

Bo Liu

Assitant Professor, Auburn University
URL: www.eng.auburn.edu/~bzl0056/

Phone: 551-333-9638
Email: boliu@auburn.edu

EDUCATION

2010 -2015 **Ph.D.** Computer Science, University of Massachusetts Amherst
Advisor: Sridhar Mahadevan
2008 -2010 **M.S.** Computer Engineering, Stevens Institute of Technology
2005 -2008 **M.S.** Control Engineering, University of Science and Technology of China

INDUSTRY EXPERIENCE

07/2016 - present	Assistant Professor	Auburn University, Auburn, AL
09/2015 - 04/2016	Research Staff Member	Philips Research, Cambridge, MA
06/2015 - 08/2015	Research Scientist Intern	Amazon Research, Seattle, WA
06/2013 - 09/2013	Machine Learning Scientist Intern	Amazon Research, Seattle, WA
06/2011 - 09/2011	Applied Research Intern	eBay Search Science, San Jose, CA

RESEARCH INTEREST

Interactive Machine Learning

Reinforcement learning, Active learning, Online learning, and multi-armed bandit with application to online recommendation and spoken dialog system

Deep Learning

Deep Neural Networks theory, design, and applications

Large-Scale Efficient Optimization

Online/stochastic optimization with application to time-series problems including demand forecasting and transaction risk management

Transfer Learning

Multi-task learning and domain adaptation with application to multi-lingual machine translation and social network analysis

SELECTED PROJECTS

Proximal Gradient Reinforcement Learning, 01/2013 - 05/2015 [UAI-2015, IJCAI-2016]

- **Goal:** For 30 years, researchers in reinforcement learning have been attempting to design a true stochastic gradient temporal difference learning method. Another long-term attempt is to provide sample complexity analysis of a temporal difference learning algorithm.
- This work is the first work that establishes a first-order stochastic optimization framework for temporal difference learning, which enables acceleration, regularization, and sample complexity analysis.
- This work received high recognition from Prof. Richard Sutton as “the best attempts to make TD methods with the robust convergence properties of stochastic gradient descent.”

Sparse Learning Models, 01/2014 - 03/2016 [UAI-2016, AAAI-2016]

- **Goal:** This project explores the improvement of the learning ability of several notable sparse supervised learning models, including Lasso and Dantzig Selector.

- Dantzig Selector with an Approximately Optimal Denoising Matrix: Dantzig Selector is notable for feature selection and sparse signal recovery. Is it possible to improve the sparse signal recovery ability of the vanilla Dantzig Selector with very little extra effort?
- Uncorrelated Group Lasso: Group Lasso captures "sparsity among groups", how to capture "sparsity inside each group"?

Transfer Learning, Domain Adaptation and Multi-task Learning with Sparsity and Geometric Structure, 01/2012 - 01/2014 [UMCS-2012]

- **Goal:** Explore two types of intrinsic structure of data: sparsity/low-rank structure and manifold geometry in transfer learning, domain adaptation, and multi-task learning.
- Sparse Manifold Alignment: Aim to reach a better trade-off between preserving cross-domain similarity and uniqueness among different tasks.
- Manifold learning is used to help preserve the cross-domain latent intrinsic structure, and sparsity is introduced to help prune out domain-specific features. The algorithm is friendly to MapReduce implementation. The work is applied to multi-lingual machine translation, image alignment, social network analysis, etc.

Sparse Reinforcement Learning, 09/2010 - 09/2014 [UAI-2016, UAI-2012, NIPS-2012, NIPS-2010]

- **Goal:** How can modern optimization help design regularized reinforcement learning algorithms?
- SparseQ: we use stochastic variational inequality to propose the first sparse Q-learning algorithm
- ROTD: Dual norm representation is applied to enable regularized off-policy TD learning
- ODDS-TD: Dantzig Selector with an approximate optimal denoising matrix is applied to improve DS-TD, and the performance is better than previous DS-TD, BPDN-TD for sparse reinforcement learning

Winner-take-all modeling and recurrent neural networks, 09/2009 - 03/2010 [IEEE TNNLS, Neurocomputing, etc.]

- **Goal:** The winner-take-all (WTA) competition is widely observed in both inanimate and biological media and society. This work aims to present a simple, explicit and interpretable WTA model with recurrent neural networks (RNN), along with other applications of RNN
- A recurrent neural network model to generate the WTA competition
- Decentralized kinematic control using RNN

SELECTED PUBLICATIONS

Google Scholar: https://scholar.google.com/citations?hl=en&user=8MliTo4AAAAJ&view_op=list_works&sortby=pubdate

Stats: H-index: 11. Conference paper: 10, Journal paper: 10.

Topics: **TL:** transfer learning, **DL:** deep learning, **RL:** reinforcement learning, **SL:** sparse learning.

Conference Publications

C1 **Bo Liu**, Ji Liu, Luwan Zhang. Dantzig Selector with an Approximately Optimal Denoising Matrix. *Proceedings of the Conference on Uncertainty in AI (UAI)*, 2016. (SL,RL)
Acceptance rate: 31%(85/275).

- C2 **Bo Liu**, Ji Liu, Mohammad Ghavamzadeh, Sridhar Mahadevan, Marek Petrik. A Proximal Gradient Framework for Robust TD Learning. *International Joint Conferences on Artificial Intelligence (IJCAI)*, 2016. (TL,RL)
Acceptance rate: 24%(551/2294).
- C3 Deguang Kong, Ji Liu, **Bo Liu**, Xuan Bao. Uncorrelated Group Lasso. *Association for the Advancement of Artificial Intelligence (AAAI)*, 2016. (TL,SL)
Acceptance rate: 26%(549/2132).
- C4 **Bo Liu**, Ji Liu, Mohammad Ghavamzadeh, Sridhar Mahadevan, Marek Petrik. Finite-Sample Analysis of Proximal Gradient Algorithms. *Proceedings of the Conference on Uncertainty in AI (UAI)*, 2015, **Facebook Best Student Paper Award**. (TL,RL)
Acceptance rate: 1%(3/291) (Best paper award rate).
- C5 Ian Gemp, Sridhar Mahadevan, **Bo Liu**. Solving Large-Scale Sustainable Supply Chain Networks using Variational Inequalities, *AAAI Workshop on Computational Sustainability*, Austin, Texas, 2015.
- C6 **Bo Liu**, Sridhar Mahadevan, Ji Liu. Regularized Off-Policy TD-Learning. *26th Annual Conference on Neural Information Processing Systems (NIPS)*, Lake Tahoe, Nevada , 2012, December 3-6, **spotlight presentation**. (RL, SL)
Acceptance rate: 5%(72/1467) (Oral presentation rate).
- C7 Sridhar Mahadevan, **Bo Liu**. Sparse Q-learning with Mirror Descent. *Proceedings of the Conference on Uncertainty in AI (UAI)*, 2012. (RL, SL)
Acceptance rate: 31%(95/304).
- C8 Sridhar Mahadevan, **Bo Liu**. Basis Construction from Power Series Expansions of Value Functions. *24th Annual Conference on Neural Information Processing Systems (NIPS)*, Vancouver, B.C., Canada, 2010, December 6-8. (RL)
Acceptance rate: 24%(293/1219).
- C9 Haibo He, **Bo Liu**. A Hierarchical Learning Architecture with Multiple-Goal Representations Based on Adaptive Dynamic Programming *IEEE International Conference on Networking, Sensing and Control (ICNSC'10)*, Chicago, 2010. (RL)
- C10 **Bo Liu**, Haibo He, Daniel.Repperger. Two-Time-Scale Online Actor-Critic Paradigm Driven by POMDP. *IEEE International Conference on Networking, Sensing and Control (ICNSC'10)*, Chicago, 2010. (RL)

Journal Publications

- J1 Shuai Li, Yuesheng Lou, **Bo Liu**. Bluetooth aided mobile phone localization: a nonlinear neural circuit approach. *ACM Transactions on Embedded Computing Systems (ACM TECS)*, 2014.
Impact factor: 1.190.
- J2 Shuai Li, **Bo Liu**, Yangming Li. Selective Positive-negative Feedback Produces the Winner-take-all Competition in Recurrent Neural Networks. *IEEE Transactions on Neural Networks and Learning Systems (IEEE TNNLS)*, 2013. (DL)
Impact factor: 4.854.
- J3 Shuai Li, Sanfeng Chen, **Bo Liu**, Yangming Li, Yongsheng Liang Decentralized Kinematic Control of A Class of Collaborative Redundant Manipulators via Recurrent Neural Networks, *Neurocomputing*, 2012. (DL)
Impact factor: 2471.

- J4 Shuai Li, Sanfeng Chen, **Bo Liu**. Accelerating a Recurrent Neural Network to Finite-Time Convergence for Solving Time-Varying Sylvester Equation by Using a Sign-Bi-power Activation Function, *Neural Processing Letters*, pp.1-17, 2012 (DL)
Impact factor: 1.747.
- J5 Shuai Li, **Bo Liu**, Baogang Chen and Yuesheng Lou. Neural Network Based Mobile Phone Localization Using Bluetooth Connectivity. *Neural Computing and Applications*, 2012.
Impact factor: 1.569.
- J6 Shuai Li, Hongzhu Cui, Yangming Li, **Bo Liu** and Yuesheng Lou. Decentralized Control of Collaborative Redundant Manipulators with Partial Command Coverage via Locally Connected Recurrent Neural Networks. *Neural Computing and Applications*, 2012.
Impact factor: 1.569.
- J7 Shuai Li, Yunpeng Wang, Jiguo Yu, **Bo Liu**. A Nonlinear Model to Generate the Winner-take-all Competition. *Communications in Nonlinear Science and Numerical Simulation*, 2012.
Impact factor: 2.866.
- J8 **Bo Liu**, Sanfeng Chen, Shuai Li, Yongsheng Liang Intelligent control of a sensor-actuator system via kernelized least-squares policy iteration. *Sensors* 12 (3), 2632-2653, 2012.
Impact factor: 2.437.
- J9 Sanfeng Chen, Shuai Li, **Bo Liu**, Yuesheng Lou and Yongsheng Liang, Self-Learning Variable Structure Control for a Class of Sensor-Actuator Systems, *Sensors*, Vol.12, pp.6117-6128, 2012.
Impact factor: 2.437.
- J10 **Bo Liu**, Haibo He, Sheng Chen. Adaptive Dual Network Design for a Class of SIMO Systems with Nonlinear Time-variant Uncertainties. *Acta Automatica Sinica*, Vol.36, pp.564-572, 2010.
Impact factor: 1.290.

Under Review

- U1 **Bo Liu**, Sridhar Mahadevan, Ian Gemp. Regularized Reinforcement Learning Using Dual Averaging. under review. (RL,SL)
- U2 Chang Wang, **Bo Liu**, Hoa Vu, Sridhar Mahadevan. Sparse Manifold Alignment. Tech. Report UM-CS-2012-030. (TL,SL)
- U3 **Bo Liu**, Ji Liu, Sridhar Mahadevan, Deguang Kong. Minimum-Volume Multi-Task Learning. under review. (TL)
- U4 **Bo Liu**, Sridhar Mahadevan. Compressive Reinforcement Learning with Oblique Random Projections. Tech. Report UM-CS-2011-024. (RL, SL)

INVITED TALKS

- Gradient, Semi-gradient and Pseudo-gradient Reinforcement Learning, SIAM Conference on Optimization, Vancouver, July 23th, 2017
- Proximal Reinforcement Learning. SUNY Buffalo, Buffalo, NY, 2016
- Proximal Reinforcement Learning. University of Wyoming, Laramie, WY, 2016
- Proximal Reinforcement Learning. Portland State University, Portland, OR, 2016
- Proximal Reinforcement Learning. Auburn Univeristy, AL, 2016
- Proximal Reinforcement Learning. Washington State University, Pullman, WA, 2015
- Efficient Transfer Decision-making. Amazon Research, Seattle, WA, 2015
- Efficient Transfer Decision-making. Philips Research, Briarcliff, NY, 2015
- Sequential Decision Making Addressing Big Data. Washington State University, Pullman, WA, 2013
- First-Order Sparse Reinforcement Learning. Adobe Research, CA, 2012

ACADEMIC SERVICES

Journal Review:

Machine Learning (MLJ), IEEE Transactions on Neural Networks and Learning Systems (TNNLS), IEEE Signal Processing Letters, IEEE Transactions on Circuits and Systems for Video Technology (TCSVT), IEEE Transactions on Transactions on Autonomous Mental Development (TAMD), IET Image Processing, Neurocomputing, Knowledge-Based Systems, Applied Soft Computing, Artificial Intelligence Review, Evolving Systems, Expert Systems with Applications, Journal of Classification, Pattern Analysis and Applications

Conference Review:

UbiComp 2015. AAAI 2016, 2017. IJCAI 2015, 2016. NIPS 2013, 2014, 2015, 2016. AISTATS 2017.

Program Committee:

AAAI 2017. IJCAI 2015, 2016.

HONORS AND AWARDS

2017 Tencent Rhino-Bird Award, Tencent AI Lab
2015 Facebook Best Student Paper Award of UAI-2015
2013 Google Fellowship Nomination
2012 NIPS Student Award
2012 UAI Student Scholarship
2010 ICNSC Doctoral Consortium Grant
2009 ICCNS Student Fellowship

TECHNICAL SKILLS

Languages: Java, Python, C, Perl, Lua

Tools: Tensorflow, Theano, Caffe, PyTorch, R, Vowpal Wabbit, Spark, Hadoop, SQL

Platforms: Mac OSX, Unix/Linux, Windows