <http://ic.arc.nasa.gov/people/jboyan>
Phone: (617)-253-8005
Fax: (617)-253-7781

Summary

I am an applications-oriented machine learning researcher with strong mathematical skills, strong implementation skills, and significant entrepreneurial experience. My research focuses on creating efficient algorithms that use data to improve decision-making, with application to large-scale problems in optimization, control, scheduling, transportation planning, data mining and information retrieval.

I am currently a Visiting Scientist at the MIT AI Lab, on assignment from my position as a research scientist at NASA Ames Research Center.

Research Projects / Systems Developed

- 2000- **Adaptive Control of NASA Life Support Systems**
Developing learning-based controllers for life-support systems. These systems regulate the levels of air, water, food and energy for long-duration crew support in space. The associated control problems are dynamic, nonstationary, and safety-sensitive, requiring novel active-learning techniques. Work with Jeff Schneider, Leslie Kaelbling, and David Kortenkamp, in progress.

- 2000 **Internet Auction Trading Strategies**
Participated in the ICMAS-2000 Trading Agent Competition (TAC). This contest involved programming a travel agent to buy and sell airline tickets, hotel rooms, and entertainment tickets in 28 simultaneous Internet auctions, so as to construct profitable travel packages at minimum cost. Our entry included innovations in real-time resource allocation, pricing of single resources given combinatorial utilities, and risk mitigation. Work with Amy Greenwald.
Results: winner (out of 25) in preliminary round; co-winner (out of 12) in finals.

- 1999- **Optimal Route Planning Under Time-Dependent Uncertainty**
Developed new representations and solutions for time-sensitive stochastic planning problems. Applications to telescope experiment scheduling and multimodal transportation planning. Work with Michael Littman and Mike Mitzenmacher, in progress. [2] [1]

- 1996- **Learning Evaluation Functions for Global Optimization**
Combined dynamic programming, function approximation and local search techniques into an algorithm that automatically constructs high-quality evaluation functions for fast combinatorial optimization. Applications to VLSI design, medical robotics, satisfiability, Bayes net structure-finding, geographic visualization, and information retrieval. Ph.D. thesis work. [14] [3] [15] [6] [8] [10]

- 1994- **Scaling Up Reinforcement Learning / Value Function Approximation**
Developed new algorithms for Value Function Approximation, solving large-scale high-dimensional control and scheduling problems with approximate dynamic programming. Work with Andrew Moore. [5] [7] [11] [12] [16]

- 1996 **Information Retrieval / Web Indexing / Machine Learning**
Developed **LASER**, a machine-learning-based search engine for the World Wide Web. Using non-intrusive feedback gathered from system users doing Web searches, it optimizes its retrieval function so as to provide better page rankings. Work with Dayne Freitag and Thorsten Joachims.

- 1995-96 **Anonymizing Proxy System for the World Wide Web**
 Developed the **Anonymizer**, a custom proxy server that enables Web users to visit sites without revealing personal information such as their email address and Internet hostname. In April 1997, the system was sold to Infonex, Inc. and made available to the public at www.anonymizer.com. [9]
- 1993-94 **Reinforcement Learning for Network Routing / Multi-Agent Learning**
 Invented Q-routing, a parallel distributed reinforcement-learning algorithm for routing packet traffic in a computer network. The routing tables adapt to varying network topologies and traffic loads. Work with Michael Littman. [13]
- 1992-93 **Artificial Neural Network algorithms**
 Designed and tested extensions to cascading and modular neural network architectures. Work with Frank Fallside, Tony Robinson and Scott Fahlman.
- 1992-94 **Self-Learning Backgammon Software**
 Wrote software combining temporal-difference learning and modular neural networks in order to train, from scratch, an expert-level computer backgammon player. My software now serves as the playing engine of MVP Backgammon, an award-winning commercial product.
- 1986-95 **Sole proprietor, BOYAN Communications**
 Authored *BOYAN Communications*, a best-selling PC modem program. Maintained, documented, and marketed four major versions of the software over a ten-year life cycle.

Chronology / Awards

- **NASA Ames Research Center**, Mountain View, CA. Research Scientist, September 1998-present.
 - Visiting Scientist, **MIT AI Lab**, January 2000-present.
 - Group affiliation at Ames: Autonomy and Robotics Area, Computational Sciences Division
 - **Best Paper Award, ICML-99 conference**, June 1999 [5]
- **Carnegie Mellon University**, Pittsburgh, PA. **Ph.D., Computer Science**, August 1998; **M.S.**, May 1995.
 - **Outstanding Paper Award, AAAI-98 conference**, July 1998 [6]
 - NASA Graduate Student Researchers Program Fellowship, 1996-98
 - National Defense Science and Engineering Graduate Fellowship, 1992-95
 - National Science Foundation Fellowship, 1991 (declined)
 - Elected to Sigma Xi, 1995
- **University of Cambridge**, Cambridge, UK. **M.Phil.**, Computer Speech and Language Processing, August 1992.
 - Churchill Scholarship, 1991-92
- **University of Chicago**, Chicago, IL. **B.S.** with General Honors, Mathematics, June 1991.
 - Paul R. Cohen Memorial Prize, 1991
 - Elected to Phi Beta Kappa, 1990
 - University of Chicago College Honors Scholarship 1987-91

Professional Activities / Teaching

- 2000-2003 Editorial Board, *Machine Learning Journal*
- 1998 Invited speaker, Symposium on Applications of Reinforcement Learning, Stanford University
I was one of fifteen invited speakers at this symposium.
- 1996 Invited speaker, NSF Reinforcement Learning workshop, Harper's Ferry, WV
I was one of three graduate students invited to participate in this national workshop.
- 1995 Teaching Assistant, *How to Think (Like a Computer Scientist)*
- 1993 Teaching Assistant, *Artificial Neural Networks* (graduate course)

- 1992-1997 Organizer, CMU Reinforcement Learning Group
- 1986-93 Mathematics teacher, Johns Hopkins University *Center for Talented Youth* residential summer program

Selected Refereed Publications

- [1] Boyan, J. A. and M. Mitzenmacher. "Improved Results for Route Planning in Stochastic Transportation Networks." Accepted to the *12th Annual Symposium on Discrete Algorithms (SODA)*, 2001.
- [2] Boyan, J. A. and M. L. Littman. "Exact Solutions to Time-Dependent MDPs." To appear in *Advances in Neural Information Processing Systems (NIPS)*. MIT Press, 2001.
- [3] Boyan, J. A. and A. W. Moore. "Learning Evaluation Functions to Improve Local Search." *Journal of Machine Learning Research*, to appear, 2000.
- [4] Boyan, J. A. "Technical Update: Least-Squares Temporal Difference Learning." *Machine Learning Journal*, to appear, 2000.
- [5] Boyan, J. A. "Least-Squares Temporal Difference Learning." In Bratko, I., and Dzeroski, S., eds., *Machine Learning: Proceedings of the Sixteenth International Conference (ICML)*, 1999.
(Selected as ICML-99 Best Paper. One of 152 submissions received this honor.)
- [6] Boyan, J. A. and A. W. Moore. "Learning Evaluation Functions for Global Optimization and Boolean Satisfiability." *Fifteenth National Conference on Artificial Intelligence (AAAI)*, 1998.
(Selected as an AAAI-98 Outstanding Paper. Three of 475 submissions received this honor.)
- [7] Schneider, J. G., J. A. Boyan and A. W. Moore. "Value Function Based Production Scheduling." *Machine Learning: Proceedings of the Fifteenth International Conference (ICML)*, 1998.
- [8] Moore, A. W., J. G. Schneider, J. A. Boyan and M. S. Lee. "Q2: Memory-Based Active Learning for Optimizing Noisy Continuous Functions." *Machine Learning: Proceedings of the Fifteenth International Conference (ICML)*, 1998.
- [9] Boyan, J. A. "The Anonymizer: Protecting User Privacy on the Web." *Computer-Mediated Communication Magazine*, 4 (9), September 1997.
- [10] Boyan, J. A. and A. W. Moore. "Using Prediction to Improve Combinatorial Optimization Search." *Sixth International Workshop on Artificial Intelligence and Statistics (AISTATS)*, 1997.
- [11] Boyan, J. A. and A. W. Moore. "Learning Evaluation Functions for Large Acyclic Domains." In L. Saitta (ed.), *Machine Learning: Proceedings of the Thirteenth International Conference (ICML)*. Morgan Kaufmann, 1996.
- [12] Boyan, J. A., and A. W. Moore, "Generalization in Reinforcement Learning: Safely Approximating the Value Function." In Tesauro, G., D. S. Touretzky, and T. K. Leen (eds.), *Advances in Neural Information Processing Systems 7 (NIPS)*. MIT Press, 1995.
- [13] Boyan, J. A., and M. L. Littman, "Packet routing in dynamically changing networks: A reinforcement learning approach." In Cowan, J. D., Tesauro, G., and Alspector, J. (eds.), *Advances in Neural Information Processing Systems 6 (NIPS)*. Morgan Kaufmann, 1994.

Selected Unrefereed Publications

- [14] Boyan, J. A. and W. L. Buntine, eds. "Statistical Machine Learning for Large-Scale Optimization." *Neural Computing Surveys* 3, 2000.

- [15] Boyan, J. A. "Learning Evaluation Functions for Global Optimization." Ph.D. thesis, Carnegie Mellon University, August 1998.
- [16] Boyan, J. A., A. W. Moore, and R. S. Sutton, Editors. "Proceedings of the Workshop on Value Function Approximation, Machine Learning Conference 1995." Carnegie Mellon Technical Report CMU-CS-95-206.