

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] [18]

- 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. [17] [3] [19] [6] [8] [10] [12] [20]

- 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] [13] [14] [20] [21]

- 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. [11]

- 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]
- Sample of media coverage:
- *USA Today*, 6/4/96
 - *PC Magazine Internet User Site of the Day*, 6/13/96
 - *US News and World Report*, 6/17/96
 - *Washington Post*, 6/26/96
 - *Scientific American*, 10/96
 - *New York Times*, 1/13/97
 - *Pittsburgh Tribune-Review*, 4/14/97
- 1995-96 **Internet Privacy Demonstration**
Created a script to demonstrate how Internet users may reveal personal information to every web site they visit. This script won a Netscape "Bugs Bounty" award in January 1996 and was featured for several years on the home page of the Center for Democracy and Technology. [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. This study has been followed up by at least five separate groups of researchers. Work with Michael Littman. [15] [16]
- 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. [26]
- 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. [26]
- Developed the neural network brains for *MVP Backgammon*, a PC program being marketed as Shareware since fall 1995 (with Marc Ringuette). MVP-BG was nominated by the Shareware Industry Awards Foundation as one of the five best Shareware entertainment products of 1996.
 - AST International Games Olympiad, London, August 1992. My modular neural net-based program, *MAESTRO*, placed second in the backgammon competition.
- 1993- **Web site design**
- HTML: Living Green in Pittsburgh page, Japan travel diary
 - JavaScript: "SNARK" web search interface -- named "JavaScript of the Month" by Boardwatch Magazine, June 1996
 - Java: Interactive Electoral Map -- this applet is now installed at U.S.News & World Report Online
- 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.
- Software reviewed in over 100 publications, including:
 - *PC Week*, 6/2/87
 - *PC Magazine*, "The Best of 1987", 1/12/88
 - *PC World*, "The Best of Shareware", 8/88
 - *Home Computing*, "Editor's Pick", 6/89
 - Version 4F, translated into French (with Maurice Periard), released 1990

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
 - Co-winner, ICMAS-2000 Trading Agent Competition, July 2000
 - **Best Paper Award, ICML-99 conference**, June 1999 [5]
 - NASA Ames Spotlight Award, April 1999
- **Carnegie Mellon University**, Pittsburgh, PA. **Ph.D., Computer Science**, August 1998; **M.S.**, May 1995.
 - Dissertation: "Learning Evaluation Functions for Global Optimization" [19]
 - Thesis committee: Andrew Moore (co-chair), Scott Fahlman (co-chair), Tom Mitchell, Tom Dietterich (Oregon State University)
 - **Outstanding Paper Award, AAI-98 conference**, July 1998 [6]
 - NASA Graduate Student Researchers Program Fellowship, 1996-98
 - Pennsylvania Space Grant Fellowship, 1995
 - National Defense Science and Engineering Graduate Fellowship, 1992-95
 - National Science Foundation Fellowship, 1991 (declined)
 - Elected to Sigma Xi, 1995
 - First place tie, CMU Programming Contest, September 1994
 - Coursework:
 - Machine Learning (Mitchell/Blum/Furst)
 - Artificial Intelligence (Carbonell/Veloso)
 - Algorithms (Furst)
 - Computer Systems (Gross)
 - Programming Languages (Reynolds)
 - Math Fundamentals for Robotics (Erdmann)
 - Operating Systems (Johnson)
 - Statistical Inference (Junker)
 - German 1 (Brockmann)
- **University of Cambridge**, Cambridge, UK. **M.Phil.**, Computer Speech and Language Processing, August 1992.
 - Thesis: "Modular Neural Networks for Learning Context-Dependent Game Strategies" [26]
 - Advisors: Frank Fallside and Tony Robinson
 - Churchill Scholarship, awarded to ten U.S. students annually for scientific studies at Cambridge, 1991-92
 - Selected coursework:
 - Hidden Markov Models (Pulman)
 - Digital Signal Processing (Niranjan)
 - Psycholinguistics (Cutler)
 - Neural Networks (Fallside)
- **University of Chicago**, Chicago, IL. **B.S.** with General Honors, Mathematics, June 1991.
 - Cumulative undergraduate GPA: 3.93 / 4.00
 - Paul R. Cohen Memorial Prize, awarded to the graduating senior who has achieved the highest academic record in mathematics, 1991
 - Elected to Phi Beta Kappa, 1990
 - Dean's List, 1987-88, 88-89, 89-90, 90-91
 - Mensa Education & Research Foundation Scholarship, 1991
 - State Farm Foundation Exceptional Student Fellowship, 1991
 - *Time Magazine* College Achievement Award, 1990
 - University of Chicago College Honors Scholarship (merit-based half-tuition scholarship), 1987-91
- **Oakland Mills High School**, Columbia, MD. May 1987.
 - Valedictorian.
 - Maryland Math League: First place in 1986-87 individual competition
 - Independent research project: "Computer Simulations of Brownian Motion"

- Supervisor: David Thirumalai, Institute for Physical Science and Technology, University of Maryland
- First place overall, physical sciences, Baltimore Science Fair, 1987
- Third place grand award, physics division, International Science and Engineering Fair, San Juan, PR, 1987
- Computer Team -- Captain, grades 10-12
 - First place, First Maryland Computer Bowl, 1986
 - First place, Second Maryland Computer Bowl, 1987

Professional Activities / Teaching

- 2000-2003 Editorial Board, *Machine Learning Journal*
- 2000 Co-organizer, AAI-2000 Workshop on Artificial Intelligence for Web Search
- 2000 Organizer, MIT Statistical AI Reading Group
- 1999 Organizer, Workshop on Statistical Machine Learning for Large-Scale Optimization, Stockholm, Sweden
I co-organized this IJCAI workshop with Wray Buntine. An edited version of the proceedings will appear in Neural Computing Surveys. [17]
- 1999 Judge, Siemens Westinghouse Science and Technology Competition, Princeton, NJ
- 1998 Judge, Santa Clara Valley Science and Engineering Fair, San Jose, CA
- 1998 Invited speaker, Symposium on Applications of Reinforcement Learning, Stanford University
I was one of fifteen invited speakers at this symposium.
- 1997 Invited speaker, NSF/CNPq Joint Workshop on Intelligent Robotic Agents, Porto Alegre, Brazil
- 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 Organizer, Workshop on Value Function Approximation, Tahoe City, CA
I co-organized this Machine Learning Conference workshop with Andrew Moore and Rich Sutton.
- 1995 Teaching Assistant, *How to Think (Like a Computer Scientist)*
Instructor: Rudich.
- 1993 Teaching Assistant, *Artificial Neural Networks* (graduate course)
Instructors: Touretzky, Waibel, Fahlman, and Pomerleau.
- 1992-1997 Organizer, CMU Reinforcement Learning Group
- 1993- Paper Referee:
- American Association for Artificial Intelligence conference (1998, 2000)
 - Artificial Intelligence Journal (2000)
 - International Conference on Machine Learning (1993, 2000)
 - International Joint Conference on Artificial Intelligence (1995, 1997)
 - IEEE Transactions on Neural Networks (1998)
 - International Symposium on Artificial Intelligence, Robotics and Automation in Space (1999)
 - Journal of Artificial Intelligence Research (1994, 1998, 1999)
 - Machine Learning Journal (1993, 1994, 1997, 1999)
 - Neural Computation (1997)
 - Neural Networks (1994)
 - Neural Information Processing Systems conference (1993, 1994, 1998, 1999)

- 1986-93 Mathematics teacher, Johns Hopkins University *Center for Talented Youth* residential summer program
- Instructor: Lancaster site, 1993, 91, 89, 88; Norton, MA site, 1990
 - Teaching assistant: Geneva, Switzerland site, 1989; Lancaster, PA site, 1987, 86

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.
(Selected for oral presentation.)
- [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. Also presented at Symposium on Applications of Reinforcement Learning, Stanford, March 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., D. Freitag and T. Joachims. "A Machine Learning Architecture for Optimizing Web Search Engines." *Proceedings of the AAAI workshop on Internet-Based Information Systems*, AAAI Technical Report WS-96-06, 1996.
- [12] Boyan, J. A. "A Reinforcement Learning Framework for Combinatorial Optimization." (student abstract) Thirteenth National Conference on Artificial Intelligence (AAAI), 1996.
- [13] 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.
- [14] 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.
(Selected for oral presentation. Approximately 30 out of 500 submissions received this honor.)

- [15] Boyan, J. A., and M. L. Littman, "Packet routing in dynamically changing networks: A reinforcement learning approach." In Cowan, J. D., Tesauro, G., and Alspecter, J. (eds.), *Advances in Neural Information Processing Systems 6 (NIPS)*. Morgan Kaufmann, 1994.
- [16] Littman, M. L. and J. A. Boyan. "A Distributed Reinforcement Learning Scheme for Network Routing." In Alspecter, J., Goodman, R., and Brown, T. X. (eds.), *Proceedings of the International Workshop on Applications of Neural Networks to Telecommunications*, October 1993. Also appeared as: Carnegie Mellon Technical Report CMU-CS-93-165.

Unrefereed Publications/Talks

- [17] Boyan, J. A. and W. L. Buntine, eds. "Statistical Machine Learning for Large-Scale Optimization." *Neural Computing Surveys* 3, 2000.
- [18] Boyan, J. A. and M. L. Littman. "Representations and Algorithms for Time-Dependent MDPs." Presented at the workshop on "Beyond MDPs", UAI-2000.
- [19] Boyan, J. A. "Learning Evaluation Functions for Global Optimization." Ph.D. thesis, Carnegie Mellon University, August 1998. Also appeared as: Carnegie Mellon Technical Report CMU-CS-98-152. Oral defense given at CMU, May 1998. Also presented at Stanford University, UC Berkeley, AT&T Labs, SRI International, NASA Ames Research Center, Compaq CRL, MERL, Microsoft Research, i2 Research, and NEC Research.
- [20] Boyan, J. A. "Learning Evaluation Functions." Ph.D. Thesis Proposal, CMU, June 1996.
- [21] 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.
- [22] Boyan, J. A. "Active Learning for Optimal Control in Acyclic Domains." In Proceedings of AAAI Symposium on Active Learning, Autumn 1995.
- [23] Boyan, J. A. "Two Algorithms for Robust VFA By Working Backwards." Presented at ML95 workshop on Value Function Approximation, July 1995.
- [24] Boyan, J. A. "Safely Approximating the Value Function." Presented at Action Learning workshop, MIT, March 1995. Similar talks given at Brown University and the University of Massachusetts, 1995, and Stanford University, 1994.
- [25] Boyan, J. A. "MAESTRO 1.0: A Modular Neural Network for Learning Context-Dependent Backgammon Strategies by Self-Play." Presented at the Fourth International Conference on Computer Games, London, August 1992.
- [26] Boyan, J. A. "Modular Neural Networks for Learning Context-Dependent Game Strategies." Master's thesis, Department of Engineering and Computer Laboratory, University of Cambridge, 1992.