Thomas J. Walsh

Email: thomasjwalsh [at] gmail.com
Web: http://www.thomasjwalsh.net

Research Summary

My research career has crossed a gamut of domains, from robotics to educational assessment, and now to leading a team of industry researchers in workforce management. I have worked in AI, machine learning, reinforcement learning, data science, and other quantitative fields. But the common element in all of these endeavors is that data tells a story, and clever algorithms can find ways to explore, model, and act on that data.

My dissertation research was in *reinforcement learning (RL)*, a subsection of machine learning where an agent or robot must actively learn a model of its environment while also completing tasks. There, my theoretical and empirical work answered open questions about how to marry traditional AI representations with modern RL and still maintain tractable solutions. In my post-doctoral work, I focused on problems where learning is done through observations of people or another agent, and used these techniques to train systems in robotics and electronic tutoring. Currently, I am the Senior Director of Data Science at a 5,000+ company, leading a team of researchers to find and deliver cutting edge AI, ML, and parallel computing techniques for workforce management problems.

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:

Robert C. Grande, Thomas J. Walsh, Girish Chowdhary, Sarah Ferguson, and Jonathan P. How. "Online Regression for Data with Changepoints using Gaussian Processes and Reusable Models". *IEEE Transactions on Neural Networks and Learning Systems*, 2016.

Mark Cutler, Thomas J. Walsh, and Jonathan P. How. "Real-World Reinforcement Learning via Multi-Fidelity Simulators". Submitted to *IEEE Transactions on Robotics (TRO)*, 2015.

Bernard Michini, Thomas J. Walsh, Ali-akbar Agha-mohammadi, and Jonathan P. How. "Bayesian Nonparametric Reward Learning from Demonstration". *IEEE Transactions on Robotics (TRO)*, 2015.

Girish Chowdhary, Miao Liu, Robert C. Grande, Thomas J. Walsh, Jonathan P. How and Lawerence Cairn. "Off-Policy Reinforcement Learning with Gaussian Processes". *Acta Automatica Sinica, 2014.*

Alborz Geramifard, Thomas J. Walsh, Stefanie Tellex, Girish Chowdhary, Nicholas Roy, Jonathan P. How. "A Tutorial on Linear Function Approximators for Dynamic Programming and Reinforcement Learning". *Accepted in Foundations and Trends in Machine Learning.*

Thomas J. Walsh, Michael L. Littman, Alexander Borgida. "Learning Web-Service Task Descriptions from Traces". *Web Intelligence and Agent Systems: An International Journal*, 10(4):397-421, 2012.

Lihong Li, Michael L. Littman, Thomas J. Walsh, Alexander L. Strehl. **"**Knows what it Knows: A Framework For Self-Aware Learning**".** *Machine Learning*, 82(3):399-443, 2011.

Fusun Yaman, Thomas J. Walsh, Michael L. Littman, Marie desJardins. "Democratic Approximation of Lexicographic Preference Models". *Artificial Intelligence Journal*, Volume 175, Pages 1290-1307, May 2011.

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:

Robert C. Grande, Thomas J. Walsh, Jonathan P. How. "Sample Efficient Reinforcement Learning with Gaussian Processes" *In Proceedings of the International Conference on Machine Learning (ICML-14)*, Beijing, China, 2014.

Mark Cutler, Thomas J. Walsh, and Jonathan P. How. "Reinforcement Learning with Multiple-Fidelity Simulators". *IEEE International Conference on Robotics and Automation (ICRA),* 2014.

Alborz Geramifard, Thomas J. Walsh, Nicholas Roy, and Jonathan P. How. "Batch-iFDD for Representation Expansion in Large MDPs". *In Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI-13),* Bellevue, WA, 2013.

Thomas J. Walsh and Sergiu Goschin. "Dynamic Teaching in Sequential Decision Making Environments". *In Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI-12),* Catalina, CA, 2012.

Thomas J. Walsh, Daniel Hewlett and Clayton T. Morrison. "Blending Autonomous Exploration and Apprenticeship Learning". *In Proceedings of the Twenty-Fifth Annual Conference on Neural Information Processing Systems (NIPS-11),* Granada, Spain, 2011.
(Also appeared at the 2011 RSS Workshop on the State of Imitation Learning)

Derek T. Green, Thomas J. Walsh, Paul R. Cohen and Yu-Han Chang."Learning a Skill-Teaching Curriculum with Dynamic Bayes Nets". *In Proceedings of the Twenty-Third Conference on Innovative Applications of Artificial Intelligence (IAAI-11),* San Francisco, CA, 2011.

Daniel Hewlett, Thomas J. Walsh, and Paul R. Cohen. "Teaching and Executing Verb Phrases". *In Proceedings of the First Joint IEEE International Conference on Development and Learning and on Epigenetic Robotics (ICDL-Epirob-11),* Frankfurt, Germany, 2011.
(Also appeared at the 2011 RSS Workshop on the State of Imitation Learning)

Raquel Torres Peralta, Tasneem Kaochar, Ian R. Fasel, Clayton T. Morrison, Thomas J. Walsh, Paul R. Cohen. "Challenges to Decoding the Intention Behind Natural Instruction". *In Proceedings of the IEEE International Symposium on Robots and Human Interactive Communications (RO-MAN-2011),* Atlanta, GA, 2011.

Derek T. Green, Thomas J. Walsh, Paul R. Cohen, Carole R. Beal and Yu-han Chang. "Gender Differences and the Value of Choice in Intelligent Tutoring Systems*". In Proceedings of User Modeling, Adaptation and Personalization (UMAP-2011),* Girona, Spain, 2011.

Tasneem Kaochar, Raquel Torres Peralta, Ian R. Fasel, Clayton T. Morrison, Thomas J. Walsh, Paul R. Cohen.

"Towards Understanding How Humans Teach Robots". *In Proceedings of User Modeling, Adaptation and Personalization (UMAP-2011)*, Girona, Spain, 2011.

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, 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". I*n 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". I*n 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), 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 and Presentations:
Girish Chowdhary, Miao Liu, Robert C. Grande, Thomas J. Walsh, Jonathan P. How. "Off-Policy Reinforcement Learning with Gaussian Processes". *Multidisciplinary Conference on Reinforcement Learning and Decision Making*, Princeton, NJ, 2013.

Zhao, F., Walsh, T. J., & Broaddus, A. E. "Item Selection Techniques in Computer Adaptive Testing Using Bayesian Networks." Poster presented at the 2013 *Annual Convention of the American Educational Research Association (AERA) SIG Business Meeting,* San Francisco, CA, 2013.

Thomas J. Walsh, Javad Taheri, Jeremy B. Wright and Paul R. Cohen. "Leadership Games and their Application in Super-Peer Networks". *AAAI Workshop on Applied Adversarial Reasoning and Risk Modeling*, San Francisco, CA, 2011.

Daniel Hewlett, Wesley Kerr, Thomas J. Walsh and Paul Cohen. "A Framework for Recognizing and Executing Verb Phrases", *RSS Workshop on Grounding Human-Robot Dialog for Spatial Tasks, 2011.*

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.

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.

Industry Articles and White Papers:

Thomas J. Walsh [anonymously] "Workforce Management in the Age of AI". *Kronos White Paper, 2018.*

Thomas J. Walsh "Clustering with Workforce Auditor". *Kronos White Paper, 2015.*

Patents
Michael A. Scarpati, Thomas J. Walsh, Gregory Lamarre Anderson. "Predicting Upcoming Missing Clockings and Alerting Workers or Managers", provisionally filed April 18, 2018.

Work Experience

- ***Senior Director of Data Science,* Kronos Incorporated**, Lowell, MA (1/2014-present)
    - *Previously Director of Data Science and Senior Data Scientist*
    - *Kronos workforce management software is used by over 35,000 companies in 100+ countries*
    - *Built a data science program from the ground up, expanding year over year and making and AI and Machine Learning core components of corporate strategy*
    - *Leading an advanced research group of 10+ data scientists and engineers applying cutting*

edge machine learning and Big Data techniques to workforce management data.
    - ***Workforce Auditor*** *– Led the research, design, and implementation of an AI-based fraud detection system.*
        - *Winner of the Brandon Hall Bronze Award for Best Advance in HR Predictive Analytics Technology, 2015.*
    - ***ML Volume Forecasting*** *– Led the research effort to design and deploy a machine learning-based sales forecasting solution now available to hundreds of major national and international retailers.*
    - *Led various research initiatives applying machine learning techniques to workforce business problems, including retention, scheduling, and capacity planning.*
    - *Worked with product management, architecture, and developers to turn prototypes real AI into products*
- **P*ost-Doctoral Researcher***, **Massachusetts Institute of Technology (MIT)**, Cambridge, MA (2/2013-1/2014 with *Prof. Jonathan How* in the ACL lab at LIDS)
    - **Reinforcement Learning (RL) for Physical Systems** – Applied reinforcement learning techniques to real-world control system using Bayesian non-parametric methods and data from simulation.
- ***Research Associate*****, Center for Educational Testing and Evaluation, University of Kansas**, Lawrence, KS (9/2011-2/2013 with *Dr. Neal Kingston*)
    - **Educational Assessment –** Led CS and educational research and designed the software architecture for performing alternate assessments of children using a "Learning Maps".
- ***Senior Specialist,*** **Accenture Technology Labs**, Chicago, IL (*6/2011 – 9/2011*)
    - **Health Care Data Mining** – Designed algorithms for analyzing health care data.
- ***Post-Doctoral Researcher,*** **University of Arizona**, Tucson, AZ (*6/2010 – 6/2011* with *Prof. Paul Cohen*)
    - **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 an 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 math and language problems.
- ***Research Assistant,*** **Rutgers University**, New Brunswick, NJ *(8/2004 – 5/2010* with Prof. *Michael Littman)*
    - **Autonomous Learning -** Performed theoretical and experimental analyses of learning algorithms for several relational action models.
    - **Learning from Demonstration -** Designed a system for extracting relational models from traces of users interacting with web-services.
    - **Robot Learning -** Implemented several reinforcement learning algorithms on various robotics platforms, including a Sony Aibo.
- ***Intern, Siemens Corporate Research***, Princeton, NJ *(5/ 2008 – 8/ 2008)*
    - Worked with researchers at SCR to develop text-classification techniques.
- ***Teaching Assistant,*** **Rutgers University**, New Brunswick, NJ *(8/2003 – 5/2005)*
- ***Intern, Applied Signal Technology****,* Annapolis Junction, Md. *(5/2004 –12/2004)*
    - Analyzed various text mining methods with a team of industry researchers.
- ***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

Invited Talks

- *"New Tricks for Old Representation": Efficient Reinforcement Learning with Relation Models*
  - Given at Carnegie Mellon University
  - University of Pittsburgh
  - MIT Lincoln Labs
- *Lessons on Learning with Compact Domain Models*, 2011 (University of Kentucky, Accenture Technology Labs, University of Kansas)
- *Compact Models of Skill Teaching and Other Sequential Tasks*, 2011 (University of Kansas)
- *Practical Reinforcement Learning: Models, Algorithms, and Teachers*, 2013 (MIT)
- *Solving Your Big Data Conundrum with Workforce Analytics Plug-ins, 2015* (National Retail Federation Annual Conference)
- *Lessons Learned from Delivering Workforce AI,* 2018 (Hellman & Friedman LLC Data Science Summit)

Awards

- Named a Distinguished Program Committee member for IJCAI 2018
- High Scoring (top 218) reviewer for NeurIPS 2018, receiving a coveted free conference registration
- High Scoring reviewer for NeurIPS 2019, receiving a coveted free conference registration
- Co-winner (with Lihong Li) of the *ICML 2008 Best Student Paper Award* for "Knows What it Knows: A 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:

- Co-editor of an upcoming book *"AI for Human Resources"* through the Society of Industrial Organizational Psychology (SIOP) professional practice series
- Co-chair, NIPS 2013 Workshop on "*Advances in Machine Learning for Sensorimotor Control*"
- Organizing Committee, *"Special Issue on Aerospace and Mechanical Applications of Reinforcement Learning and Adaptive Learning Based Control"* for the Journal of Aerospace Information Systems (JAIS)
- Organizing Committee, 2013 CDC Workshop on "*Intelligent Planning and Control*"
- Program Committee, 2013 UAI Workshop on "*Models for Spatial, Temporal, and Network Data*"
- Organizing Committee, AAAI 2013 Spring Symposium on "*Lifelong Machine Learning*"
- Domain contributor and tester, 2011 *International Probabilistic Planning Competition*
- Senior Program Committee, *International Joint Conference on Artificial* Intelligence (IJCAI) [2011,2016]
- Associate Editor, *IEEE Transactions on Big Data*
- Reviewer for
  - *Artificial Intelligence Journal*
  - *Journal of Machine Learning Research*
  - *Journal of Artificial Intelligence Research*
  - *Machine Learning Journal*
  - *IEEE Transactions on Automatic Control*
  - *IEEE Transactions on Neural Networks and Learning Systems*
  - *IEEE Robotics and Automation Letters*
  - *Robotics and Autonomous Systems*
  - *Journal of Aerospace Information Systems*
  - *Information and Computation*
  - *Artificial Intelligence in Medicine*
  - *International Joint Conference on Artificial Intelligence (IJCAI)*
  - *Conference of the Association for the Advancement of Artificial Intelligence(AAAI)*
  - *Neural Information Processing Systems (NeurIPS)*

- *International Conference on Machine Learning (ICML)*
- *International Conference on Automated Planning and Scheduling (ICAPS)*
- *International Conference on Algorithmic Learning Theory (ALT)*
- *International Conference on Learning Representations (ICLR)*
- *IEEE Conference on Decision and Control (CDC)*
- *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*
- *International Conference on Robotics and Automation (ICRA)*
- *International Conference on Control, Automation, Robotics and Vision (*ICARCV*)*
- *IEEE Conference on Big Data*
- *International Conference on Development and Learning (ICDL-EPIROB)*
- *Asian Conference on Machine Learning (ACML)*
- *European Workshop on Reinforcement Learning (EWRL)*
- Various other conferences and workshops.

Technical Skills

Programming Languages: Python, Java, R, Scala, C, C++, MATLAB, LISP, Prolog.

AI Paradigms: **Reinforcement learning** (model-based, model-free, planning, exploration/exploitation, relational RL), **Active Learning** (Bandit Algorithms), **Regression and Classification** (random forests, GBDT, SVMs, Neural Nets), **Graphical Models** (Bayes Nets, Hidden Markov Models), **Unsupervised learning** (Clustering, dimensionality reduction)

Computing Frameworks: Apache Spark, Pandas, Cassandra (NoSQL), XGBoost, Jupyter Notebooks, ROS, scikit-learn, many others…