

Dept. of Computer Science, University of Toronto
6 King's College Rd.
Toronto, ON, M5S 3G4, Canada.
Tel: (416) 978-7391.
Fax: (416) 978-1455.
rgrosse@cs.toronto.edu
<http://www.cs.toronto.edu/~rgrosse/>
Citizenship: USA

1 Education

- Ph.D., Computer Science, 2008–2014
Massachusetts Institute of Technology
Supervised by William Freeman
Thesis: *Model Selection in Compositional Spaces*.
- M.S., Computer Science, 2007–2008
B.S., Symbolic Systems, 2003–2007
Stanford University
Graduated with distinction and honors

2 Recognitions

- Ontario Ministry of Research, Innovation and Science (MRIS) Early Researcher Award (ERA), 2018.
- Canada Research Chair in Probabilistic Inference and Deep Learning, 2017.
- Best Student Paper, Conference on Uncertainty in Artificial Intelligence (UAI), 2012.
- Best Student Presentation, Rank Prize Summit on Machine Learning and Computer Vision, 2012.
- National Defense Science and Engineering Grant (NDSEG), 2009–2012.
- Best Application Paper, International Conference on Machine Learning (ICML), 2009.

3 Employment History

- Assistant Professor, July 2016–present
University of Toronto Dept. of Computer Science
- Postdoctoral Fellow, 2014–2016
University of Toronto Dept. of Computer Science
Supervised by Ruslan Salakhutdinov
- Research Intern, 2010
Microsoft Research Cambridge

4 Research Funding History

- Title: Neural Networks for Uncertainty Modeling and Exploration
Source: Ministry of Research, Innovation and Science (MRIS) Early Researcher Award (ERA)
Years of Tenure: 2018–2023
Total Amount: \$190,000
- Canada Research Chair in Probabilistic Inference and Deep Learning
Source: NSERC Canada Research Chair, Tier 2
Years of Tenure: 2017–2022
Total Amount: \$500,000
- Title: Scalable and Flexible Bayesian Learning
Source: Connaught Fund
Role: PI
Years of Tenure: 2017-2018
Total Amount: \$33,685
- Title: Evaluating and Improving Deep Neural Networks
Source: NSERC Discovery Grant
Role: PI
Years of Tenure: 2017-2022
Total Amount: \$155,000
- Title: Neural Network Optimization and Structure Search
Source: Google
Role: PI
Total Amount: \$65,000
- Title: Sequential Model-Based Optimization for Deep Learning
Source: Defense Advanced Research Projects Agency (DARPA)
Role: co-PI
Status: declined
Total Amount: \$800,000
- Title: Deep Reversible Networks
Source: University of Toronto Excellence Awards
Role: PI
Years of Tenure: 2017
Total Amount: \$6,000
- University of Toronto Startup Funds
Total Amount: \$175,000
- Banting Postdoctoral Fellowship
Years of Tenure: 2015-2016
Total Amount: \$70,000
- Marsden Postdoctoral Fellowship
Years of Tenure: 2015

Total Amount: \$25,000

5 Activities

5.1 Supervision

5.1.1 Current Graduate Students

1. Yuhuai Wu
PhD, 2016–present
2. James Lucas
MSc, 2017–present
Co-supervised by Richard Zemel
3. Shengyang Sun
PhD, 2017–present
4. Guodong Zhang
MSc, 2017–present
Co-supervised by David Duvenaud
5. Jesse Bettencourt
MSc, 2017–present
Co-supervised by David Duvenaud
6. William Saunders
PhD, 2017–present
7. Matthew MacKay
MSc, 2017–present
Co-supervised by David Duvenaud
8. Paul Vicol
PhD, 2018–present

5.1.2 Past Graduate Students

1. Siddharth Ancha
MSc, 2016–2017
Co-supervised by Daniel Roy
Currently PhD student at Carnegie Mellon University

5.1.3 Current Undergraduates

1. Aidan Gomez, 2016–present
2. Tiantian Fang, 2017–present
3. Sicong Huang, 2017–present
4. Qiyang Li, 2017–present

5. Cem Anil, 2017–present
6. Xuechen Li, 2018–present
7. Xuchan Bao, 2018–present
8. SiQi Hao, 2018–present

5.1.4 Past Undergraduates

1. Behzad Abghari, 2016–2017
2. Jingxing Qian, 2017
3. Weijie Xu, 2017
4. Yeming Wen, 2017
5. Matthew MacKay, 2017
6. Jing Yao Li, 2016–2017

5.2 Reviewing and Editorial Activity

5.2.1 Senior Program Committee (Area Chair)

1. Neural Information Processing Systems (NIPS), 2018.
2. Conference on Uncertainty in Artificial Intelligence (UAI), 2018.
3. Association for the Advancement of Artificial Intelligence (AAAI) Conference on Artificial Intelligence, 2018.
4. International Conference on Machine Learning (ICML), 2017.
5. Conference on Artificial Intelligence and Statistics (AISTATS), 2017.
6. International Conference on Learning Representations (ICLR), 2016, 2018.

5.2.2 Conference and Journal Reviewing

- Program Committee:
 - International Conference on Learning Representations (ICLR), 2015, 2017.
 - Neural Information Processing Systems (NIPS), 2013, 2014, 2015, 2016, 2017.
 - International Conference on Machine Learning (ICML), 2014, 2015, 2016, 2018.
 - Artificial Intelligence and Statistics (AISTATS), 2012, 2013, 2014, 2016.
 - Conference on Uncertainty in Artificial Intelligence (UAI), 2015, 2017.
 - European Conference on Computer Vision (ECCV), 2014.
 - Computer Vision and Pattern Recognition (CVPR), 2012, 2013, 2014.
 - International Conference on Computer Vision (ICCV), 2013.

- Reviewed journal submissions:
 - Journal of Machine Learning Research (JMLR), 2014, 2015, 2016, 2017, 2018.
 - Neural Computation, 2014, 2015, 2016.
 - Transactions on Pattern Analysis and Machine Intelligence (PAMI), 2013, 2014, 2015, 2017.
 - SIAM Journal on Applied Algebra & Geometry (SIAGA), 2016.

5.3 Leadership

- Co-organizer of the 2018 Deep Learning Summer School, a selective event attended by hundreds of researchers.
- Co-organizer of the 2016 Neural Information Processing Systems (NIPS) Deep Learning Symposium, an event attended by over 1000 researchers which highlighted the year's most exciting developments in deep learning.

5.4 Committee Memberships

- Undergraduate Education Committee, 2017–present
- Working Group for the Engineering Science Machine Intelligence Major, 2017–present
- Working Group for the Undergraduate Program in Data Science, 2016–present

6 Contributions

6.1 Presentations

- Data, Learning, and Inference (DALI), Lanzarote, Canary Islands, Spain, April 3, 2018.
- Borealis AI Research, Toronto, ON, November 9, 2017.
- Vector Institute Endless Summer School, Toronto, ON, November 1, 2017.
- Google Study Group on Domain Specific Architectures. Mountain View, CA (delivered remotely), August 24, 2017.
- Smith School of Business. Kingston, ON, July 19, 2017.
- OpenAI. San Francisco, CA, July 11, 2017.
- Center for Human-Compatible Artificial Intelligence. Berkeley, CA, July 7, 2017.
- Google. Mountain View, CA, Nov. 17, 2016.
- Berkeley Computer Vision Seminar. Berkeley, CA, Sept. 20, 2016.
- World Congress on Probability and Statistics. Toronto, ON, July 15, 2016.

- Invited seminar talks at 10 universities: University of Cambridge, University of Toronto, Harvard University, Massachusetts Institute of Technology, University of Washington, University of California Los Angeles, New York University, Princeton University, Cornell University, Stanford University, Feb.–Apr. 2016.
- Google DeepMind. London, UK, March 30, 2016.
- Future of Humanity Institute. Oxford, UK, March 28, 2016.
- Google Brain. Mountain View, CA, March 18, 2016.
- OpenAI. San Francisco, CA, March 11, 2016.
- Closing Conference on Statistical and Computational Analytics for Big Data. Halifax, NS. June 12, 2015.
- Harvard Intelligent Probabilistic Systems Seminar. Cambridge, MA. Apr. 30, 2015.
- Next.ML Deep Learning Workshop. Cambridge, MA. Apr. 27, 2015.
- Fields Institute Thematic Program on Statistical Inference, Learning, and Models for Big Data. Toronto, ON. Jan. 26, 2015.
- Neural Information Processing Systems (NIPS) Workshop on Software Engineering for Machine Learning. Montreal, QC. Dec. 13, 2014.
- Harvard Intelligent Probabilistic Systems Seminar. Cambridge, MA. Oct. 2, 2014.
- Cognitive Computing Forum. San Jose, CA. Aug. 21, 2014.
- CIFAR NCAP Summer School. Toronto, ON. Aug. 15, 2014.
- Brown Pattern Theory Seminar. Providence, RI, Oct. 9, 2013.
- Google Machine Learning Seminar. Mountain View, CA, Aug. 8, 2013.
- University of Toronto, March 13, 2013.
- University of Edinburgh Machine Learning Seminar. Edinburgh, UK, Nov. 6, 2012.
- Gatsby Unit Seminar. London, UK, October 31, 2012.
- Representation Learning Workshop, International Conference on Machine Learning. Edinburgh, UK, July 1, 2012.
- University of Massachusetts Amherst Machine Learning and Friends Lunch. Amherst, MA, Apr. 26, 2012.
- SRI International Sarnoff. Princeton, NJ, Apr. 19, 2012.
- IARPA Automated Machine Learning RFI Workshop. Arlington, VA, Apr. 6, 2012.

6.2 Publications

* indicates a student of the applicant

6.2.1 Refereed Conference Publications

1. Yeming Wen*, Paul Vicol*, Jimmy Ba, Dustin Tran, and Roger Grosse (2018)
Flipout: efficient pseudo-independent weight perturbations on mini-batches.
International Conference on Learning Representations (ICLR).
2. Yuhuai Wu*, Mengye Ren, Renjie Liao, and Roger Grosse (2018)
Understanding short-horizon bias in stochastic meta-optimization.
International Conference on Learning Representations (ICLR).
3. Aidan Gomez*, Mengye Ren, Raquel Urtasun, and Roger Grosse (2017)
The Reversible Residual Network: Backpropagation Without Storing Activations.
Neural Information Processing Systems (NIPS).
4. Yuhuai Wu*, Elman Mansimov, Roger Grosse, Shun Liao, and Jimmy Ba (2017)
Scalable trust-region method for deep reinforcement learning using Kronecker-factored approximation.
Neural Information Processing Systems (NIPS).
5. Yuhuai Wu*, Yuri Burda, Ruslan Salakhutdinov, and Roger Grosse (2017)
On the quantitative analysis of decoder-based generative models.
International Conference on Learning Representations (ICLR).
6. Jimmy Ba, Roger Grosse, and James Martens (2017)
Distributed second-order optimization using Kronecker-factored approximations.
International Conference on Learning Representations (ICLR).
7. Jacob Gardner, Chuan Guo, Kilian Weinberger, Roman Garnett, and Roger Grosse (2017)
Discovering and exploiting additive structure for Bayesian optimization.
Artificial Intelligence and Statistics (AISTATS).
8. Roger Grosse, Siddharth Ancha*, and Daniel Roy (2016)
Measuring the reliability of MCMC inference with bidirectional Monte Carlo.
Neural Information Processing Systems (NIPS).
9. Roger Grosse and James Martens (2016)
A Kronecker-factored approximate Fisher matrix for convolution layers.
International Conference on Machine Learning (ICML).
10. Yuri Burda, Roger Grosse, and Ruslan Salakhutdinov (2016)
Importance Weighted Autoencoders.
International Conference on Learning Representations (ICLR).
11. Jimmy Ba, Roger Grosse, Ruslan Salakhutdinov, and Brendan Frey (2015)
Learning Wake-Sleep Recurrent Attention Models.
Neural Information Processing Systems (NIPS).
12. Roger Grosse and Ruslan Salakhutdinov (2015)
Scaling Up Natural Gradient by Sparsely Factorizing the Inverse Fisher Matrix.
International Conference on Machine Learning (ICML).

13. James Martens and Roger Grosse (2015)
Optimizing Neural Networks with Kronecker-factored Approximate Curvature.
International Conference on Machine Learning (ICML).
14. Yuri Burda†, Roger Grosse†, and Ruslan Salakhutdinov (2015)
Accurate and Conservative Estimates of MRF Log-likelihood Using Reverse Annealing.
Artificial Intelligence and Statistics (AISTATS).
† denotes equal contribution.
15. James Lloyd, David Duvenaud, Roger Grosse, Joshua Tenenbaum, and Zoubin Ghahramani (2014)
Automatic Construction and Natural-language Description of Nonparametric Regression Models.
AAAI Conference on Artificial Intelligence (AAAI).
16. Roger Grosse, Christopher Maddison, and Ruslan Salakhutdinov (2013)
Annealing Between Distributions by Averaging Moments.
Neural Information Processing Systems (NIPS).
(Oral presentation.)
17. David Duvenaud, James Lloyd, Roger Grosse, Joshua Tenenbaum, and Zoubin Ghahramani (2013)
Structure Discovery in Nonparametric Regression through Compositional Kernel Search.
International Conference on Machine Learning (ICML).
18. Roger Grosse, Ruslan Salakhutdinov, William Freeman, and Joshua Tenenbaum (2012)
Exploiting Compositionality to Explore a Large Space of Model Structures.
Conference on Uncertainty in AI (UAI).
(Best Student Paper.)
19. Roger Grosse, Micah Johnson, Edward Adelson, and William Freeman (2009)
A Ground-truth Dataset and Baseline Evaluations for Intrinsic Image Algorithms.
International Conference on Computer Vision (ICCV).
20. Honglak Lee, Roger Grosse, Rajesh Ranganath, and Andrew Ng (2009)
Convolutional Deep Belief Networks for Scalable Unsupervised Learning of Hierarchical Representations.
International Conference on Machine Learning (ICML).
(Best Application Paper.)
21. Roger Grosse, Rajat Raina, Helen Kwong, and Andrew Ng (2007)
Shift-invariant Sparse Coding for Audio Classification.
Conference on Uncertainty in AI (UAI).

6.2.2 Refereed Journal Publications

1. Honglak Lee, Roger Grosse, Rajesh Ranganath, and Andrew Ng (2011)
Unsupervised Learning of Hierarchical Representations with Convolutional Deep Belief Networks.
Communications of the ACM.

6.2.3 Papers Under Submission to Conferences

1. Tian Qi Chen, Xuechen Li, Roger Grosse, and David Duvenaud
Isolating the Sources of Disentanglement in the ELBO
Under submission to the International Conference on Machine Learning (ICML), 2018.
2. Guodong Zhang*, Shengyang Sun*, David Duvenaud, and Roger Grosse
Noisy Natural Gradient as Variational Inference
Under submission to the International Conference on Machine Learning (ICML), 2018.
3. Kuan-Chieh Wang, Paul Vicol*, James Lucas*, Li Gu, Roger Grosse, and Richard Zemel
Distilling the Posterior in Bayesian Neural Networks
Under submission to the International Conference on Machine Learning (ICML), 2018.
4. Shengyang Sun*, Guodong Zhang*, Chaoqi Wang, Wenyuan Zeng, Jiaman Li, and Roger Grosse
Differentiable Compositional Kernel Learning for Gaussian Processes
Under submission to the International Conference on Machine Learning (ICML), 2018.
5. James Lucas*, Richard Zemel, and Roger Grosse
Aggregated Momentum: Stability Through Passive Damping
Under submission to the International Conference on Machine Learning (ICML), 2018.

6.2.4 Workshop Publications and Unrefereed Technical Reports

1. Roger Grosse, Zoubin Ghahramani, and Ryan Adams (2015)
Sandwiching the Marginal Likelihood using Bidirectional Monte Carlo.
arXiv:1511.02543
2. Roger Grosse and David Duvenaud (2014)
Testing MCMC Code.
Neural Information Processing Systems (NIPS) Workshop on Software Engineering for Machine Learning.
3. Roger Grosse (2014)
Model Selection in Compositional Spaces.
Ph.D. thesis, Massachusetts Institute of Technology.