

CURRICULUM VITAE

Le Song

Associate Professor

Computational Science and Engineering

College of Computing

Associate Director

Center for Machine Learning

Georgia Institute of Technology

1340 Klaus Building

266 Ferst Drive

Atlanta, GA 30332, USA

Email: lsong@cc.gatech.edu

Homepage: <http://www.cc.gatech.edu/~lsong>

CONTENTS

1	Earned Degrees	3
2	Employment History	3
3	Honor and Awards	3
4	Research, Scholarship, and Creative Activities	4
4.1	Summary	4
4.2	Submitted Articles	5
4.3	Published and Accepted Journal Articles	5
4.4	Conference Presentations with Proceedings (Refereed)	6
5	Keynotes and Invited Talks	12
6	Grants and Contracts	15
6.1	As Principal Investigator	15
6.2	As Co-Principal Investigator	15
6.3	As Senior Personnel or Contributor	16
6.4	Other Professional Activities	16
7	Teaching	16
7.1	Courses Developed	16
7.2	Courses Taught	17
7.3	Individual Student Guidance	17
8	Service	19
8.1	Professional Contributions	19
8.2	Public and Community Service	20
8.3	Institute Contributions	22

EARNED DEGREES

1. 2004 – 2008 *Ph.D. in Computer Science*
University of Sydney and National ICT Australia, Australia
Advisor: Alex Smola
2. 2002 – 2004 *M.Sc. in Computer Science*
University of Sydney, Australia
Advisor: Masahiro Takatsuka and Peter Eades
3. 1998 – 2002 *B.S. in Computer Science*
South China University of Technology, China

EMPLOYMENT HISTORY

1. 2017.2 – Present *Associate Professor, CSE, College of Computing, Georgia Institute of Technology*
Developing machine learning methods and algorithms for complex and dynamic data
2. 2016.9 – Present *Associate Director, Center for Machine Learning, Georgia Institute of Technology*
Directing machine learning and cross campus multi-disciplinary research involving 90 faculty
3. 2011.9 – 2017.2 *Assistant Professor, CSE, College of Computing, Georgia Institute of Technology*
Developing machine learning methods and algorithms for complex and dynamic data
4. 2011.3 – 2011.8 *Research Scientist, Google Research*
Developing large scale machine learning package for Internet data
5. 2008.7 – 2011.3 *Postdoc Research Fellow, School of Computer Science, Carnegie Mellon University*
Developing nonparametric probabilistic graphical models for complex social and biological data
6. 2008.1 – 2008.6 *Visiting Researcher, Statistical Machine Learning, National ICT Australia, Australia*
Analyzing sensor time series data using kernel methods

HONOR AND AWARDS

1. 2016 **Best Paper Award**, ACM Conference on Recommendation System (Recsys) Workshop on Deep Learning for Recommendation System
2. 2016 **Best Student Paper Award**, Artificial Intelligence and Statistics (AISTATS)
3. 2015 **Best Paper Award**, IEEE International Parallel & Distributed Processing Symposium (IPDPS)
4. 2014 National Science Foundation **CAREER Award**
5. 2014 Outstanding Junior Faculty Research Award, College of Computing, Georgia Institute of Technology
6. 2014 Lockheed Martin Inspirational Young Faculty Award, College of Computing, Georgia Institute of Technology

7. 2013 **Best Paper Award**, Neural Information Processing Systems (NIPS)
8. 2013 Raytheon Faculty Fellowship, College of Computing, Georgia Institute of Technology
9. 2016 **Best Paper Award**, NIPS Workshop on Frontiers of Network Analysis: Methods, Models and Applications
10. 2010 **Best Paper Award**, International Conference on Machine Learning (ICML)
11. 2008 – 2011 Lane Fellowship, School of Computer Science, Carnegie Mellon University
12. 2004 First Student Prize, IEEE Information Visualization Contest
13. 2001 Meritorious, SIAM International Mathematical Contest in Modeling for Undergraduates
14. 2001 First Prize, Chinese National Mathematical Contest in Modeling for Undergraduates

RESEARCH, SCHOLARSHIP, AND CREATIVE ACTIVITIES

Summary

My principal research interests lie in the development of efficient algorithms and intelligent systems which can learn from a massive volume of complex (high dimensional, nonlinear, multi-modal, skewed, and structured) data arising from both artificial and natural systems, reveal trends and patterns too subtle for humans to detect, and automate decision making processes in uncertain and dynamic possible world. I develop core machine learning methodology, including kernel methods, deep learning, graphical models, probabilistic and stochastic modeling, scalable learning algorithms.

I am also interested in developing machine learning models and algorithms to address interdisciplinary problems. For instance, I've conducted research on the management of information diffusion networks and recommendation systems, the discovery of time-varying gene regulatory networks, the understanding of disease progression, the extraction of topics based on online document feeds, the prediction of materials properties, and the predictive modeling of robotic systems.

My research work center around four major themes:

- Embedding: A Learning Framework for Complex Distributions, Structures and Dynamics
- Dynamic Processes over Networks: Representation, Modeling, Learning, Inference, Optimization and Control
- Large Scale Machine Learning: Efficient Algorithms, Distributed Learning, and High Performance Computing Techniques
- Interdisciplinary Problems: Social Network Analysis, Healthcare Analytics, Computational Biology and Neuroscience, Materials Science

Asterisks indicate work done at Georgia Tech and the names of student co-authors are in boldface.

Submitted Articles

1. * **B. Dai**, R. Guo, S. Kumar, and L. Song (2016). Stochastic Generative Hashing. *submitted*.
2. * **H. Dai**, **Y. Wang**, **R. Trivedi** and L. Song (2016). Recurrent coevolutionary feature embedding processes for recommendation. *submitted*.
3. * **E. Choi**, **M. Bahadori**, L. Song, W. Stewart and J. Sun (2016). GRAM: graph-based attention model for healthcare representation learning. *submitted*.
4. * **Y. Chen**, E. Theodorou and L. Song (2016). Stochastic control of information diffusion processes. *submitted*.
5. * **S. Li**, Y. Xie, **H. Dai** and L. Song (2016). M-statistics for kernel change point detection. *submitted to Journal of Machine Learning Researches*.

Published and Accepted Journal Articles

1. * **M. Farajtabar**, **Y. Wang**, M. Rodriguez, **S. Li**, H. Zha and L. Song (2017). COEVOLVE: a joint point process model for information diffusion and network co-evolution. *Journal of Machine Learning Researches*. (accepted)
2. * L. Song, H. Liu, A. Parikh and E. Xing (2017). Nonparametric Latent Tree Graphical Models: Inference, Estimation, and Structure Learning. *Journal of Machine Learning Researches*. (accepted)
3. * **S. Li**, **M. Farajtabar**, **A. Verma**, Y. Xie and L. Song (2017). Detecting weak changes in dynamic events over networks. *IEEE Trans. on Signal Processing*. (accepted)
4. * **N. Du**, **Y. Liang**, M. Balcan, M. Rodriguez, H. Zha and L. Song (2016). Scalable influence maximization for multiple products. *Journal of Machine Learning Researches*. (in press)
5. * M. Rodriguez, L. Song, **N. Du**, H. Zha and B. Schölkopf (2016). Influence estimation and maximization in continuous-time diffusion networks. *ACM Transactions on Information Systems*. (in press)
6. * N. Daneshmand, M. Rodriguez, L. Song and B. Schölkopf (2015). Estimating diffusion network structure: recovery conditions, sample complexity, and a soft-thresholding algorithm. *Journal of Machine Learning Researches*. (in press)
7. * L. Song, K. Fukumizu, A. Gretton. (2013). Kernel embedding of conditional distributions. *IEEE Signal Processing Magazine*, 30(4), 98 – 111.
8. * K. Fukumizu, L. Song, A. Gretton. (2013). Kernel Bayes rule. *Journal of Machine Learning Researches*. 14:3385 – 3415.
9. * **B. Xie**, B. Jankovic, V. Bajic, L. Song and X. Gao (2013). Poly(A) motif prediction using spectral latent features from human DNA sequences. *Bioinformatics*. 29(13): i316-i325.
10. * L. Song, A. Smola, A. Getton, J. Bedo and K. Borgwardt. (2012). Feature selection via dependence maximization. *Journal of Machine Learning Researches*, 13:1393 – 1434.

11. E. Xing, W. Fu, and L. Song (2010). A state-space mixed membership blockmodel for dynamic network tomography. *Annals of Applied Statistics*, 4(2), pp. 535–566.
12. M. Kolar, L. Song, A. Ahmed, and E. Xing (2010). Estimating time-varying networks. *Annals of Applied Statistics*, 4(1), pp. 94–123.
13. N. Quadrianto, A. Smola, L. Song and T. Tuytelaars (2010). Kernelized sorting. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 32(10), pp. 1809–1821.
14. L. Song, M. Kolar, E. Xing. (2009). KELLER: estimating time-varying interactions between genes. *Bioinformatics*, 25(12), pp. i128–i136.
15. L. Williams, J. Gatt, S. Kuan, C. Dobson-Stone, D. Palmer, R. Paul, L. Song, P. Costa, P. Schofield and E. Gordon (2009). A polymorphism of the maoa gene is associated with emotional brain markers and personality traits on an antisocial index. *Nature Neuropsychopharmacology*, vol 34, pp. 1797–1809.
16. L. Song, J. Bedo, K. Borgwardt, A. Getton and A. Smola. (2007). Gene selection via the BAHSIC family of algorithms. *Bioinformatics* 23(13), i490–i498.
17. L. Williams, D. Palmer, B. Liddell, L. Song and E. Gordon. (2006). The ‘when’ and ‘where’ of perceiving signals of threat versus non-threat. *NeuroImage*, vol 31, pp. 458–467.
18. S.Q. Liu, and L. Song. (2005). Curvature relation of wave front and wave changing in external field. *Applied Mathematics and Mechanics*, 26(7), pp. 929–937.
19. S.Q. Liu, and L. Song. (2004). Numerical analysis of Lobster stomatogastric nervous system. *Acta Biophysica Sinica*, 20(3), pp. 217–224.

Conference Presentations with Proceedings (Refereed)

1. * **W. Liu**, Y. Wen, Z. Yu, B. Raj, L. Song (2016). SphereFace: deep hypersphere embedding for face recognition. *Computer Vision and Pattern Recognition (CVPR 2017)*.
2. * **H. Dai, B. Dai, Y. Zhang, S. Li** and L. Song (2016). Recurrent hidden semi-Markov model. *International Conference on Learning Representations (ICLR 2017)*.
3. * **B. Xie**, Y. Liang and L. Song (2017). Diversity neural network learns true target function. *Artificial Intelligence and Statistics (AISTATS 2017)*.
4. * **B. Dai, N. He, Y. Pan** and L. Song (2017). Learning from conditional distributions via dual embedding. *Artificial Intelligence and Statistics (AISTATS 2017)*.
5. * **Y. Chen**, X. Ye, H. Zha and L. Song (2017). Fokker-Planck inference machines: linking microscopic event history to macroscopic prediction. *Artificial Intelligence and Statistics (AISTATS 2017)*.
6. B. Tabibian, I. Velera, M. Farajtabar, L. Song, B. Scholkopf and M. Rodriguez (2017). Distilling Information Reliability and Source Trustworthiness from Digital Traces. *World Wide Web (WWW 2017)*.

7. ***H. Dai, Y. Wang, R. Trivedi** and L. Song (2016). Recurrent coevolutionary latent feature embedding for recommendation. ACM Conference on Recommendation System (Recsys) Workshop on Deep Learning for Recommendation System. (**Best Paper Award**)
8. ***Y. Wang, N. Du, R. Trivedi** and L. Song (2016). Coevolutionary latent feature processes for continuous-time user-item interactions. *Advances in Neural Information Processing Systems 26 (NIPS 2016)*.
9. ***M. Farajtabar, X. Ye, S. Harati, L. Song** and H. Zha (2016). Multistage campaigning in social networks. *Advances in Neural Information Processing Systems 26 (NIPS 2016)*.
10. * **B. Dai, N. He, H. Dai** and L. Song (2016). Provable Bayesian inference via particle mirror descent. *Artificial Intelligence and Statistics (AISTATS 2016)*. (**Best Student Paper Award**)
11. * Y. Nishiyama, **A. Afsharinejad, B. Boots** and L. Song (2016). Nonparametric kernel Bayes smoother. *Artificial Intelligence and Statistics (AISTATS 2016)*.
12. * **H. Dai, B. Dai** and L. Song (2016). Discriminative embedding of latent variables for structured data. *International Conference on Machine Learning (ICML 2016)*.
13. * **Y. Wang, B. Xie, N. Du** and L. Song (2016). Isotonic Hawkes processes. *International Conference on Machine Learning (ICML 2016)*.
14. * **N. Du, H. Dai, R. Trivedi, U. Upadhyay, M. Rodriguez** and L. Song (2016). Recurrent marked temporal point processes. *Knowledge Discovery and Data Mining (KDD 2016)*.
15. * M. Balcan, Y. Liang, L. Song, D. Woodruff and **B. Xie** (2016) Communication efficient distributed kernel principal component analysis. *Knowledge Discovery and Data Mining (KDD 2016)*. (**Plenary Long Talk**)
16. * **M. Karimi, E. Tavakoli, M. Farajtabar, L. Song** and M. Rodriguez (2016). Smart broadcasting: Do you want to be seen? *Knowledge Discovery and Data Mining (KDD 2016)*.
17. * **E. Khalil, B. Dilkina, L. Song, P. Le Bodic** and G. Nemhauser (2016). Learning to branch in mixed integer programming. *30th AAAI Conference on Artificial Intelligence (AAAI 2016)*.
18. * **M. Farajtabar, Y. Wang, M. Rodriguez, S. Li, H. Zha** and L. Song (2015). COEVOLVE: a joint point process model for information diffusion and network co-evolution. *Advances in Neural Information Processing Systems 25 (NIPS 2015)*. (**Plenary Long Talk**)
19. * **T. Long, M. Farajtabar, L. Song** and H. Zha (2015). NetCodec: community detection from individual activities. *SIAM International Conference on Data Mining (SDM 2015)*.
20. * **M. Farajtabar, N. Du, M. Zamani, M. Rodriguez** and L. Song (2015). Back to the past: source identification in diffusion networks. *Artificial Intelligence and Statistics (AISTATS 2015)*. (**Plenary Long Talk**)
21. * **N. Du, M. Farajtabar, A. Ahmed, A. Smola** and L. Song (2015). Dirichlet-Hawkes processes with applications to clustering continuous-time document streams. *Knowledge Discovery and Data Mining (KDD 2015)*.

22. * **Y. Liu, F. Li, S. Li**, L. Song and J. Rehg (2015). Efficient continuous-time hidden Markov model for disease modeling. *Advances in Neural Information Processing Systems 25 (NIPS 2015)*.
23. * **N. Du, Y. Wang, N. He** and L. Song (2015). Time-sensitive recommendation for recurrent activities. *Advances in Neural Information Processing Systems 25 (NIPS 2015)*.
24. * **E. Choi, N. Du, R. Chen**, L. Song and J. Sun (2015). Constructing disease network and temporal progression model via context-sensitive Hawkes processes. *International Conference on Data Mining (ICDM 2015)*.
25. * **S. Li, Y. Xie, H. Dai** and L. Song (2015). M-statistics for kernel change point detection. *Advances in Neural Information Processing Systems 25 (NIPS 2015)*.
26. * **B. Xie, Y. Liang** and L. Song (2015). Scale up nonlinear component analysis with doubly stochastic gradients. *Advances in Neural Information Processing Systems 25 (NIPS 2015)*.
27. * **Y. You**, J. Demmel, K. Czechowski, L. Song and R. Vuduc (2015). CA-SVM: communication-avoiding kernel support vector machines on clusters. *IEEE International Parallel & Distributed Processing Symposium (IPDPS 2015)*. (**Best Paper Award**)
28. * **Z. Yang, M. Moczulski, M. Denil**, N. de Freitas, A. Smola, L. Song and Z. Wang (2015). Deep fried convnets. *International Conference on Computer Vision (ICCV, 2015)*.
29. * **A. Shaban, M. Farajtabar, B. Xie**, L. Song and B. Boots (2015). Learning latent variable models via method of moments and exterior point optimization. *Uncertainty in Artificial Intelligence (UAI 2015)*.
30. * **Z. Yang**, A. Smola, L. Song and A. Wilson (2015). A la Carte – Learning fast kernels. *Artificial Intelligence and Statistics (AISTATS 2015)*. (**Plenary Long Talk**)
31. * **M. Farajtabar, N. Du**, M. Rodriguez, I. Valera, H. Zha and L. Song (2014). Social activity shaping by incentivizing users. *Advances in Neural Information Processing Systems 24 (NIPS 2014)*.
32. * **N. Du, Y. Liang**, M. Balcan and L. Song (2014). Learning time-varying coverage functions. *Advances in Neural Information Processing Systems 24 (NIPS 2014)*.
33. * N. Balcan, **C. Berling**, A. Blum, E. Cohen and **K. Patnaik** and L. Song (2014). Active learning and best-response dynamics. *Advances in Neural Information Processing Systems 24 (NIPS 2014)*.
34. * N. Daneshmand, M. Rodriguez, L. Song and B. Schölkopf (2014). Estimating diffusion network structure: recovery conditions, sample complexity, and a soft-thresholding algorithm. *31th International Conference on Machine Learning (ICML 2014)*. (**Invited to JMLR**)
35. * **N. Du, Y. Liang**, N. Balcan and L. Song (2014). Influence function learning in information diffusion networks. *31th International Conference on Machine Learning (ICML 2014)*.
36. * **E. Khalil**, B. Dilkina and L. Song (2014). Scalable diffusion-aware optimization of network topology. *Knowledge Discovery and Data Mining (KDD 2014)*.
37. M. Rodriguez, L. Song and B. Schölkopf (2014). Finding good cascade sampling processes for the network inference problem. *Computational Learning Theory (COLT 2014)*.

38. * **B. Dai, B. Xie, N. He, Y. Liang, A. Raj**, M. Balcan and L. Song (2014). Scalable kernel methods via doubly stochastic gradients. *Advances in Neural Information Processing Systems 24 (NIPS 2014)*.
39. * L. Song, A. Anandkumar, **B. Dai and B. Xie** (2014). Nonparametric estimation of multi-view latent variable models. *31th International Conference on Machine Learning (ICML 2014)*.
40. * **N. Du**, L. Song, H. Zha and M. Rodriguez (2013). Scalable influence maximization in continuous time diffusion networks. *Advances in Neural Information Processing Systems 23 (NIPS 2013)*. (**Best Paper Award**)
41. * **K. Zhou**, H. Zha and L. Song (2013). Learning the influence kernel in Hawkes process. *30th International Conference on Machine Learning (ICML 2013)*.
42. * **N. Du**, L. Song, H. Woo and H. Zha (2013). Uncover topic-sensitive information diffusion networks. *Artificial Intelligence and Statistics (AISTATS 2013)*. (**Plenary Long Talk**)
43. * **K. Zhou**, H. Zha and L. Song (2013). Learning social infectivity in sparse low-rank networks using multi-dimensional Hawkes processes. *Artificial Intelligence and Statistics (AISTATS 2013)*.
44. * L. Song and **B. Dai** (2013). Robust low rank kernel embedding of multivariate distributions. *Advances in Neural Information Processing Systems 23 (NIPS 2013)*.
45. * L. Song, M. Ishteva, H. Park, A. Parikh, E. Xing (2013). Hierarchical tensor decomposition for latent variable models. *30th International Conference on Machine Learning (ICML 2013)*.
46. * M. Ishteva, H. Park, L. Song (2013). Unfolding latent tree structure using 4th order tensors. *30th International Conference on Machine Learning (ICML 2013)*.
47. ***E. Khalil**, B. Dilkina and L. Song (2013). CuttingEdge: influence minimization in networks. NIPS Workshop on Frontiers of Network Analysis: Methods, Models and Applications. (**Best Paper Award**)
48. * A. Parikh, L. Song, M. Ishteva and E. Xing (2012). A spectral algorithm for latent junction trees. *Uncertainty in Artificial Intelligence 28 (UAI 2012)*. (**Plenary Long Talk**)
49. * **N. Du**, L. Song, A. Smola and M. Yuan (2012). Learning networks of heterogeneous influence. *Advances in Neural Information Processing Systems 22 (NIPS 2012)*. (**Plenary Spotlight**)
50. Q. Ho, L. Song and E. Xing (2011). Evolving cluster mixed-membership blockmodel for time-varying networks. *Artificial Intelligence and Statistics (AISTATS 2011)*.
51. Q. Ho, A. Parikh, L. Song and E. Xing (2011). Multiscale community blockmodel for network exploration. *Artificial Intelligence and Statistics (AISTATS 2011)*.
52. K. Fukumizu, L. Song and A. Gretton (2011). Kernel Bayes rules. *Advances in Neural Information Processing Systems 21 (NIPS 2011)*.
53. L. Song, A. Parikh and E. Xing (2011). Hilbert space embedding of latent tree graphical models. *Advances in Neural Information Processing Systems 21 (NIPS 2011)*.

54. A. Anandkumar, K. Chaudhuri, D. Hsu, S. Kakade, L. Song and T. Zhang (2011). Spectral methods for learning multivariate latent tree structure. *Advances in Neural Information Processing Systems 21 (NIPS 2011)*.
55. A. Parikh, L. Song and E. Xing (2011). A spectral algorithm for latent tree graphical models. *28th International Conference on Machine Learning (ICML 2011)*.
56. L. Song, A. Gretton, D. Bickson, Y. Low and C. Guestrin (2011). Kernel belief propagation. *Artificial Intelligence and Statistics (AISTATS 2011)*.
57. L. Song, B. Boots, S. Siddiqi, G. Gordon and A. Smola (2010). Hilbert space embedding of hidden Markov model. *27th International Conference on Machine Learning (ICML 2010)* (**Best Paper Award**).
58. L. Song, A. Gretton and C. Guestrin (2010). Nonparametric tree graphical models. *Artificial Intelligence and Statistics (AISTATS 2010)*.
59. T. Huang, L. Song and J. Schneider (2010). Learning nonlinear dynamic models from non-sequenced data. *Artificial Intelligence and Statistics (AISTATS 2010)*.
60. L. Song, J. Huang, A. Smola and K. Fukumizu (2009). Hilbert space embedding of conditional distribution with applications to dynamical systems. *26th International Conference on Machine Learning (ICML 2009)*.
61. A. Smola, L. Song and C. Teo. (2009). Relative novelty detection. *Artificial Intelligence and Statistics (AISTATS 2009)*.
62. M. Thoma, H. Cheng, A. Gretton, J. Han, H. Kriegel, A. Smola, L. Song, P. Yu, X. Yan and K. Borgwardt. (2009). Near-optimal supervised feature selection among frequent subgraphs. *SIAM International Conference on Data Mining (SDM 2009)*.
63. M. Kolar, L. Song and E. Xing (2009). Sparsistent learning of varying-coefficient models with structural changes. *Advances in Neural Information Processing Systems 22 (NIPS 2009)*. (**Plenary Spotlight**)
64. W. Fu, L. Song and E. Xing (2009). Dynamic mixed membership blockmodel for evolving networks. *26th International Conference on Machine Learning (ICML 2009)*.
65. L. Song, M. Kolar and E. Xing (2009). Time-varying dynamic Bayesian networks. *Advances in Neural Information Processing Systems 22 (NIPS 2009)*. (**Plenary Spotlight**)
66. L. Song, M. Kolar, E. Xing. (2009). KELLER: estimating time-varying interactions between genes. *Intelligent Systems in Molecular Biology (ISMB 2009)*.
67. X. Zhang, L. Song, A. Gretton and A. Smola. (2008). Kernel measures of independence for non-iid data. *Advances in Neural Information Processing Systems 21 (NIPS 2008)*. (**Plenary Spotlight**)
68. N. Quadrianto, L. Song and A. Smola. (2008). Kernelized sorting. *Advances in Neural Information Processing Systems 21 (NIPS 2008)*. (**Plenary Spotlight**)
69. L. Song, X. Zhang, A. Smola, A. Gretton and B. Schölkopf. Tailoring density estimation via reproducing kernel moment matching. (2008). *25th International Conference on Machine Learning (ICML 2008)*.

70. L. Song, A. Smola, K. Borgwardt and A. Getton. (2007). Colored maximum variance unfolding. *Advances in Neural Information Processing Systems 20 (NIPS 2007)*. (**Plenary Long Talk**)
71. A. Gretton, K. Fukumizu, C.H. Teo, L. Song, B. Schölkopf and A. Smola. (2007). A kernel statistical test of independence. *Advances in Neural Information Processing Systems 20 (NIPS 2007)*. (**Plenary Spotlight**)
72. A. Smola, A. Gretton, L. Song and B. Schölkopf. (2007). A Hilbert space embedding for distributions. *18th International Conference on Algorithmic Learning Theory (ALT 2007)*.
73. L. Song, A. Smola, Arthur Gretton, K. Borgwardt and J. Bedo. (2007). Supervised feature selection via dependence estimation. *24th International Conference on Machine Learning (ICML 2007)*.
74. L. Song, A. Smola, Arthur Gretton and K. Borgwardt. (2007). A dependence maximization view of clustering. *24th International Conference on Machine Learning (ICML 2007)*.
75. L. Song, J. Bedo, K. Borgwardt, A. Getton and A. Smola. (2007). Gene selection via the BAHASIC family of algorithms. *Intelligent Systems in Molecular Biology (ISMB 2007)*.
76. L. Song, and J. Epps. (2006). Classifying EEG for brain-computer interfaces: learning optimal filters for dynamical system features. *23rd International Conference on Machine Learning (ICML 2006)*.
77. L. Song, and J. Epps. (2006). Improving the separability of EEG signals during motor imagery with an efficient circular Laplacian. *31st IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2006)*.
78. L. Song, E. Gordon, and E. Gysels. (2005). Phase synchrony rate for the recognition of motor imagery in brain-computer interface. *Advances in Neural Information Processing Systems 18 (NIPS 2005)*.
79. L. Song. (2005). Desynchronization network analysis for the recognition of imagined movement. *27th IEEE International Conference of the Engineering in Medicine and Biology Society (EMBC 2005)*.
80. W. Huang, C. Murray, X. Shen, L. Song, Y.X. Wu, and L. Zheng. (2005). Visualization and analysis of network motifs. *9th International Conference on Information Visualization (IV 2005)*.
81. A. Ahmed, T. Dywer, S.H. Hong, C. Murray, L. Song, and Y.X. Wu. (2005). Visualization and analysis of large and complex scale-free networks. *7th IEEE VGTC Symposium on Visualization (EUROGRAPHICS 2005)*.
82. L. Zheng, L. Song and P. Eades. (2005). Crossing minimization problems of drawing bipartite graphs in two clusters. *4th Asian-Pacific Symposium on Information Visualization (APVIS 2005)*.
83. L. Song, and M. Takatsuka. (2005). Real-time 3D finger pointing for an augmented desk. *6th Australasian User Interface Conference (AUIC 2005)*.
84. A. Ahmed, T. Dywer, S.H. Hong, C. Murray, L. Song, and Y.X. Wu. (2004). Wilmascope graph visualization. *10th IEEE Symposium on Information Visualization (IEEE INFOVIS 2004)*.

KEYNOTES AND INVITED TALKS

1. December 2016, Co-coevolutionary feature embedding processes for recommendation and knowledge reasoning. Tsinghua University, China.
2. December 2016, Co-coevolutionary feature embedding processes for recommendation and knowledge reasoning. MIFS, Chinese Academy of Engineering, China.
3. December 2016, Co-coevolutionary feature embedding processes for recommendation and knowledge reasoning. Toutiao Lab, China.
4. December 2016, Co-coevolutionary feature embedding processes for recommendation and knowledge reasoning. Microsoft Research Asia, China.
5. November 2016, Co-coevolutionary feature embedding processes for recommendation and knowledge reasoning. Amazon, Seattle, USA.
6. November 2016, Discriminative embedding of latent variable models for structured data. Berkeley, UC Berkeley, USA.
7. November 2016, Discriminative embedding of latent variable models for structured data. Palo Alto, Stanford University, USA.
8. September 2016, Discriminative embedding of latent variable models for structured data. Caltech, Pasadena, USA.
9. September 2016, Discriminative embedding of latent variable models for structured data. USC, Los Angeles, USA.
10. September 2016, Discriminative Embedding of Molecular Structures for Property Prediction. UCLA IPAM workshop on Machine Learning Meets Many-Particle Problems, Los Angeles, USA.
11. September 2016, Understanding deep learning for big data. MLconf machine learning conference series, Atlanta, USA.
12. September 2016, Discriminative embedding of latent variable models for structured data. MIT, Boston, USA.
13. September 2016, Coevolutionary latent feature processes for continuous-time user-item interactions. Recsys workshop on deep learning for recommendation systems. Boston, USA.
14. August 2016, Dynamic processes over network: representation, modeling, learning and inference. **Keynote Talk** at KDD workshop on time-series analysis, San Francisco, USA.
15. July 2016, Discriminative embedding of latent variable models for structured data. Cambridge University, Cambridge, UK.
16. July 2016, Discriminative embedding of latent variable models for structured data. University College of London, London, UK.
17. July 2016, Time-sensitive recommendation. BBVA Data & Analytics, Madrid, Spain.

18. May 2016, Dynamic processes over network: representation, modeling, learning and inference. Machine Learning Summer School, Cardiz, Spain.
19. March 2016, Kernel methods for graphical models. Workshop on probabilistic graphical models, ISM, Tokyo, Japan.
20. March 2016, Understanding deep learning via kernel methods. University of Tokyo, Tokyo, Japan.
21. March 2016, Machine learning for big nonlinear models. Symposium on big data in analytical sciences: challenges and solutions, Pittcon, Atlanta, USA.
22. March 2016, Understanding deep learning via kernel methods. Google Machine Learning Workshop, New York, USA.
23. January 2016, COEVOLVE: a joint model of information diffusion and network evolution. Institute of Advanced Study, Center for the Statistical Science, Hong Kong University of Science and Technology, Hongkong.
24. December 2015, Kernel methods for learning big nonlinear models. Workshop on modern questions and challenges of feature extraction, Neural Information Processing Systems, Montreal, Canada.
25. October 2015, Large scale and distributed algorithms for nonparametric and nonlinear machine learning models. China National Computer Congress, Hefei, China.
26. October 2015, Scalable kernel methods for big nonlinear models. Microsoft Research Asia, Beijing, China.
27. October 2015, Scalable kernel methods for big nonlinear models. General Electric, Atlanta, USA.
28. October 2015, Estimating networks from high dimensional temporal data. Dept. of Biochemistry and Molecular Biology, University of Georgia, Athena, USA.
29. September 2015, Scalable kernel methods for big nonlinear models. Booth School of Business, University of Chicago, Chicago, USA.
30. September 2015, Estimating networks from high dimensional temporal data. Rollin School of Public Health, Emory University, Atlanta, USA.
31. August 2015, Modeling Temporal Information. Homedepot. Atlanta, USA
32. March 2015, Shaping social activity by incentivizing users. Department of Statistics, University of Washington, Seattle, USA.
33. March 2015, Shaping social activity by incentivizing users. Twitter, San Francisco, USA.
34. March 2015, Scalable kernel methods via doubly stochastic gradients. Microsoft Research Redmond, Seattle, USA.
35. March 2015, Scalable kernel methods via doubly stochastic gradients. Industrial and System Engineering, Georgia Institute of Technology, Atlanta, USA
36. October 2014, Scalable learning of nonparametric latent variable models. Department of Statistics, Oxford University, Oxford, UK.

37. October 2014, Scalable learning of nonparametric latent variable models. Gatsby Computational Unit, University of College London, London, UK.
38. October 2014, Scalable learning of nonparametric latent variable models. Department of Computer Science, Department of Statistical Science, Duke University, Durham, USA.
39. June 2014, Learning and inference of high dimensional “asynchronous” and “interdependent” event data. School of Mathematical Sciences, Fudan University, Shanghai, China.
40. April 2014, Learning and inference of high dimensional “asynchronous” and “interdependent” event data. School of Mathematics, Georgia Institute of Technology, Atlanta, USA.
41. March 2014, Learning and inference of high dimensional “asynchronous” and “interdependent” event data. Graduate School of Informatics, Kyoto University, Kyoto, Japan.
42. December 2013, Scalable influence estimation in continuous-time diffusion models. Randomized Machine Learning Algorithm Workshop, Neural Information Processing Systems, Lake Tahoe, USA.
43. December 2013, Learning nonparametric latent variable models with kernels. Spectral Learning Workshop, Neural Information Processing Systems, Lake Tahoe, USA
44. November 2013, Learning nonparametric latent variable models via Hilbert space embedding of distributions. Max Planck Institute for Intelligent Systems, Tübingen, Germany.
45. October 2013, Learning high-dimensional multivariate point processes. Microsoft Research Asia, Beijing, China.
46. October 2013, Learning high-dimensional multivariate point processes. Computer Science Department, Tsinghua University, Beijing, China.
47. October 2013, Learning high-dimensional multivariate point processes. Machine Learning Department, Carnegie Mellon University, Pittsburgh, USA.
48. June 2013 Learning nonparametric latent variable models with kernels. Spectral Learning Workshop, International Conference on Machine Learning, Atlanta, USA
49. Sept 2012, Learning Networks of Heterogeneous Influence. Workshop on “Networks: Processes and Causality”, Menorca, Spain.
50. April 2012, Hilbert space embedding of distributions. Computer Science Department, King Abdullah University of Science and Technology, Thuwal, Saudi Arabia
51. April 2011, Modeling Rich Structured Data via Kernel Distribution Embeddings. Electrical Engineering and Computer Science Department, Massachusetts Institute of Technology, Cambridge, USA.
52. March 2011, Modeling Rich Structured Data via Kernel Distribution Embeddings. Computer Science Department, Stanford University, Stanford, USA.
53. March 2011, Modeling Rich Structured Data via Kernel Distribution Embeddings. Computational Science and Engineering Department, Georgia Institute of Technology, Atlanta, USA.

54. February 2011, Modeling Rich Structured Data via Kernel Distribution Embeddings. Department of Biostatistics, Harvard School of Public Health, Cambridge, USA.
55. June 2010, Graphical Models via Kernels. Google Research, New York, USA.
56. May 2010, Graphical Models via Kernels. Yahoo! Research, Sunnyvale, USA.
57. December 2007, Learning via Hilbert space embedding of distributions. Microsoft Research, Seattle, USA.
58. December 2007, Colored maximum variance unfolding. Machine Learning Department, Carnegie Mellon University, Pittsburgh, USA.

GRANTS AND CONTRACTS

As Principal Investigator

1. 2016 Nvidia Hardware grant. Machine Learning for Materials Structure-Property Linkage. **\$3,000**.
2. 2015 – 2018 ONR N00014-15-1-2340: Modeling, Estimation, and Management of High Dimensional Asynchronous Event Data. **\$395,000**.
3. 2014 – 2019 NSF CAREER IIS-1350983: New Representations of Probability Distributions to Improve Machine Learning — A Unified Kernel Embedding Framework for Distributions. **\$499,000**.
4. 2013 – 2016 NIH BIGDATA 1R01GM108341: Social Behavior Driven Modeling and Optimization of Information Diffusion. **\$450,000**. Collaborator: Hongyuan Zha (Co-PI). Le Song's share: 50%
5. 2013 – 2014 Raytheon Faculty Fellowship. **\$75,000**. Collaborator: Nina Balcan (Co-PI). Le Song's share: 50%
6. 2012 – 2015 NSF IIS1218749: Efficient, Nonparametric and Local-Minimum-Free Latent Variable Models: With Application to Large-Scale Computer Vision and Genomics. **\$500,000**. Collaborator: Eric Xing (Co-PI). Le Song's share: 60%
7. 2008 – 2011 Lane Fellowship. **\$198,000**.

As Co-Principal Investigator

1. 2016 – 2019 Intel: MLsploit: A Framework for Evaluating and Fortifying ML-Based Security Analytics. **\$1,500,000**. Collaborators: Wenke Lee (PI), Taesoo Kim (Co-PI), Polo Chau (Co-PI). Le Song's share: 25%
2. 2016 – 2018 NSF IIS-1639792 EAGER: Asynchronous Event Models for State-Topology Co-Evolution of Temporal Networks. **\$200,000**. Collaborators: Hongyuan Zha (PI). Le Song's share: 50%
3. 2016 Data science platform. **\$75,000** from Georgia Tech and 200 server donation from Yahoo Labs, as part of Yahoo Servers to Academic Researchers program. Collaborator: Polo Chau (PI). Le Song's share: 50%

4. 2011 – 2014 US Government: Multi-scale Characterization of Document Themes, Granularity, and Triage. **\$450,000**. Collaborator: Haesun Park (PI). Le Song's share: 50%

As Senior Personnel or Contributor

1. 2013 – 2018 Senior Personnel: NSF IGERT-CIF21: Computation-Enabled Design and Manufacturing of High Performance Materials. **\$2,800,000**.
2. 2010 – 2015 Senior Personnel: NIH 1R01GM093156: Time/Space-Varying Networks of Molecular Interactions: A New Paradigm for Studying Dynamic Biological Regulation and Pathways. **\$2,237,000**.

Other Professional Activities

1. 2017–present. Consulting Amazon on project related to online recommendation.
2. 2016–2016. Consulting Homedepot on project related to online recommendation.
3. 2015–2016. Consulting TrustVesta on projects related fraud detection.
4. 2011–2012. Consulting Microsoft on projection related to deep learning.

TEACHING

Courses Developed

1. Spring 2015 CX4240: Introductory Computational Data Analysis at Georgia Institute of Technology (class size: 41)
This undergraduate course focuses on basic computational techniques for data analysis. Students have very diverse backgrounds with majors other than computer science.
2. Spring 2014 CX4240: Introductory Computational Data Analysis at Georgia Institute of Technology (class size: 47)
This undergraduate course focuses on basic computational techniques for data analysis. Students have very diverse backgrounds with majors other than computer science.
3. Spring 2013 ML8803: Advanced Machine Learning at Georgia Institute of Technology (class size: 12)
This course focuses on advanced machine learning methodologies, including specialized topics in graphical models and kernel methods. Students have very diverse background, including CS, CSE, ISYE and Biology.
4. Spring 2012 ML8803: Advanced Machine Learning at Georgia Institute of Technology (class size: 35)
This course focuses on advanced machine learning methodologies, including specialized topics in graphical models and kernel methods. Students have very diverse backgrounds, including CS, CSE, ISYE and Biology.

Courses Taught

1. Fall 2016 CSE6740/ISYE6740/CS7641: Machine Learning and Data Analytics at Georgia Institute of Technology (class size: 280)
This course is cross-listed across the campus, and it focuses on basic techniques and concepts in machine learning. Students have very diverse background, including CS, CSE, ISYE, Biology, Biomedical Engineering, Civil Engineering, Aerospace Engineering and Architecture.
2. Fall 2015 CSE6740/ISYE6740/CS7641: Machine Learning and Data Analytics at Georgia Institute of Technology (class size: 210)
This course is cross-listed across the campus, and it focuses on basic techniques and concepts in machine learning. Students have very diverse background, including CS, CSE, ISYE, Biology, Biomedical Engineering, Civil Engineering, Aerospace Engineering and Architecture.
3. Fall 2014 CSE6740/ISYE6740/CS7641: Machine Learning and Data Analytics at Georgia Institute of Technology (class size: 173)
This course is cross-listed across the campus, and it focuses on basic techniques and concepts in machine learning. Students have very diverse background, including CS, CSE, ISYE, Biology, Biomedical Engineering, Civil Engineering, Aerospace Engineering and Architecture.
4. Fall 2013 CSE6740/ISYE6740/CS7641: Machine Learning and Data Analytics at Georgia Institute of Technology (class size: 117)
This course is cross-listed across the campus, and it focuses on basic techniques and concepts in machine learning. Students have very diverse background, including CS, CSE, ISYE, Biology, Biomedical Engineering, Civil Engineering, Aerospace Engineering and Architecture.
5. Fall 2012 CSE6740/ISYE6740/CS7641: Machine Learning and Data Analytics at Georgia Institute of Technology (class size: 67)
This course is cross-listed across the campus, and it focuses on basic techniques and concepts in machine learning. Students have very diverse backgrounds, including CS, CSE, ISYE, Biology, Biomedical Engineering, Civil Engineering, Aerospace Engineering and Architecture.

Individual Student Guidance

Ph.D. Students

Primary advisor:

1. 2014–present, Hanjun Dai at Georgia Tech
Award: Best Paper at Recyc workshop on deep learning, Best Student Paper at AISTATS 2016
2. 2015–present, Yichen Wang at Georgia Tech
Award: Best graduate research assistant 2016
3. 2013–present, Bo Dai at Georgia Tech
Award: Best Student Paper at AISTATS 2016
4. 2012–present, Bo Xie (with S. Aluru and H. Park) at Georgia Tech
Award: ARC fellowship 2015

5. 2011–2016, Nan Du at Georgia Tech (**Graduated**)
 Position: Researcher at Google
 Thesis title: Learning, inference and optimization of high dimensional asynchronous event data
Award: Best Paper at NIPS 2013, Facebook Fellowship 2014, Outstanding Graduate Research Assistant at Georgia Tech 2015, Microsoft Research Asia Outstanding Ph.D. student 2015.

Co-advisor:

1. 2013–present, Mehrdad Farajtabar (with Hongyuan Zha) at Georgia Tech
2. 2015–present, Shuang Li (with Yao Xie at ISYE) at Georgia Tech
Award: Best graduate research award ISYE 2016
3. 2014–present, Elias Kahlil (with Bistra Dilkina) at Georgia Tech
Award: Best Student Paper at NIPS social network workshop 2013, Marshall D. Williamson Fellowship 2012, Donald V. Jackson Fellowship 2013
4. 2011–2015, Chris Berlind (with Nina Balcan) at Georgia Tech (**Graduated**)
 Thesis title: New theoretical results on active learning
 Position: CTO of Oncora Medical
5. 2009–2014, Ankur Parikh (with E. Xing) at CMU (**Graduated**)
 Position: Researcher at Google
6. 2009–2014, Qirong Ho (with E. Xing) at CMU (**Graduated**)
 Position: Scientist at Institute for Infocomm Research, A*STAR, Singapore
7. 2004–2009, Mladen Kolar (with E. Xing) at CMU (**Graduated**)
 Position: Assistant Professor at University of Chicago

M.S. Students (Indicate thesis option for each student)

1. 2016–present, Apurv Verman at Georgia Tech
 Thesis title: Automatic question answering using deep learning in knowledge graph
2. 2013–2015, Kaushik Patnaik at Georgia Tech (**Graduated**)
 Position: Data Scientist at Element Analytics
 Thesis title: Active learning for graphical models
Award: Master Research Award at Georgia Tech 2015
3. 2012–2014, Elias Khalil at Georgia Tech (**Graduated**)
 Thesis title: Optimizing the Structure of Diffusion Networks: Theory and Algorithms **Award:** Best Student Paper at NIPS social network workshop 2013, Marshall D. Williamson Fellowship 2012, Donald V. Jackson Fellowship 2013
4. 2011–2013, Hai Shang at Georgia Tech (**Graduated**)
 Position: Software Engineer at Google

Undergraduate Students

1. 2016–present, Jenna Kwon at Georgia Tech
2. 2016–present, Yuxuan Zhuge at Georgia Tech

Mentorship of postdoctoral fellows or visiting scholars

- 2016 – 2017 Yamin Zhang (visiting scholar from Chinese Academy of Science)
- 2011 – 2012 Maria Ishteva (with Haesun Park)

SERVICE

*Professional Contributions***Tutorials**

1. Dynamic Processes over Networks: Representation, Modeling, Learning and Inference, Machine Learning Summer School, Spain, 2016.
2. Diffusion in Social and Information Networks: Research Problems, Probabilistic Models & Machine Learning Methods, Knowledge Discovery and Data Mining (KDD), 2015
3. Diffusion in Social and Information Networks: Research Problems, Probabilistic Models & Machine Learning Methods, World Wide Web (WWW), 2015
4. Spectral Approaches to Learning Latent Variable Models, International Conference on Machine Learning (ICML), 2012.

Workshop Organizer or Program Chair

1. Organizer, Workshop on Large-Scale Kernel Learning: Challenges and New Opportunities, International Conference on Machine Learning (ICML), 2015
2. Organizer, Workshop on Diffusion, Activity and Events in Networks: Models, Methods & Applications, World Wide Web (WWW), 2015
3. Organizer, Workshop on Modern Nonparametrics: Automating the Learning Pipeline, Neural Information Processing Systems (NIPS), 2014
4. Program Chair, Workshop on Spectral Algorithms for Latent Variable Models, Neural Information Processing Systems (NIPS), 2013
5. Program Chair, Workshop on Spectral Algorithms for Latent Variable Models, International Conference on Machine Learning (ICML), 2013
6. Organizer, Workshop on Spectral Algorithms for Latent Variable Models, Neural Information Processing Systems (NIPS), 2012
7. Organizer, Workshop on Confluence between Kernel Methods and Graphical Models, Neural Information Processing Systems (NIPS), 2012

8. Organizer, Workshop on Transfer Learning for Structured Data, Neural Information Processing Systems (NIPS), 2009

Public and Community Service

Conference Area Chair or Program Committee

1. Area Chair, International Joint Conference on Artificial Intelligence (IJCAI), 2017
2. Area Chair, Neural Information Processing Systems (NIPS), 2017
3. Area Chair, International Conference on Machine Learning (ICML), 2017
4. Area Chair, Artificial Intelligence and Statistics (AISTATS), 2017
5. Area Chair, Association for the Advancement of Artificial Intelligence conference (AAAI), 2017
6. Area Chair, International Conference on Machine Learning (ICML), 2016
7. Area Chair, Artificial Intelligence and Statistics (AISTATS), 2016
8. Area Chair, Neural Information Processing Systems (NIPS), 2015
9. Area Chair, Artificial Intelligence and Statistics (AISTATS), 2015
10. Senior Program Committee, International Joint Conference on Artificial Intelligence (IJCAI), 2015
11. Program Committee, AAAI Conference on Artificial Intelligence (AAAI), 2015
12. Program Committee, Web Search and Data Mining (WSDM), 2015
13. Program Committee, Neural Information Processing Systems (NIPS), 2014
14. Program Committee, Uncertainty on Artificial Intelligence (UAI), 2014
15. Program Committee, Knowledge Discovery and Data Mining (KDD), 2014
16. Area Chair, International Conference on Machine Learning (ICML), 2014
17. Program Committee, Artificial Intelligence and Statistics (AISTATS), 2014
18. Area Chair, Neural Information Processing Systems (NIPS), 2013
19. Program Committee, International Conference on Machine Learning (ICML), 2013
20. Program Committee, Uncertainty on Artificial Intelligence (UAI), 2013
21. Program Committee, Neural Information Processing Systems (NIPS), 2012
22. Program Committee, International Conference on Machine Learning (ICML), 2012
23. Program Committee, Neural Information Processing Systems (NIPS), 2011
24. Program Committee, International Conference on Machine Learning (ICML), 2011

25. Program Committee, multiple year between 2008–2011
 - Neural Information Processing Systems (NIPS)
 - International Conference on Machine Learning (ICML)
 - Knowledge Discovery and Data Mining (KDD)
 - Research on Computational Biology (RECOMB)
 - International Conference on Intelligent Systems for Molecular Biology (ISMB)
 - International Conference on Data Mining (ICDM)

Journals Reviewer and Editor

1. Journal of Machine Learning Researches (JMLR) (**Action Editor**)
2. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)
3. IEEE Transactions on Neural Networks (TNN)
4. IEEE Transactions on Knowledge and Data Engineering (TKDE)
5. IEEE Transactions on Information Theory (TIT)
6. IEEE Transactions on Signal Processing (TSP)
7. Transactions on Computational Biology and Bioinformatics (TCBB)
8. Journal of Artificial Intelligence Research (JAIR)
9. Journal of Graphical and Computational Statistics (JGCS)
10. Machine Learning
11. Pattern Recognition
12. Bioinformatics

Books Reviewer

1. Cambridge University Press
2. Oxford University Press
3. Springer
4. Chapman & Hall/CRC

Research Project Reviewer

1. National Science Foundation
2. Israel National Science Foundation
3. Czech National Science Foundation
4. National Security Agency Mathematical Sciences Grant Program
5. Austrian Academy of Sciences

Institute Contributions

1. 2014–2016 Machine Learning Seminar Committee, College of Computing, Georgia Institute of Technology
2. 2013–2016 CSE Seminar Committee, Georgia Institute of Technology
3. 2014–present Machine Learning Taskforce, College of Computing, Georgia Institute of Technology
4. 2012–2016 PhD Qualification Exam Committee on Machine Learning Area, CSE, Georgia Institute of Technology
5. 2011–2013 PhD Admission Committee, CSE, Georgia Institute of Technology
6. 2011–2013 Master’s Admission Committee, CSE, Georgia Institute of Technology