

Thomas J. Walsh

Thomas J. Walsh

Email: thomasjwalsh [at] gmail.com

Web: www.cs.arizona.edu/~twalsh

Department of Computer Science
University of Arizona
1040 E. 4th Street
Tucson, AZ 85721

Primary Research

My dissertation research focused on representation issues for agents in the *Reinforcement Learning (RL)* subsection of machine learning, where an agent must efficiently learn to act optimally in an environment based on feedback. Specifically, I designed agents to learn models in environments that can be compactly described at a relational level (examples range from “blocks world” to web-service domains). By using modern machine learning techniques, we were able to preserve important theoretical guarantees in terms of sample efficiency and near-optimal behavior, even with these more powerful representations. In my post-doctoral work, I have focused on similar problems, but now with more realistic tasks (including robotics, and electronic tutoring) that usually include a human teacher demonstrating examples of good behavior. I’ve also done work in other machine learning areas, including reinforcement learning on Aibo robots, text-mining on web-documents, and preference learning.

Education

Rutgers University, New Brunswick, New Jersey

PhD in Computer Science, 2010

Dissertation: “Efficient Learning of Relational Models for Sequential Decision Making”

Research Advisor: Dr. Michael L. Littman

University of Maryland, Baltimore County, Baltimore, Maryland

Bachelor of Science in Computer Science, 2003

Research Advisor: Dr. Marie desJardins

Publications

Journal Papers:

Lihong Li, Michael L. Littman, Thomas J. Walsh, Alexander L. Strehl. “Knows what it Knows: A Framework For Self-Aware Learning” (accepted in Machine Learning Journal)

Fusun Yaman, Thomas J. Walsh, Michael L. Littman, Marie desJardins. “Democratic Approximation of Lexicographic Preference Models” (accepted to Artificial Intelligence Journal)

Thomas J. Walsh, Ali Nouri, Lihong Li, and Michael L. Littman. “Learning and Planning in Environments with Delayed Feedback”. *Journal of Autonomous Agents and Multi-Agent Systems*, Volume 18, Issue1, 83-101, February, 2009.

Dennis D.Y. Kim, Thomas T.Y. Kim, Thomas Walsh, Yoshifumi Kobayashi, Tara C. Matise, Steven Buyske, and Abram Gabriel. “Widespread RNA Editing of Embedded Alu Elements in the Human Transcriptome” *Genome Research*. 2004 14 (September): 1719-1725.

Conference Papers:

Thomas J. Walsh, Sergiu Goshin, and Michael L. Littman Integrating “Sample-based Planning and Model-based Reinforcement Learning”. *In Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI-10)*, Atlanta, GA, 2010.

Thomas J. Walsh

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 (ICML-10)*, Haifa, Israel, 2010.

Thomas J. Walsh, István Szita, Carlos Diuk, and Michael L. Littman. “Exploring Compact Reinforcement-Learning Representations with Linear Regression” In *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence (UAI-09)*, Montreal, Quebec, 2009.

Thomas J. Walsh and Michael L. Littman. “Efficient Learning of Action Schemas and Web-Service Descriptions”. In *Proceedings of Twenty-Third AAAI Conference on Artificial Intelligence*. Chicago, IL, 2008.
(also published and presented at the IJCAI 2009 Workshop on Learning Structural Knowledge From Observations)

Lihong Li, Michael L. Littman, Thomas J. Walsh. “Knows what it Knows: A Framework For Self-Aware Learning”. To appear in *Proceedings of the 25th International Conference on Machine Learning*. Helsinki, Finland, 2008.
(also presented at the 2008 European Workshop on Reinforcement Learning (EWRL-08))

Fusun Yaman, Thomas J. Walsh, Michael L. Littman, Marie desJardins. “Democratic Approximation of Lexicographic Preference Models”. In *Proceedings of the 25th International Conference on Machine Learning*. Helsinki, Finland, 2008.
(also published and presented at the 4th Multidisciplinary Workshop on Advances in Preference Handling at AAAI 2008)

Thomas J. Walsh, Ali Nouri, Lihong Li, and Michael L. Littman. “Planning and Learning in Environments with Delayed Feedback”, In *Proceedings of the Eighteenth European Conference on Machine Learning*, Warsaw, Poland, 2007.

Lihong Li, Thomas J. Walsh, and Michael L. Littman. “Towards a Unified Theory of State Abstraction for MDPs”. *Proceedings of the Ninth International Symposium on Artificial Intelligence and Mathematics (AIMA06)*, Ft. Lauderdale, FL, 2006.

Bethany R. Leffler, Michael L. Littman, Alexander L. Strehl, Thomas J. Walsh. “Efficient Exploration With Latent Structure”. In *Proceedings of Robotics: Science and Systems*. Cambridge, Massachusetts, 2005.

Dissertation:

Thomas J. Walsh, *Efficient Learning of Relational Models for Sequential Decision Making* PhD Thesis, Rutgers university, 2010.

Book Chapters:

Fusun Yaman, Thomas J. Walsh, Michael L. Littman, Marie desJardins. “Learning Lexicographic Preference Models”. In *Preference Learning* (ed. Eyke Hüllermeier and Johannes Fürnkranz), To be published by Springer-Verlang, 2010.

Magazine Articles:

Thomas J. Walsh and D. Richard Kuhn. “Challenges in Securing Voice over IP” *IEEE Security & Privacy*. Vol 3(3) 2005 (May/June) : 44-49.

Government Publications:

D. Richard Kuhn, Thomas J. Walsh, Steffen Fries. *Security Considerations for Voice Over IP Systems* Special Publication from the National Institute of Standards and Technology, 2005

Workshop/Symposia Papers:

Daniel Hewlett, Thomas J. Walsh, and Paul R. Cohen. “A Framework for Teaching and Executing Verb Phrases”, *AAAI Spring Symposium on Bridging the Gaps in Human-Agent Collaboration*, 2011

Thomas J. Walsh

Tasneem Kaochar, Raquel Torres Peralta, Clayton T. Morrison, Thomas J. Walsh, Ian R. Fasel, Sumin Beyon, Anh Tran, Jeremy Wright and Paul R. Cohen. "Human Natural Instruction of a Simulated Electronic Student" *AAAI Spring Symposium on Bridging the Gaps in Human-Agent Collaboration*, 2011

Thomas J. Walsh and Michael L. Littman. "A Multiple representation approach to learning dynamical Systems" *AAAI Fall Symposium on Representation Change*, Washington D.C., 2007

Thomas J. Walsh and Michael L. Littman. "Planning with Conceptual Models Mined from User Behavior" *AAAI Workshop on Acquiring Planning Knowledge via Demonstration*, Vancouver, BC 2007.

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*, Pittsburgh, PA, 2006.

Alex Borgida, Thomas J. Walsh, and Haym Hirsh. "Towards Measuring Similarity in Description Logics." *In Proceedings of the 2005 International Workshop on Description Logics (DL2005)*, Edinburgh, Scotland, 2005.

Working Papers

In Submission:

Thomas J. Walsh, Michael L. Littman, Alexander Borgida. "Learning Web-Service task Descriptions from Traces" (submitted to *Web Intelligence and Agent Systems: An International Journal*)

Work Experience

- *Post-Doctoral Researcher*, University of Arizona, Tucson, AZ (6/2010 – present)
 - **Learning Verb Meanings** – Helped design and build a system that learned to enact commands based on human demonstrations of "verb actions".
 - **Human/Teacher Interfaces** – Helped develop a teaching interface for human teachers to interact with an electronic student flying a simulated UAV.
 - **Teaching Skills**- Developed a framework and learning agent for tutoring humans in high-level skills in math and language problems.
- *Research Assistant*, Rutgers University, New Brunswick, NJ (8/2004 – 5/2010)
 - Performed theoretical and experimental analyses of learning algorithms for several relational action models.
 - Designed a system for extracting relational models from traces of users interacting with web-services.
 - Implemented several reinforcement learning algorithms from the academic literature on a Sony Aibo robot so that it could learn to complete a series of real-world tasks.
- *Intern, Siemens Corporate Research*, Princeton, NJ (5/2008 – 8/2008)
 - Worked with researchers at SCR to develop web-site monitoring software using advanced text-classification techniques.
- *Teaching Assistant*, Rutgers University, New Brunswick, NJ (8/2003 – 5/2005)
 - Taught recitation sessions twice a week for "Principles of Programming Languages"
- *Intern, Applied Signal Technology*, Annapolis Junction, Md. (5/2004 – 12/2004)
 - With a team of industry researchers, analyzed various text mining methods for their applicability in stream processing.
- *Guest Researcher*, National Institute of Standards and Technology (NIST), Gaithersburg, Md. (5/2003 – 8/2003).
 - Co-authored the NIST Special Publication on Voice Over IP Security

Awards

- Co-winner (with Lihong Li) of the *ICML 2008 Best Student Paper Award* for "Knows What it Knows: A

Thomas J. Walsh

- Framework for Self-Aware Learning” (Li, Littman, Walsh).
- Co-Winner (with Ali Nouri and Lihong Li) of the First Annual Reinforcement Learning Competition Pentathlon (2006).

PC / Referee / Organizational Experience:

- Senior Program Committee member for International Joint Conference on Artificial Intelligence (IJCAI-2011)
- Domain contributor and tester for the 2011 International Probabilistic Planning Competition (IPPC-2011)
- Reviewer for
 - *Artificial Intelligence Journal*
 - *Journal of Machine Learning Research*
 - *IEEE Transactions on Automatic Control*
 - *IEEE Conference on Decision and Control*
 - *Information and Computation*
 - *Neural Information Processing Systems (NIPS) Conference*
 - Various workshops.