

RUSLAN SALAKHUTDINOV

100 St. George Street, Sidney Smith Hall, Room 6002,
Toronto, Ontario M5S 3G3, Canada
rsalaku@utstat.toronto.edu
<http://www.utstat.toronto.edu/~rsalaku/>

Academic Career

Assistant Professor, University of Toronto
Department of Statistics and Computer Science August 2011 - Present

Postdoctoral Research Associate, Brain and Cognitive Sciences (BCS)
and Computer Science and Artificial Intelligence Lab (CSAIL), MIT. Sep. 2009 - July 2011
Advised by Josh Tenenbaum.

PhD, Department of Computer Science, University of Toronto. Sep. 2005 - Aug. 2009
Thesis: Learning Deep Generative Models, advised by Geoffrey Hinton.

Master of Science, Department of Computer Science, University of Toronto. Sep. 2001 - Aug. 2003
Thesis: Optimization Algorithms for Learning, advised by Sam Roweis.

Bachelor of Science, High Point University, NC, USA. Aug. 1998 - May. 2001
Double major in Computer Science and Mathematics, Honors Degree.

Teaching Experience

- STAD 37H, Statistical Multivariate Analysis. Winter 2012
- STA 4273H, Research Topics In Statistical Machine Learning. Fall 2011
- 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

Yahoo Research, New York, USA. Summer Intern. Jun. 2008 - Aug 2008

Canadian Imperial Bank of Commerce (CIBC), Toronto, Canada. Sep. 2003 - Aug 2005
Sr. Risk Analyst, Risk Management.

Grants, Scholarships & Awards

- Connaught New Researcher Award (2012 - 2014)
- Early Researcher Award (2012-2017)
- NSERC Individual Discovery Grant, (2012-2017)
along with NSERC Early Career Researcher Supplement
- Canadian Institute for Advanced Research (2011-2014)
Scholar of the Neural Computation and Adaptive Perception Program.
- Natural Sciences and Engineering Research Council of Canada:
Postdoctoral Fellowship. (CAD \$80,000, 2009-2011)
- The UK Engineering and Physical Sciences Research Council:
Postdoctoral Fellowship in Theoretical Computer Science (Declined): (£230,000, 2009-2012)

- Natural Sciences and Engineering Research Council of Canada:
Canada Graduate Scholarship. (CAD \$105,000, 2006-2009)
- Ontario Graduate Scholarship (Declined). (CAD \$15,000, 2006-2007)
- Precarn Scholar, Canada. (CAD \$4,300, 2002-2003)
- Outstanding Computer Science Major, High Point University. (US \$500, 2001)
- Member of Alpha Chi Honor Society (top 5% in class), High Point University.
- International Phi Theta Kappa Full Tuition Scholarship (US \$33,000, 1998-2001)
High Point University, NC, USA.
- Youth For Understanding Fellowship (US \$26,000, 1997-1998)
Sponsored by the US government, Washington DC, USA.

Refereed scientific publications:

1. Ruslan Salakhutdinov and Geoffrey Hinton (2012)
Efficient Learning of Deep Boltzmann Machines
To appear in Neural Computation.
2. Ruslan Salakhutdinov, Josh Tenenbaum, and Antonio Torralba (2012)
One-Shot Learning with a Hierarchical Nonparametric Bayesian Model
Journal of Machine Learning Research (JMLR) WC&P Unsupervised and Transfer Learning, 2012,
3. Yichuan Tang , Ruslan Salakhutdinov, and Geoffrey Hinton (2012)
Deep Lambertian Networks
The 29th International Conference on Machine Learning (ICML 2012)
4. Yichuan Tang , Ruslan Salakhutdinov, and Geoffrey Hinton (2012)
Deep Mixtures of Factor Analysers
The 29th International Conference on Machine Learning (ICML 2012)
5. Brenden Lake , Ruslan Salakhutdinov, and Josh Tenenbaum (2012)
Concept learning as motor program induction: A large-scale empirical study.
Proceedings of the 34rd Annual Conference of the Cognitive Science Society, 2012
6. Yichuan Tang , Ruslan Salakhutdinov, and Geoffrey Hinton (2012)
Robust Boltzmann Machines for Recognition and Denoising
IEEE Computer Vision and Pattern Recognition (CVPR) 2012
7. Yaodong Zhang, Ruslan Salakhutdinov, Hung-An Chang, and James Glass (2012)
Resource Configurable Spoken Query Detection using Deep Boltzmann Machines
37th International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2012)
8. Dean Foster, Sham Kakade, and Ruslan Salakhutdinov (2012)
Domain Adaptation: A Small Sample Statistical Approach
Journal of Machine Learning Research W&CP 15 (AISTATS 2012)
9. Ruslan Salakhutdinov, Josh Tenenbaum , Antonio Torralba (2012)
Learning to Learn with Compound Hierarchical-Deep Models
Advances in Neural Information Processing Systems 25, (NIPS 25)
10. Joseph Lim , Ruslan Salakhutdinov Antonio Torralba (2012)
Transfer Learning by Borrowing Examples
Advances in Neural Information Processing Systems 25, (NIPS 25)
11. Rina Foygel, Ruslan Salakhutdinov, Ohad Shamir, Nathan Srebro (2012)
Learning with the Weighted Trace-norm under Arbitrary Sampling Distributions
Advances in Neural Information Processing Systems 25, (NIPS 25)

12. 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,
13. Ruslan Salakhutdinov, Antonio Torralba, and Josh Tenenbaum (2011)
Learning to Share Visual Appearance for Multiclass Object Detection
IEEE Computer Vision and Pattern Recognition (CVPR) 2011
14. Ruslan Salakhutdinov and Nathan Srebro (2011)
Collaborative Filtering in a Non-Uniform World: Learning with the Weighted Trace Norm
Advances in Neural Information Processing Systems 24, (NIPS 24)
15. 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)
16. Geoffrey Hinton and Ruslan Salakhutdinov (2010)
Discovering Binary Codes for Documents by Learning Deep Generative Models
Topics in Cognitive Science.
17. Ruslan Salakhutdinov (2010)
Learning in Deep Boltzmann Machines using Adaptive MCMC
In 27th International Conference on Machine Learning (ICML 2010), Haifa, Israel.
18. Ruslan Salakhutdinov (2010)
Learning in Markov Random Fields using Tempered Transitions
Advances in Neural Information Processing Systems 23 (NIPS 23), Vancouver, Canada.
19. Ruslan Salakhutdinov and Geoffrey Hinton (2010)
Replicated Softmax: an Undirected Topic Model
Advances in Neural Information Processing Systems 23 (NIPS 23), Vancouver, Canada.
20. Ilya Sutskever, Ruslan Salakhutdinov, and Josh Tenenbaum (2010)
Modelling Relational Data using Bayesian Clustered Tensor Factorization
Advances in Neural Information Processing Systems 23 (NIPS 23), Vancouver, Canada.
21. Ruslan Salakhutdinov and Hugo Larochelle (2010)
Efficient Learning of Deep Boltzmann Machines
13th International Conference on Artificial Intelligence and Statistics, AISTATS 2010, Sardinia, Italy.
22. Ruslan Salakhutdinov and Geoffrey Hinton (2009)
Semantic Hashing
International Journal of Approximate Reasoning.
23. John Langford, Ruslan Salakhutdinov and Tong Zhang (2009)
Learning Nonlinear Dynamic Models
26th International Conference on Machine Learning (ICML 2009), Montreal, Canada.
24. Hanna M. Wallach, Iain Murray, Ruslan Salakhutdinov and David Mimno (2009)
Evaluation Methods for Topic Models
26th International Conference on Machine Learning (ICML 2009), Montreal, Canada.
25. Ruslan Salakhutdinov and Geoffrey Hinton (2009)
Deep Boltzmann Machines
12th International Conference on Artificial Intelligence and Statistics, AISTATS 2009, Clearwater, Florida.
26. Iain Murray and Ruslan Salakhutdinov (2009)
Evaluating probabilities under high-dimensional latent variable models
Advances in Neural Information Processing Systems 22 (NIPS 22), Vancouver, Canada.

27. Ruslan Salakhutdinov & Andriy Mnih (2008)
Bayesian Probabilistic Matrix Factorization using MCMC
25th International Conference on Machine Learning (ICML 2008), Helsinki, Finland.
28. Ruslan Salakhutdinov & Iain Murray (2008)
On the Quantitative Analysis of Deep Belief Networks
25th International Conference on Machine Learning (ICML 2008), Helsinki, Finland.
29. Ruslan Salakhutdinov & Andriy Mnih (2008)
Probabilistic Matrix Factorization
Advances in Neural Information Processing Systems 21 (NIPS 21), Vancouver, Canada.
30. Ruslan Salakhutdinov & Geoffrey Hinton (2008)
Using Deep Belief Nets to Learn Covariance Kernels for Gaussian Processes
Advances in Neural Information Processing Systems 21 (NIPS 21), Vancouver, Canada.
31. Ruslan Salakhutdinov, Andriy Mnih, & Geoffrey Hinton (2007)
Restricted Boltzmann Machines for Collaborative Filtering
24th International Conference on Machine Learning (ICML 2007), Corvallis, Oregon, USA.
32. Ruslan Salakhutdinov & Geoffrey Hinton (2007)
Learning a Nonlinear Embedding by Preserving Class Neighbourhood Structure
11th International Conference on Artificial Intelligence and Statistics, AISTATS 2007, San Juan, Puerto Rico.
33. Geoffrey Hinton & Ruslan Salakhutdinov (2006)
Reducing the Dimensionality of Data with Neural Networks
SCIENCE 28 July 2006: Vol. 313. no. 5786, pp. 504 - 507.
34. Sam Roweis & Ruslan Salakhutdinov (2005)
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.
35. Jacob Goldberger, Sam Roweis, Geoffrey Hinton, Ruslan Salakhutdinov (2005)
Neighbourhood Component Analysis
Advances in Neural Information Processing Systems 18 (NIPS 18), Vancouver, Canada.
36. Grigoris Karakoulas & Ruslan Salakhutdinov (2004)
Semi-Supervised Mixture-of-Experts Classification
The Fourth IEEE International Conference on Data Mining, ICDM 2004 Brighton, UK.
37. Ruslan Salakhutdinov, Sam Roweis & Zoubin Ghahramani (2003)
Optimization with EM and Expectation-Conjugate-Gradient
International Conference on Machine Learning (ICML 2003), Washington DC, USA.
38. Ruslan Salakhutdinov & Sam Roweis (2003)
Adaptive Overrelaxed Bound Optimization Methods
International Conference on Machine Learning (ICML 2003), Washington DC, USA.
39. Ruslan Salakhutdinov, Sam Roweis & Zoubin Ghahramani (2003)
On the Convergence of Bound Optimization Algorithms
Uncertainty in Artificial Intelligence (UAI 2003), Acapulco, Mexico.

Refereed workshop publications:

1. Nitish Srivastava and Ruslan Salakhutdinov (2012)
Learning Representations for Multimodal Data with Deep Belief Nets
ICML workshop on Representation Learning, 2012

2. Ruslan Salakhutdinov (2009)
Learning Feature Hierarchies by Learning Generative Models
NIPS 2009 Workshop on the Generative and Discriminative Learning Interface. Dec 12, 2009.
3. Ruslan Salakhutdinov (2009)
Undirected Topic Models
NIPS 2009 Workshop on Applications for Topic Models: Text and Beyond Dec 11, 2009.
4. Ruslan Salakhutdinov & Geoffrey Hinton (2009)
Learning Deep Boltzmann Machines
The Snowbird learning workshop, Clearwater, Florida.
5. Hanna M. Wallach, Iain Murray, Ruslan Salakhutdinov and David Mimno (2009)
Evaluation Methods for Topic Models
The Snowbird learning workshop, Clearwater, Florida.
6. Ruslan Salakhutdinov & Geoffrey Hinton (2007)
Semantic Hashing
Proceedings of the SIGIR Workshop on Information Retrieval and Applications of Graphical Models, Amsterdam.
7. Ruslan Salakhutdinov & Geoffrey Hinton (2007)
Deep Belief Networks
The Snowbird learning workshop, San Juan, Puerto Rico.
8. 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.

Invited Talks

- Guest Lecture, IPAM summer school, UCLA. July, 2012
- ISBA World Meeting, Special Topic Session on Adaptive Monte Carlo, Kyoto, Japan. June, 2012
- Tutorial Speaker, CVPR, Providence, Rhode Island, Deep Learning Tutorial. June, 2012
- Workshop on Perspectives on High-dimensional Data Analysis, University of Montreal. May, 2012
- Workshop on Statistical Machine Learning for Speech Processing, Kyoto, Japan. April, 2012
- Tokyo Institute of Technology. March, 2012
- Seminar Series in Computational Statistics. University of Guelph. March, 2012
- NIPS 2011 Workshop on Beyond Mahalanobis: Supervised Large-Scale Learning of Similarity. Dec, 2011
- Contributed talk, NIPS 2011 Workshop on Domain Adaptation: Theory and Applications. Dec, 2011
- CIFAR NCAP Workshop, Granada, Spain. Dec, 2011
- 2011 Symposium on Advances in Intelligent Systems, University of Waterloo. Dec 2011
- CIFAR Summer School, University of Toronto August, 2011
- IPAM Graduate Summer School: Probabilistic Models of Cognition, UCLA July, 2011
- ICML 2011 Workshop on Unsupervised and Transfer Learning July, 2011
- Workshop on Infusing Statistics and Engineering, Harvard University June, 2011
- Department of Statistics, Stanford April, 2011
- Machine Learning Department, CMU April, 2011
- EECS, MIT March, 2011
- Computer Science Department, Princeton University March, 2011
- Computer Science Department, Stanford March, 2011
- Department of Statistics, University of Toronto March, 2011
- Department of Computer Science, Johns Hopkins University March, 2011
- Toyota Technological Institute at Chicago February, 2011
- Department of Computer Science, University of Rochester February, 2011
- Computer Science Department, UMass Amherst February, 2011
- Computer Science & Engineering, University of Washington February, 2011
- Computer Science, Harvard University February, 2011
- Brown University, Pattern Theory Group, Division of Applied Mathematics October, 2010
- Cornell University, AI seminar October, 2010
- The 43rd Annual Meeting of the Society for Mathematical Psychology (invited talk) August 2010
- Cognitive and Neural Models for Automated Processing of Speech and Text (CONAS 2010) invited workshop seminar, Ghent, Belgium. July 2010
- Harvard University, Computational Neuroscience Seminar, Center for Brain Science May 2010
- Harvard University, 24th New England Statistics Symposium April 2010
- Cognitive Machines Group, MIT Media Lab April 2010
- University of Massachusetts, Amherst, Machine Learning Lunch March 2010
- Toyota Technological Institute at Chicago (invited seminar) January 2010
- NIPS 2009 Workshop on Approximate Learning of Large Scale Graphical Models December 2009
- NIPS 2009 Workshop on the Generative and Discriminative Learning (contributed talk) December 2009
- Stochastic Systems Group, LIDS, MIT December 2009

- CBCL Seminar, Brain and Cognitive Sciences, MIT October 2009
- Vision Seminar, CSAIL, MIT October 2009
- Snowbird workshop (contributed talk) April 2009
- Microsoft Research, Redmond March 2009
- NIPS 2008 Workshop on Approximate inference - How far have we come? (invited talk) December 2008
- MIT Computer Science and AI Laboratory November 2008
- Yahoo Research Labs, New York August 2008
- Google Labs, New York February 2008
- Institute for Pure and Applied Mathematics, UCLA (invited talk) October 2007
- MS-MITACS Joint Conference (invited seminar) June 2007
- Snowbird workshop (contributed talk) April 2007
- NIPS06 workshop on Novel Applications of Dimensionality Reduction (invited talk) December 2006

Program Committee Member or Chair

- Area Chair for the NIPS 2012 program committee
- Guest editor for IEEE Transactions on Pattern Analysis and Machine Intelligence, Special Issue on Learning Deep Architectures
- Workshop Chair for the UAI 2012
- Area Chair for the ICML 2012 program committee
- NIPS 2011 Workshop on Challenges in Learning Hierarchical Models: Transfer Learning and Optimization,
- Area Chair for the NIPS 2011 program committee
- Area Chair for the ICML 2011 program committee, Deep Learning, Optimization Algorithms, Recommendation and Matrix Factorization
- NIPS Workshop on Transfer Learning Via Rich Generative Models, 2010, Co-chair
- NIPS Workshop on Deep Learning, 2010, program committee
- NIPS Workshop on Approximate Learning of Large Scale Graphical Models, 2009, Co-chair
- ICML Workshop on Learning Feature Hierarchies, 2009, Co-chair
- International Conf. on Artificial Intelligence and Statistics, Program Committee (AISTATS 2007,2009,2010)
- International Conf. on Machine Learning, Program Committee (ICML 2007, 2008, 2009, 2010), Session Chair (2010)
- International Joint Conf. on Artificial Intelligence, Program Committee (IJCAI 2009)
- NIPS mini-symposium, Deep Learning: Foundations and Future Directions, 2007, Co-chair

Reviewing Activity

- Journal of Machine Learning Research (JMLR), 2006, 2008, 2009, 2010, 2011
- Artificial Intelligence, 2011, 2012
- Machine Learning Journal, 2006, 2007, 2009
- Statistics and Computing, 2010
- Science, 2009
- The Journal of Computer and System Sciences (JCSS), 2010
- IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI), 2005, 2006, 2008, 2011
- IEEE Transactions on Signal Processing, 2007
- Neural Computation, 2006, 2007, 2008
- Advances in Neural Information Processing Systems (NIPS), 2006, 2007, 2008, 2009, 2010