

# Lihong Li

Yahoo! Research  
4401 Great America Parkway  
Santa Clara, CA, USA 95054

lihong@yahoo-inc.com  
<http://www.cs.rutgers.edu/~lihong>

---

## RESEARCH INTERESTS

My primary research interests are in three closely related subfields of artificial intelligence: *machine learning*, *reinforcement learning*, and *decision-theoretic planning*. I am most interested in creating new algorithms with nice theoretical guarantees and practical values (especially in large-scale applications), as well as developing novel theoretical insights into existing algorithms and problems. Recently, I am working on machine-learning problems with an interactive nature, including online learning, bandit learning, and reinforcement learning.

## EDUCATION

- |                   |        |  |
|-------------------|--------|--|
| 01/2005 – 05/2009 | Ph.D.  | Computer Science, Rutgers University, USA<br>Advisor: Michael L. Littman                         |
| 09/2002 – 07/2004 | M.Sc.  | Computing Science, University of Alberta, Canada<br>Co-advisors: Vadim Bulitko & Russell Greiner |
| 09/1998 – 07/2002 | B.Eng. | Computer Science and Technology, Tsinghua University, China                                      |

## RESEARCH & INDUSTRY EXPERIENCE

- |                   |  |
|-------------------|--|
| 06/2009 – present | Postdoctoral Scientist at Yahoo! Research<br>Playing with machine-learning techniques, theory, and applications.   |
| 06/2008 – 08/2008 | Research Intern at AT&T Shannon Labs<br>Developed and implemented algorithms for scaling up reinforcement learning and feature selection in spoken dialog management.  |
| 05/2007 – 08/2007 | Research Intern at Yahoo! Research NYC<br>Researched and implemented efficient online-learning algorithms for large-scale problems (Vowpal Wabbit project). Developed learning algorithms for sponsored search.          |
| 05/2006 – 08/2006 | Engineering Intern at Google NYC<br>Designed and implemented software modules for an object identification task using naïve Bayes net.   |
| 01/2005 – 05/2009 | Graduate Research Assistant at the Rutgers University<br>Researched, developed, and analyzed models and algorithms for reinforcement learning, decision-theoretic planning, and machine learning.                        |
| 09/2002 – 07/2004 | Research Assistant at the University of Alberta<br>Researched, developed, and analyzed algorithms for reinforcement learning. Explored and tested machine-learning techniques in an automatic object recognition system. |
-

**TEACHING EXPERIENCE**

- Spring 2009            Guest lecturer for a graduate-level course at the Rutgers University  
Taught the least-squares policy iteration (LSPI) algorithm in the course “Learning and Sequential Decision Making”.
- 09/2007 – 12/2007    Co-organizer for a graduate seminar at the Rutgers University  
Compiled reading materials, arranged weekly meetings, and presented papers for “Planning in Learned Environments” (w/ Michael Littman).
- 05/2005 – 08/2005    Organizer for a graduate seminar at the Rutgers University  
Compiled reading materials, arranged weekly meetings, presented papers, and invited an external speaker for “Abstractions and Hierarchies for Learning and Planning”.
- 09/2002 – 07/2004    Teaching Assistant at the University of Alberta  
Taught seminar sessions and graded homework for the undergraduate course on discrete mathematics: “Formal Systems and Logic in Computing Science”.

**SELECTED RECENT AWARDS**

- 2003    Canada    J. Gordin Kaplan Scholarship, University of Alberta
- 2003    Canada    GSA Professional Development Grant, University of Alberta
- 2004    Canada    Teaching Assistant Award, University of Alberta
- 2006    USA        Winner in PENTATHLON and first place in PUDDLEWORLD, First Annual Competition of Reinforcement Learning (w/ A. Nouri, T.J. Walsh, & M.L. Littman)
- 2006    USA        Best Student Poster Award, New York Academy of Sciences
- 2008    USA        Best Student Paper Award, ICML’2008
- 2008    USA        Google Student Award Winner, New York Academy of Sciences

**PROFESSIONAL ACTIVITIES**

- Open source contributions
  - Vowpal Wabbit: an open source project with John Langford and Alexander L. Strehl for fast online-learning in large-scale prediction problems. URL: <http://www.hunch.net/~vw> .
- Organization
  - Funding Chair, 2009 Reinforcement Learning Competition (ICML/UAI/COLT’09 Workshop)
- Reviewer for journals
  - Journal of Machine Learning Research (2005–2008)
  - Journal of Artificial Intelligence Research (2008–2009)
  - Data Mining and Knowledge Discovery (2008)
  - Machine Learning (2009)
  - Neural Computation (2009)
  - Neurocomputing (2009)
  - IEEE Transactions on Wireless Communications (2009)
  - Journal of Computer Science and Technology (2009)
  - Artificial Intelligence Communications (2009)
- Reviewer or program committee member for conferences:
  - National Conferences on Artificial Intelligence (AAAI): 2006, 2008, 2010
  - International Joint Conferences on Artificial Intelligence (IJCAI): 2007

- Annual Meetings on Neural Information Processing Systems (NIPS): 2008, 2009
- International Conferences on Machine Learning (ICML): 2009, 2010
- European Conferences on Machine Learning (ECML): 2009
- Conference presentations (oral/spotlight/poster) at: IJCAI’03, NIPS’03 (workshop), ECML’04, AAAI’05, ICML’06, AI&Math’08, ICML’08, NIPS’08, Asilomar SSC’09.

### **INVITED TALKS**

- “Provably Efficient Exploration in Reinforcement Learning”
  - AT&T Shannon Labs, Florham Park, NJ, USA. January 2008.
- “Go as a Testbed for Advancing Reinforcement Learning Research”
  - DARPA Information Processing Technology meeting, Arlington, VA, USA. February 2008.
- “Sparse Online Learning via Truncated Gradient”
  - Text Analysis and Machine Learning Group, University of Ottawa, Ottawa, ON, Canada. May 2008.
  - Department of Information Analysis & Management, NEC Laboratories America, Cupertino, CA, USA. April 2009.
  - eBay Research Labs, San Jose, CA, USA. April 2009.
  - Asilomar Conference on Signals, Systems, and Computers, Pacific Grove, CA, USA. November 2009.
- “Knows What It Knows: A Framework for Self-Aware Learning”
  - Reasoning and Learning Laboratory, McGill University, McGill, QC, Canada. May 2008.
- “A Unifying Framework for Computational Reinforcement Learning Theory”
  - Yahoo! Research, New York, NY, USA. January 2009.
  - Google Research, New York, NY, USA. April 2009.
  - Yahoo! Research, Sunnyvale, CA, USA. April 2009.

### **SELECTED PUBLICATIONS**

#### **Journal Papers**

1. A.L. Strehl, *L. Li*, and M.L. Littman: Reinforcement learning in finite MDPs: PAC analysis. In the *Journal of Machine Learning Research*, 10:2413–2444, 2009.
2. E. Brunskill, B.R. Leffler, *L. Li*, M.L. Littman, and N. Roy: Provably efficient learning with typed parametric models. In the *Journal of Machine Learning Research*, 10:1955–1988, 2009.
3. J. Langford, *L. Li*, J. Wortman, and Y. Vorobeychik: Maintaining equilibria during exploration in sponsored search auctions. To appear in *Algorithmica* (Special Issue on WINE-07), 2009.
4. J. Langford, *L. Li*, and T. Zhang: Sparse online learning via truncated gradient. In the *Journal of Machine Learning Research*, 10:777–801, 2009.
5. T.J. Walsh, A. Nouri, *L. Li*, and M.L. Littman: Planning and learning in environments with delayed feedback. In the *Journal of Autonomous Agents and Multi-Agent Systems*, 18(1):83–105, 2009.
6. *L. Li*, V. Bulitko, and R. Greiner: Focus of attention in reinforcement learning. In the *Journal of Universal Computer Science*, 13(9):1246–1269, 2007.
7. *L. Li*, M. Shao, Z. Zheng, C. He, and Z.-H. Du: Typical XML document transformation methods and an application system (in Chinese). *Computer Science*, 30(2):40–44, February, 2003.

#### **Conference Papers**

1. Y. Xie, Y. Zhang, and *L. Li*: Neuro-fuzzy reinforcement learning for adaptive intersection traffic signal control. In the *Annual Meeting of Transportation Research Board*, 2010.

2. L. Li, J.D. Williams, and S. Balakrishnan: Reinforcement learning for spoken dialog management using least-squares policy iteration and fast feature selection. In *the Tenth Annual Conference of the International Speech Communication Association (INTERSPEECH)*, 2009.
3. C. Diuk, L. Li, and B.R. Leffler: The adaptive  $k$ -meteorologists problem and its application to structure learning and feature selection in reinforcement learning. In *the Twenty-Sixth International Conference on Machine Learning (ICML)*, 2009.
4. J. Asmuth, L. Li, M.L. Littman, A. Nouri, and D. Wingate: A Bayesian sampling approach to exploration in reinforcement learning. In *the Twenty-Fifth International Conference on Uncertainty in Artificial Intelligence (UAI)*, 2009.
5. L. Li and M.L. Littman and C.R. Mansley: Online exploration in least-squares policy iteration. In *the Eighth International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2009.
6. L. Langford, L. Li, and T. Zhang: Sparse online learning via truncated gradient. In *Advances in Neural Information Processing Systems 21 (NIPS)*, 2009.
7. L. Li: A worst-case comparison between temporal difference and residual gradient. In *the Twenty-Fifth International Conference on Machine Learning (ICML)*, 2008.
8. L. Li, M.L. Littman, and T.J. Walsh: Knows what it knows: A framework for self-aware learning. In *the Twenty-Fifth International Conference on Machine Learning (ICML)*, 2008. **Co-winner of the Best Student Paper award of ICML. A Google Student Award winner at the New York Academy of Sciences Symposium on Machine Learning, 2008.**
9. R. Parr, L. Li, G. Taylor, C. Painter-Wakefield, and M.L. Littman: An analysis of linear models, linear value function approximation, and feature selection for reinforcement learning. In *the Twenty-Fifth International Conference on Machine Learning (ICML)*, 2008.
10. E. Brunskill, B.R. Leffler, L. Li, M.L. Littman, and N. Roy: CORL: A continuous-state offset-dynamics reinforcement learner. In *the Twenty-Fourth Conference on Uncertainty in Artificial Intelligence (UAI)*, 2008.
11. L. Li and M.L. Littman: Efficient value-function approximation via online linear regression. In *the Tenth International Symposium on Artificial Intelligence and Mathematics (AI&Math)*, 2008.
12. J. Wortman, Y. Vorobeychik, L. Li, and J. Langford: Maintaining equilibria during exploration in sponsored search auctions. In *the Third International Workshop on Internet and Network Economics (WINE)*, LNCS 4858, 2007.
13. T.J. Walsh, A. Nouri, L. Li, and M.L. Littman: Planning and learning in environments with delayed feedback. In *the Eighteenth European Conference on Machine Learning (ECML)*, LNCS 4701, 2007.
14. R. Parr, C. Painter-Wakefield, L. Li, and M.L. Littman: Analyzing feature generation for value-function approximation. In *the Twenty-Fourth International Conference on Machine Learning (ICML)*, 2007.
15. A.L. Strehl, L. Li, E. Wiewiora, J. Langford, and M.L. Littman: PAC model-free reinforcement learning. In *the Twenty-Third International Conference on Machine Learning (ICML)*, 2006. **Best Student Poster Award winner at the New York Academy of Sciences Symposium on Machine Learning, 2006.**
16. A.L. Strehl, L. Li, and M.L. Littman: Incremental model-based learners with formal learning-time guarantees. In *the Twenty-Second Conference on Uncertainty in Artificial Intelligence (UAI)*, 2006.
17. L. Li, T.J. Walsh, and M.L. Littman: Towards a unified theory of state abstraction for MDPs. In *the Ninth International Symposium on Artificial Intelligence and Mathematics (AI&Math)*, 2006.
18. L. Li, M.L. Littman: Lazy approximation for solving continuous finite-horizon MDPs. In *the Twentieth National Conference on Artificial Intelligence (AAAI)*, 2005.
19. L. Li, V. Bulitko, and R. Greiner: Batch reinforcement learning with state importance (extended abstract). In *the Fifteenth European Conference on Machine Learning (ECML)*, LNCS 3201, 2004.
20. V. Bulitko, L. Li, R. Greiner, and I. Levner: Lookahead pathologies for single agent search (poster paper). In *the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI)*, 2003.

21. I. Levner, V. Bulitko, L. Li, G. Lee, and R. Greiner: Towards automated creation of image interpretation systems. In *the Sixteenth Australian Joint Conference on Artificial Intelligence*, LNCS 2903, 2003.
22. L. Li, V. Bulitko, R. Greiner, and I. Levner: Improving an adaptive image interpretation system by leveraging. In *the Eighth Australian and New Zealand Intelligent Information System Conference (ANZIIS)*, 2003.

### Book

1. M. Shao, L. Li, Z. Zheng, and C. He: Practical Programming in XML. *Tsinghua University Press*, Beijing, China, December, 2002. ISBN 7-900643-85-0.

### Theses

1. L. Li: A unifying framework for computational reinforcement learning theory. *Doctoral dissertation*, Department of Computer Science, Rutgers University, New Brunswick, NJ, USA, May, 2009.
2. L. Li: Focus of attention in reinforcement learning. *MSc thesis*, Department of Computing Science, University of Alberta, Edmonton, Alberta, Canada, July, 2004.
3. L. Li: Design and implementation of an agent communication module based on KQML. *Bachelor degree thesis*, Department of Computer Science and Technology, Tsinghua University, Beijing, China, June, 2002.

### Working Papers (Under Review)

1. L. Li, W. Chu, J. Langford, K. Papineni, and R.E. Schapire: A contextual-bandit approach to personalized news article recommendation. Submitted to *WWW-10*, 2010.
2. L. Li, M.L. Littman, T.J. Walsh, and A.L. Strehl: Knows what it knows: A framework for self-aware learning. Invited for submission to the *Machine Learning* journal.
3. L. Li and M.L. Littman: Prioritized sweeping converges to the optimal value function. Technical report DCS-TR-631, Department of Computer Science, Rutgers University, May 2008. To be submitted to the *Machine Learning* journal.
4. L. Li and M.L. Littman: Efficient exploration with linear value-function approximation. Under review by the *Annals of Mathematics and Artificial Intelligence*.
5. J. Langford, L. Li, R.P. McAfee, and K. Papineni: Voluntary admission control for Internet traffic management. Working paper.

### ACADEMIC COLLABORATORS (COAUTHORS)

- *AT&T Shannon Labs*: Suhrid Balakrishnan, Jason D. Williams
- *Duke University*: Christopher Painter-Wakefield, Ronald Parr, Gavin Taylor
- *Massachusetts Institute of Technology*: Emma Brunskill, Nicholas Roy, David Wingate
- *Princeton University*: Robert E. Schapire
- *Rutgers University*: John Asmuth, Carlos Diuk, Bethany R. Leffler, Michael L. Littman, Christopher R. Mansley, Ali Nouri, Alexander L. Strehl, Thomas J. Walsh, Tong Zhang
- *Texas A&M University*: Yuanchang Xie, Yunlong Zhang
- *Tsinghua University*: Zhi-Hui Du, Chuan He, Min Shao, Zhengkun Zheng
- *University of Alberta*: Vadim Bulitko, Russell Greiner, Greg Lee, Ilya Levner
- *University of California at San Diego*: Eric Wiewiora
- *University of Michigan*: Yevgeniy Vorobeychik
- *University of Pennsylvania*: Jennifer Wortman
- *Yahoo! Research*: Wei Chu, John Langford, R. Preston McAfee, Kishore Papineni, Alexander L. Strehl, Tong Zhang