

RUSLAN SALAKHUTDINOV

Gates Hillman Center, Office 8017,
5000 Forbes Avenue, Pittsburgh, PA 15213
rsalakhu@cs.cmu.edu
<http://www.cs.cmu.edu/~rsalakhu/>

Professional Career

Associate Professor, Carnegie Mellon University Machine Learning Department	Feb 2016 -
Assistant Professor, University of Toronto Department of Computer Science and Department of Statistical Sciences	August 2011 - January, 2016
Postdoctoral Research Associate, Brain and Cognitive Sciences (BCS) and Computer Science and Artificial Intelligence Lab (CSAIL), MIT.	Sep. 2009 - July 2011
PhD, Department of Computer Science, University of Toronto. Thesis: Learning Deep Generative Models.	Sep. 2005 - Aug. 2009
Master of Science, Department of Computer Science, University of Toronto. Thesis: Optimization Algorithms for Learning.	Sep. 2001 - Aug. 2003
Bachelor of Science, High Point University, NC, USA. Double major in Computer Science and Mathematics, Honors Degree.	Aug. 1998 - May. 2001

Teaching Experience

• 10703: Deep Reinforcement Learning and Control.	Spring 2017, 2018
• 10807, Topics in Deep Learning.	Fall 2016, 2017
• CSC 411, Introduction to Machine Learning.	Fall 2015
• STA 4273, Large-Scale Machine Learning (graduate course at the Fields Institute).	Winter 2015
• STA 414/2104, Statistical Methods for Machine Learning and Data Mining.	Winter 2015, Fall 2016
• STAD 68, Advanced Machine Learning and Data Mining.	Winter 2014
• STAD 37, Statistical Multivariate Analysis.	Winter 2012, Winter 2013
• STA 4273, Research Topics In Statistical Machine Learning.	Fall 2011, Fall 2012, Fall 2013
• Two Guest Lectures, MIT 9.520: Statistical Learning Theory and Applications.	Spring 2010
• Substitute Lecture, MIT 9.660: Computational Cognitive Science.	Fall 2009
• Teaching Assistant, Department of Computer Science, University of Toronto. CSC321, Introduction to Neural Networks and Machine Learning.	Spring 2009.
CSC2515, Machine Learning.	Fall 2006, 2007, 2008
CSC2506, Probabilistic Reasoning.	Spring 2006
CSC412, Uncertainty and Learning in Artificial Intelligence.	Spring 2003

Professional Experience

• Visiting Scientist, Stanford University.	Jul 2012 - Aug 2012
• Yahoo Research, New York, USA. Summer Intern.	Jun 2008 - Aug 2008
• Canadian Imperial Bank of Commerce (CIBC), Toronto, Canada.	Sep 2003 - Aug 2005

Awards

• Google Focused Award (co-PI),	(2017 - 2019)
• Pioneers of AI Research, Nvidia.	April 2016
• Canada Research Chair in Statistical Machine Learning.	March 2016

- Google Faculty Award. April 2014
- Microsoft Research Faculty Fellowship. April 2013
- Sloan Research Fellowship. May 2013
- Best Student Paper Award, Conference on Uncertainty in Artificial Intelligence. July 2012
- Connaught New Researcher Award. Aug 2012
- Early Researcher Award. Aug 2012
- Canadian Institute for Advanced Research
Fellow of the Neural Computation and Adaptive Perception Program. Sep 2011
- Precarn Scholar, Canada. Sep 2002

Grants and Scholarships

- DARPA Award (co-PI), (2017 - 2019)
- Apple Award (PI), (2017 - 2019)
- IARPA (co-PI), (2016 - 2017)
- AFOSR, jointly with BAE Systems (2016 - 2019)
- Canada Research Chair in Statistical Machine Learning. March 2016
- Disney Research Grant, Disney. Nov 2015
- Samsung, Deep Learning Grant (PI). (2015-2017)
- Raytheon BBN Technologies, Deep Learning for Video Analysis. (2014-2016)
- Samsung, Deep Learning for Face Recognition. (2014-2015)
- Office of Naval Research Grant (Co-PI). (2014 - 2018)
- NSERC Engage Grant. (2014)
- NSERC Individual Discovery Grant,
along with NSERC Early Career Researcher Supplement. (2012-2017)
- Natural Sciences and Engineering Research Council of Canada:
Postdoctoral Fellowship. (2009-2011)
- The UK Engineering and Physical Sciences Research Council:
Postdoctoral Fellowship in Theoretical Computer Science (Declined). (2009-2012)
- Natural Sciences and Engineering Research Council (NSERC) of Canada:
Canada Graduate Scholarship. (2006-2009)
- Ontario Graduate Scholarship (Declined). (2006-2007)

Selected Invited Tutorials and Short Courses

- Foundations of Deep Learning and its Recent Advances, ML Summer School, Spain. July 2017
- Foundations of Deep Learning, STOC, Montreal. June 2017
- Deep Learning Tutorial, ML Summer School, Tbingen, Germany. June 2017
- Deep Learning, Simons Institute for the Theory of Computing, Berkeley. March 2017
- Foundations of Unsupervised Deep Learning, Stanford. May 2016
- Learning Deep Generative Models, CIFAR Summer School, Montreal. Aug 2016
- Deep Learning Tutorial, CIFAR Summer School, Montreal. Aug 2015
- Deep Learning Tutorial, Machine Learning and Intelligence School, Russia. June 2015
- Deep Learning Tutorial, KDD, New York. Aug 2014
- Deep Learning Tutorial, Machine Learning Summer School, Beijing. June 2014
- Deep Learning Tutorial, SORA-BN-TABA Workshop, Biostatistics Research Day. May 2014
- Deep Learning Tutorial, IEEE International Symposium on Biomedical Imaging, invited tutorial. April 2013
- IPAM Graduate Summer School on Deep Learning, UCLA, invited tutorial. July 2012
- CVPR Tutorial on Deep Learning Methods for Vision, Providence, RI. June 2012
- CIFAR Graduate Summer School on Machine Learning, University of Toronto, invited tutorial. Aug 2011
- IPAM Graduate Summer School on Probabilistic Models of Cognition, UCLA, invited tutorial. July 2011

Professional Activities

Senior Program Committees:

- Area Chair, NIPS 2017, NIPS 2016, 2015, 2012, 2011; ICML 2016, 2015, 2013, 2012, 2011; CVPR 2015; ICLR 2016, 2017.
- Workshop Chair, UAI 2012, ICML 2016.
- Demonstration Chair, NIPS 2013.
- Tutorial Chair, ICML 2018, ICML 2014.

Associate Editor and Editorial Board

- Action Editor, Journal of Machine Learning Research (JMLR).
- Editorial Board, Journal of Artificial Intelligence Research (JAIR).
- Guest editor for IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI), Special Issue on Learning Deep Architectures, 2014.

Workshops Organized:

- Co-chair, ICML 2018 workshop on Theoretical Foundations of Deep Generative Models.
- Co-chair, NIPS 2017 workshop on Deep Learning: Bridging Theory and Practice.
- Co-chair, NIPS 2013 Deep Learning Workshop.
- Co-chair, AAI 2014 Workshop on Cognitive Computing for Augmented Human Intelligence.
- NIPS 2011 Workshop on Learning Hierarchical Models: Transfer Learning and Optimization, Co-chair.
- NIPS 2010 Workshop on Transfer Learning Via Rich Generative Models, Co-chair.
- NIPS 2010 Workshop on Deep Learning, program committee.
- NIPS 2009 Workshop on Approximate Learning of Large Scale Graphical Models, Co-chair.
- ICML 2009 Workshop on Learning Feature Hierarchies, Co-chair.
- NIPS mini-symposium, Deep Learning: Foundations and Future Directions, 2007, Co-chair.

Reviewing Activity:

- Journal of Machine Learning Research, 2006, 2008, 2009, 2010, 2011, 2012, 2013, 2014, 2015, 2016.
- Journal of Artificial Intelligence Research, 2011, 2012, 2014.
- Machine Learning Journal, 2006, 2007, 2009.
- Statistics and Computing, 2010, 2013, 2014.
- Transactions on Modeling and Computer Simulation, 2012.
- Science, 2009.
- The Journal of Computer and System Sciences (JCSS), 2010.
- IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI), 2005, 2006, 2008, 2011, 2012, 2013, 2014.
- IEEE Transactions on Signal Processing, 2007.
- Neural Computation, 2006, 2007, 2008, 2011, 2012, 2013, 2014.
- Advances in Neural Information Processing Systems (NIPS), 2006, 2007, 2008, 2009, 2010, 2013, 2014.
- International Conf. on Artificial Intelligence and Statistics (AISTATS 2007,2009,2010, 2014).
- International Conf. on Machine Learning (ICML 2007, 2008, 2009, 2010, 2014).
- International Joint Conf. on Artificial Intelligence, Program Committee (IJCAI 2009).

Refereed scientific publications:

Citations according to Google Scholar: 34404, H-index: 58

1. Style Transfer Through Back-Translation
Shrimai Prabhumoye, Yulia Tsvetkov, Ruslan Salakhutdinov and Alan Black
ACL 2018

2. Neural Models for Reasoning over Multiple Mentions using Coreference
Bhuwan Dhingra, Qiao Jin, Zhilin Yang, William W. Cohen, Ruslan Salakhutdinov
NAACL, 2018 (Short Paper)
3. Breaking the Softmax Bottleneck: A High-Rank RNN Language Model
Zhilin Yang, Zihang Dai, Ruslan Salakhutdinov, William W. Cohen
ICLR 2018, oral.
4. Active Neural Localization
Devendra Singh Chaplot, Emilio Parisotto, Ruslan Salakhutdinov
ICLR 2018.
5. Neural Map: Structured Memory for Deep Reinforcement Learning
Emilio Parisotto, Ruslan Salakhutdinov
ICLR 2018.
6. On Unifying Deep Generative Models
Zhiting Hu, Zichao Yang, Ruslan Salakhutdinov, Eric P. Xing
ICLR 2018.
7. A Generic Approach for Escaping Saddle Points
Sashank J Reddi, Manzil Zaheer, Suvrit Sra, Barnabas Poczos, Francis Bach, Ruslan Salakhutdinov, Alexander J Smola
AISTats 2018.
8. Gated-Attention Architectures for Task-Oriented Language Grounding
Devendra Singh Chaplot, Kanthashree Mysore Sathyendra, Rama Kumar Pasumarthi, Dheeraj Rajagopal, Ruslan Salakhutdinov
AAAI 2018, oral.
9. Knowledge-based Word Sense Disambiguation using Topic Models
Devendra Singh Chaplot, Ruslan Salakhutdinov
AAAI 2018, oral.
10. Good Semi-supervised Learning that Requires a Bad GAN
Zihang Dai, Zhilin Yang, Fan Yang, William W. Cohen, Ruslan Salakhutdinov
NIPS 2017.
11. Deep Sets
Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Ruslan Salakhutdinov, Alexander Smola
NIPS 2017, oral.
12. Geometry of Optimization and Implicit Regularization in Deep Learning
Behnam Neyshabur, Ryota Tomioka, Ruslan Salakhutdinov, Nathan Srebro
Survey paper.
13. Deep Conditional Determinantal Point Process for Large-Scale Multi-Label Classification
P. Xie, R. Salakhutdinov, L. Mou and E.P. Xing
ICCV 2017.
14. Learning Robust Visual-Semantic Embeddings
Yao-Hung Hubert Tsai, Liang-Kang Huang, Ruslan Salakhutdinov
ICCV 2017.
15. Improved Variational Autoencoders for Text Modeling using Dilated Convolutions
Zichao Yang, Zhiting Hu, Ruslan Salakhutdinov, Taylor Berg-Kirkpatrick
ICML 2017.

16. Controllable Text Generation
Zhiting Hu, Zichao Yang, Xiaodan Liang, Ruslan Salakhutdinov, Eric P. Xing
ICML 2017.
17. Semi-Supervised QA with Generative Domain-Adaptive Nets
Zhilin Yang, Junjie Hu, Ruslan Salakhutdinov, William W. Cohen
ACL 2017.
18. Gated-Attention Readers for Text Comprehension
Bhuwan Dhingra, Hanxiao Liu, Zhilin Yang, William W. Cohen, Ruslan Salakhutdinov
ACL 2017.
19. The More You Know: Using Knowledge Graphs for Image Classification
Kenneth Marino, Ruslan Salakhutdinov, Abhinav Gupta
CVPR 2017.
20. Spatially Adaptive Computation Time for Residual Networks
Michael Figurnov, Maxwell D. Collins, Yukun Zhu, Li Zhang, Jonathan Huang, Dmitry Vetrov, Ruslan Salakhutdinov
CVPR 2017, [arXiv], [Code].
21. Transfer Learning for Sequence Tagging with Hierarchical Recurrent Networks
Zhilin Yang, Ruslan Salakhutdinov, William W. Cohen
ICLR 2017.
22. On the Quantitative Analysis of Decoder-Based Generative Models
Yuhuai Wu, Yuri Burda, Ruslan Salakhutdinov, Roger Grosse
ICLR 2017.
23. Words or Characters? Fine-grained Gating for Reading Comprehension
Zhilin Yang, Bhuwan Dhingra, Ye Yuan, Junjie Hu, William W. Cohen, Ruslan Salakhutdinov
ICLR 2017.
24. Deep Neural Networks with Massive Learned Knowledge
Zhiting Hu, Zichao Yang, Ruslan Salakhutdinov, and Eric Xing
Conference on Empirical Methods in Natural Language Processing (EMNLP' 16).
25. Iterative Refinement of Approximate Posterior for Training Directed Belief Networks
Devon Hjelm, Ruslan Salakhutdinov, Kyunghyun Cho, Nebojsa Jojic, Vince Calhoun, Junyoung Chung
NIPS 2016.
26. Path-Normalized Optimization of Recurrent Neural Networks with ReLU Activations
Behnam Neyshabur, Yuhuai Wu, Ruslan Salakhutdinov, Nathan Srebro
NIPS 2016.
27. Stochastic Variational Deep Kernel Learning
Andrew Gordon Wilson, Zhiting Hu, Eric Xing, Ruslan Salakhutdinov
NIPS 2016.
28. On Multiplicative Integration with Recurrent Neural Networks
Yuhuai Wu, Saizheng Zhang, Ying Zhang, Yoshua Bengio, Ruslan Salakhutdinov
NIPS 2016.
29. Encode, Review, and Decode: Reviewer Module for Caption Generation
Zhilin Yang, Ye Yuan, Yuexin Wu, William W. Cohen, Ruslan Salakhutdinov
NIPS 2016.
30. Architectural Complexity Measures of Recurrent Neural Networks
Saizheng Zhang, Yuhuai Wu, Tong Che, Zhouhan Lin, Roland Memisevic, Ruslan Salakhutdinov, Yoshua Bengio
NIPS 2016.

31. Revisiting Semi-Supervised Learning with Graph Embeddings
Zhilin Yang, William Cohen, Ruslan Salakhutdinov
International Conference on Machine Learning (ICML) 2016
32. Importance Weighted Autoencoders
Yuri Burda, Roger Grosse, Ruslan Salakhutdinov
International Conference on Learning Representations (ICLR), 2016
33. Actor-Mimic: Deep Multitask and Transfer Reinforcement Learning
Emilio Parisotto, Jimmy Lei Ba, Ruslan Salakhutdinov
International Conference on Learning Representations (ICLR), 2016
34. Generating Images from Captions with Attention
Elman Mansimov, Emilio Parisotto, Jimmy Lei Ba, Ruslan Salakhutdinov
International Conference on Learning Representations (ICLR), 2016
35. Data-Dependent Path Normalization in Neural Networks
Behnam Neyshabur, Ryota Tomioka, Ruslan Salakhutdinov, Nathan Srebro
International Conference on Learning Representations (ICLR), 2016
36. Action Recognition using Visual Attention
Shikhar Sharma, Ryan Kiros, Ruslan Salakhutdinov
International Conference on Learning Representations (ICLR) workshop, 2016
37. Deep Kernel Learning
Andrew Gordon Wilson, Zhiting Hu, Ruslan Salakhutdinov, Eric Xing
AI and Statistics, 2016
38. Human-level concept learning through probabilistic program induction
Brenden Lake, Ruslan Salakhutdinov, and Joshua Tenenbaum (2015),
Science, 350(6266), 1332-1338.
39. Learning Wake-Sleep Recurrent Attention Models
Lei Jimmy Ba, Roger Grosse, Ruslan Salakhutdinov, Brendan Frey
NIPS 2015
40. Skip-Thought Vectors
Ryan Kiros, Yukun Zhu, Ruslan Salakhutdinov, Richard S. Zemel, Antonio Torralba, Raquel Urtasun, Sanja Fidler
NIPS 2015
41. Path-SGD: Path-Normalized Optimization in Deep Neural Networks
Behnam Neyshabur, Ruslan Salakhutdinov, Nathan Srebro
NIPS 2015
42. Aligning Books and Movies: Towards Story-like Visual Explanations by Watching Movies and Reading Books
Yukun Zhu, Ryan Kiros, Richard Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, Sanja Fidler
ICCV 2015
43. Predicting Deep Zero-Shot Convolutional Neural Networks using Textual Descriptions
Jimmy Ba, Kevin Swersky, Sanja Fidler, Ruslan Salakhutdinov
ICCV 2015
44. Learning Deep Generative Models
Ruslan Salakhutdinov
Annual Review of Statistics and Its Application, Vol. 2, pp. 361-385, 2015
45. Roger Grosse and Ruslan Salakhutdinov
Scaling up Natural Gradient by Sparsely Factorizing the Inverse Fisher Matrix
International Conference on Machine Learning (ICML 2015)

46. Nitish Srivastava, Elman Mansimov, Ruslan Salakhutdinov
Unsupervised Learning of Video Representations using LSTMs
International Conference on Machine Learning (ICML 2015)
47. Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhutdinov, Richard Zemel, Yoshua Bengio
Show, Attend and Tell: Neural Image Caption Generation with Visual Attention
International Conference on Machine Learning (ICML 2015)
48. Ruslan Salakhutdinov
Learning Deep Generative Models
Annual Review of Statistics and Its Application, Vol. 2, 2015
49. Y. Zhu, R. Urtasun, R. Salakhutdinov and S.Fidler
segDeepM: Exploiting Segmentation and Context in Deep Neural Networks for Object Detection
Conference on Computer Vision and Pattern Recognition (CVPR), Boston, MA, USA, June 2015,
50. Yuri Burda, Roger B. Grosse, and Ruslan Salakhutdinov
Accurate and Conservative Estimates of MRF Log-likelihood using Reverse Annealing
AI and Statistics (AISTATS), 2015.
51. Ryan Kiros, Ruslan Salakhutdinov, Richard Zemel
Unifying Visual-Semantic Embeddings with Multimodal Neural Language Models
Transactions of the Association for Computational Linguistics (TACL), 2015.
52. Yichuan Tang, Nitish Srivastava, and Ruslan Salakhutdinov
Learning Generative Models with Visual Attention
Neural Information Processing Systems (NIPS 28), **oral**.
53. Ryan Kiros, Richard Zemel, Ruslan Salakhutdinov
A Multiplicative Model for Learning Distributed Text-Based Attribute Representations
Neural Information Processing Systems (NIPS 28).
54. Nitish Srivastava and Ruslan Salakhutdinov
Multimodal Learning with Deep Boltzmann Machines
Journal of Machine Learning Research (JMLR), 2014.
55. S. Plis, D. Hjelm, R. Salakhutdinov, E. Allen, H. Bockholt, J. Long, H. Johnson, J. Paulsen, J. Turner, and V. Calhoun
Deep Learning for Neuroimaging: a Validation Study
Frontiers in Neuroscience, 2014.
56. D Hjelm, V Calhoun, EA Allen, T Adali, R Salakhutdinov, S Plis.
Restricted Boltzmann Machines for Neuroimaging: an Application in Identifying Intrinsic Networks
NeuroImage, in Press
57. Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, Ruslan Salakhutdinov
Dropout: A simple way to prevent neural networks from overfitting
Journal of Machine Learning Research (JMLR), 2014.
58. Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov
Multimodal Neural Language Models
International Conference on Machine Learning (ICML 2014)
59. Roger Grosse, Chris Maddison, and Ruslan Salakhutdinov
Annealing between Distributions by Averaging Moments
In Neural Information Processing Systems (NIPS 27), **oral**.
60. Nitish Srivastava and Ruslan Salakhutdinov
Discriminative Transfer Learning with Tree-based Priors
In Neural Information Processing Systems (NIPS 27).

61. Yichuan Tang and Ruslan Salakhutdinov
Learning Stochastic Feedforward Neural Networks
In Neural Information Processing Systems (NIPS 27).
62. Brenden Lake, Ruslan Salakhutdinov, and Josh Tenenbaum
One-shot Learning by Inverting a Compositional Causal Process
In Neural Information Processing Systems (NIPS 27).
63. B. Neyshabur, N. Srebro, R. Salakhutdinov, Y. Makarychev, and P. Yadollahpour
The Power of Asymmetry in Binary Hashing
In Neural Information Processing Systems (NIPS 27).
64. Ruslan Salakhutdinov, Josh Tenenbaum, and Antonio Torralba
Learning with Hierarchical-Deep Models
In IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), 2013
65. Nitish Srivastava Ruslan Salakhutdinov, and Geoffrey Hinton
Modeling Documents with Deep Boltzmann Machines
In Uncertainty in Artificial Intelligence (UAI) 2013, **oral**.
66. Yichuan Tang, Ruslan Salakhutdinov, and Geoffrey Hinton (2013)
Tensor Analyzers
The 29th International Conference on Machine Learning (ICML 2013).
67. Nitish Srivastava and Ruslan Salakhutdinov
Multimodal Learning with Deep Boltzmann Machines
In Neural Information Processing Systems (NIPS 26), **oral**.
68. Mohammad Norouzi, David Fleet, and Ruslan Salakhutdinov
Hamming Distance Metric Learning
In Neural Information Processing Systems (NIPS 26)
69. Ruslan Salakhutdinov and Geoffrey Hinton
A Better Way to Pretrain Deep Boltzmann Machines
In Neural Information Processing Systems (NIPS 26)
70. Rina Foygel, Nathan Srebro, Ruslan Salakhutdinov
Matrix reconstruction with the local max norm
In Neural Information Processing Systems (NIPS 26)
71. Kevin Swersky, Daniel Tarlow, Ilya Sutskever, Ruslan Salakhutdinov, Richard Zemel, and Ryan Adams
Cardinality Restricted Boltzmann Machines
In Neural Information Processing Systems (NIPS 26)
72. Roger Grosse, Ruslan Salakhutdinov, William Freeman, and Joshua Tenenbaum (2012)
Exploiting Compositionality to Explore a Large Space of Model Structures
In Uncertainty in Artificial Intelligence (UAI) 2012, **Best student paper award**
73. Ruslan Salakhutdinov and Geoffrey Hinton
Efficient Learning of Deep Boltzmann Machines
Neural Computation, August 2012, Vol. 24, No. 8: 1967-2006.
74. Ruslan Salakhutdinov, Josh Tenenbaum, and Antonio Torralba
One-Shot Learning with a Hierarchical Nonparametric Bayesian Model
Journal of Machine Learning Research (JMLR) WC&P Unsupervised and Transfer Learning, 2012,
75. Yichuan Tang , Ruslan Salakhutdinov, and Geoffrey Hinton
Deep Lambertian Networks
The 29th International Conference on Machine Learning (ICML 2012)

76. Yichuan Tang , Ruslan Salakhutdinov, and Geoffrey Hinton
Deep Mixtures of Factor Analyzers
The 29th International Conference on Machine Learning (ICML 2012)
77. Brenden Lake , Ruslan Salakhutdinov, and Josh Tenenbaum
Concept learning as motor program induction: A large-scale empirical study
Proceedings of the 34rd Annual Conference of the Cognitive Science Society, 2012
78. Yichuan Tang , Ruslan Salakhutdinov, and Geoffrey Hinton
Robust Boltzmann Machines for Recognition and Denoising
IEEE Computer Vision and Pattern Recognition (CVPR) 2012
79. Yaodong Zhang, Ruslan Salakhutdinov, Hung-An Chang, and James Glass
Resource Configurable Spoken Query Detection using Deep Boltzmann Machines
37th International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2012)
80. Dean Foster, Sham Kakade, and Ruslan Salakhutdinov
Domain Adaptation: A Small Sample Statistical Approach
Journal of Machine Learning Research W&CP 15 (AISTATS 2012)
81. Ruslan Salakhutdinov, Josh Tenenbaum , Antonio Torralba
Learning to Learn with Compound Hierarchical-Deep Models
Advances in Neural Information Processing Systems 25, (NIPS 25)
82. Joseph Lim , Ruslan Salakhutdinov Antonio Torralba
Transfer Learning by Borrowing Examples
Advances in Neural Information Processing Systems 25, (NIPS 25)
83. Rina Foygel, Ruslan Salakhutdinov, Ohad Shamir, Nathan Srebro
Learning with the Weighted Trace-norm under Arbitrary Sampling Distributions
Advances in Neural Information Processing Systems 25, (NIPS 25)
84. Brenden Lake , Ruslan Salakhutdinov, Jason Gross, and Josh Tenenbaum (2011)
One-shot Learning of Simple Visual Concepts
Proceedings of the 33rd Annual Conference of the Cognitive Science Society,
85. Ruslan Salakhutdinov, Antonio Torralba, and Josh Tenenbaum
Learning to Share Visual Appearance for Multiclass Object Detection
IEEE Computer Vision and Pattern Recognition (CVPR) 2011
86. Ruslan Salakhutdinov and Nathan Srebro
Collaborative Filtering in a Non-Uniform World: Learning with the Weighted Trace Norm
Advances in Neural Information Processing Systems 24, (NIPS 24)
87. Jason Lee, Ben Recht, Ruslan Salakhutdinov, Nathan Srebro, and Joel Tropp (2011)
Practical Large-Scale Optimization for Max-Norm Regularization
Advances in Neural Information Processing Systems 24, (NIPS 24)
88. Geoffrey Hinton and Ruslan Salakhutdinov
Discovering Binary Codes for Documents by Learning Deep Generative Models
Topics in Cognitive Science, 2010.
89. Ruslan Salakhutdinov
Learning in Deep Boltzmann Machines using Adaptive MCMC
In 27th International Conference on Machine Learning (ICML 2010), Haifa, Israel.
90. Ruslan Salakhutdinov
Learning in Markov Random Fields using Tempered Transitions
Advances in Neural Information Processing Systems 23 (NIPS 23), Vancouver, Canada.

91. Ruslan Salakhutdinov and Geoffrey Hinton
Replicated Softmax: an Undirected Topic Model
Advances in Neural Information Processing Systems 23 (NIPS 23), Vancouver, Canada.
92. Ilya Sutskever, Ruslan Salakhutdinov, and Josh Tenenbaum
Modelling Relational Data using Bayesian Clustered Tensor Factorization
Advances in Neural Information Processing Systems 23 (NIPS 23), Vancouver, Canada.
93. Ruslan Salakhutdinov and Hugo Larochelle
Efficient Learning of Deep Boltzmann Machines
13th International Conference on Artificial Intelligence and Statistics, AISTATS 2010, Sardinia, Italy.
94. Ruslan Salakhutdinov and Geoffrey Hinton
Semantic Hashing
International Journal of Approximate Reasoning, 2009.
95. John Langford, Ruslan Salakhutdinov and Tong Zhang
Learning Nonlinear Dynamic Models
26th International Conference on Machine Learning (ICML 2009), Montreal, Canada.
96. Hanna M. Wallach, Iain Murray, Ruslan Salakhutdinov and David Mimno
Evaluation Methods for Topic Models
26th International Conference on Machine Learning (ICML 2009), Montreal, Canada.
97. Ruslan Salakhutdinov and Geoffrey Hinton
Deep Boltzmann Machines
12th International Conference on Artificial Intelligence and Statistics, AISTATS 2009, Clearwater, Florida.
98. Iain Murray and Ruslan Salakhutdinov
Evaluating probabilities under high-dimensional latent variable models
Advances in Neural Information Processing Systems 22 (NIPS 22), Vancouver, Canada.
99. Ruslan Salakhutdinov & Andriy Mnih
Bayesian Probabilistic Matrix Factorization using MCMC
25th International Conference on Machine Learning (ICML 2008), Helsinki, Finland.
100. Ruslan Salakhutdinov & Iain Murray
On the Quantitative Analysis of Deep Belief Networks
25th International Conference on Machine Learning (ICML 2008), Helsinki, Finland.
101. Ruslan Salakhutdinov & Andriy Mnih
Probabilistic Matrix Factorization
Advances in Neural Information Processing Systems 21 (NIPS 21), Vancouver, Canada, **oral**.
102. Ruslan Salakhutdinov & Geoffrey Hinton
Using Deep Belief Nets to Learn Covariance Kernels for Gaussian Processes
Advances in Neural Information Processing Systems 21 (NIPS 21), Vancouver, Canada.
103. Ruslan Salakhutdinov, Andriy Mnih, & Geoffrey Hinton
Restricted Boltzmann Machines for Collaborative Filtering
24th International Conference on Machine Learning (ICML 2007), Corvallis, Oregon, USA.
104. Ruslan Salakhutdinov & Geoffrey Hinton
Learning a Nonlinear Embedding by Preserving Class Neighbourhood Structure
11th International Conference on Artificial Intelligence and Statistics, AISTATS 2007, San Juan, Puerto Rico.
105. Geoffrey Hinton & Ruslan Salakhutdinov
Reducing the Dimensionality of Data with Neural Networks
SCIENCE 28 July 2006: Vol. 313. no. 5786, pp. 504 - 507.

106. Sam Roweis & Ruslan Salakhutdinov
Simultaneous Localization and Surveying with Multiple Agents
In R. Murray-Smith, R. Shorten (eds), *Switching and Learning in Feedback Systems*
(Springer LNCS vol 3355). pp. 313–332, 2005.
107. Jacob Goldberger, Sam Roweis, Geoffrey Hinton, Ruslan Salakhutdinov (2005)
Neighbourhood Component Analysis
Advances in Neural Information Processing Systems 18 (NIPS 18), Vancouver, Canada.
108. Grigoris Karakoulas & Ruslan Salakhutdinov
Semi-Supervised Mixture-of-Experts Classification
The Fourth IEEE International Conference on Data Mining, ICDM 2004 Brighton, UK.
109. Ruslan Salakhutdinov, Sam Roweis & Zoubin Ghahramani
Optimization with EM and Expectation-Conjugate-Gradient
International Conference on Machine Learning (ICML 2003), Washington DC, USA.
110. Ruslan Salakhutdinov & Sam Roweis
Adaptive Overrelaxed Bound Optimization Methods
International Conference on Machine Learning (ICML 2003), Washington DC, USA.
111. Ruslan Salakhutdinov, Sam Roweis & Zoubin Ghahramani
On the Convergence of Bound Optimization Algorithms
Uncertainty in Artificial Intelligence (UAI 2003), Acapulco, Mexico.

Refereed workshop publications:

1. Devendra Singh Chaplot, Guillaume Lample, Kanthashree Mysore Sathyendra, Ruslan Salakhutdinov
Transfer Deep Reinforcement Learning in 3D Environments: An Empirical Study
Deep Reinforcement Learning Workshop, NIPS 2016
2. Ryan Kiros, Ruslan Salakhutdinov, Richard Zemel
Unifying Visual-Semantic Embeddings with Multimodal Neural Language Models
NIPS Deep Learning Workshop, 2014.
3. Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov (2013)
Multimodal Neural Language Models
In NIPS 2013 Deep Learning Workshop
4. Nitish Srivastava and Ruslan Salakhutdinov (2012)
Learning Representations for Multimodal Data with Deep Belief Nets
ICML workshop on Representation Learning, 2012
5. Ruslan Salakhutdinov (2009)
Learning Feature Hierarchies by Learning Generative Models
NIPS 2009 Workshop on the Generative and Discriminative Learning Interface. Dec 12, 2009.
6. Ruslan Salakhutdinov (2009)
Undirected Topic Models
NIPS 2009 Workshop on Applications for Topic Models: Text and Beyond Dec 11, 2009.
7. Ruslan Salakhutdinov & Geoffrey Hinton (2009)
Learning Deep Boltzmann Machines
The Snowbird learning workshop, Clearwater, Florida.
8. Hanna M. Wallach, Iain Murray, Ruslan Salakhutdinov and David Mimno (2009)
Evaluation Methods for Topic Models
The Snowbird learning workshop, Clearwater, Florida.

9. Ruslan Salakhutdinov & Geoffrey Hinton (2007)
Semantic Hashing
Proceedings of the SIGIR Workshop on Information Retrieval and Applications of Graphical Models, Amsterdam.
10. Ruslan Salakhutdinov & Geoffrey Hinton (2007)
Deep Belief Networks
The Snowbird learning workshop, San Juan, Puerto Rico.
11. Ruslan Salakhutdinov & Geoffrey Hinton (2006)
Nonlinear Dimensionality Reduction
NIPS 2006 workshop on Novel Applications of Dimensionality Reduction, Dec 2006.

Unrefereed technical reports

1. Ruslan Salakhutdinov and Geoffrey Hinton (2010)
An Efficient Learning Procedure for Deep Boltzmann Machines
MIT Technical Report MIT-CSAIL-TR-2010-037 (submitted for journal publication)
2. Ruslan Salakhutdinov, Josh Tenenbaum, and Antonio Torralba (2010)
One-Shot Learning with a Hierarchical Nonparametric Bayesian Model
MIT Technical Report MIT-CSAIL-TR-2010 (submitted for journal publication)
3. Ruslan Salakhutdinov (2009)
Learning Deep Generative Models
PhD Thesis, Sep 2009, Dept. of Computer Science, University of Toronto
4. Ruslan Salakhutdinov (2008)
Learning and Evaluating Boltzmann Machines
Technical Report UTML TR 2008-002, Dept. of Computer Science, University of Toronto
5. Iain Murray and Ruslan Salakhutdinov (2008)
Notes on the KL-divergence between a Markov chain and its equilibrium distribution
Technical Report UTML TR 2008, Dept. of Computer Science, University of Toronto
6. Ruslan Salakhutdinov, Sam Roweis, and Zoubin Ghahramani (2002)
Relationship between gradient and EM steps in latent variable models
Technical Report, University of Toronto.
7. Ruslan Salakhutdinov, Sam Roweis, and Zoubin Ghahramani (2002)
Expectation Conjugate-Gradient: An Alternative to EM
Technical Report, University of Toronto.