

Radford M. Neal

Curriculum Vitae, 25 April 2021

Biographical Information

Personal Data

Date of Birth: 12 Sept 1956	Citizenship: Canadian
Home address: 21 Shaftesbury Avenue, #106 Toronto, Ontario M4T 3B4	University address: Department of Statistical Sciences University of Toronto M5S 3G3
U of T web page:	http://www.cs.utoronto.ca/~radford/
Blog (not affiliated with U of T):	https://radfordneal.wordpress.com
E-mail addresses:	radford@stat.utoronto.ca , radfordneal@gmail.com

Degrees Awarded

BSc (Honours) in Computer Science, University of Calgary, 1977

MSc in Computer Science, University of Calgary, 1980

Thesis title: *An Editor for Trees*

Supervisor: David Hill

Ph.D in Computer Science, University of Toronto, 1995

Thesis title: *Bayesian Learning for Neural Networks*

Supervisor: Geoffrey Hinton

Employment and Other Relevant Experience

Present position:

Professor Emeritus, Department of Statistical Sciences and Department of Computer Science, University of Toronto,

Appointed as Lecturer 1 July 1994.

Appointed as Associate Member (Limited Term) of the graduate faculty 13 January 1995.

Promoted to Assistant Professor 1 July 1995.

Promoted to Associate Professor and full member of the graduate faculty 1 July 1999.

Promoted to Full Professor 1 July 2001.

Awarded a Canada Research Chair in Statistics and Machine Learning 1 July 2003 (held to 31 December 2016).

Cross-appointed to the Dalla Lana School of Public Health 14 February 2006.

Retired as of 1 January 2017.

Earlier teaching experience:

Sessional instructor in computer science at the University of Calgary for the following courses:

Summers 1979-1981 Taught half of the core second-year course on machine architecture and programming.

Winter 1981 Taught a fourth-year course on computer graphics.

1986/87 & 1987/88 Taught a fourth-year project course on system software design and development.

Teaching assistant in computer science at the University of Calgary (1977-1979) and the University of Toronto (1989).

Statistical consulting experience:

I have provided statistical consulting for a Canadian company regarding the analysis of mass spectroscopy data in 1990–1991, for a Canadian financial services company in 1997, for a US biomedical research company in 2004–2006, and for the Canadian Department of National Defence in 2006–2007.

Software engineering experience:

I did contract work as a software engineer on a variety of industrial projects, totaling approximately three years work, at various times between 1980 and 1989. These projects involved audio signal processing, computer graphics, data acquisition, oil and gas production accounting, data base design, and local-area networks.

From 1981 to 1985 I was a software engineer and research associate at the University of Calgary, first on a project to implement a code generator for a Simula compiler, and later as the system architect and team leader for software development in support of a research project in distributed computer systems.

I worked part-time as a programmer supporting researchers at the University of Calgary in computer science, management, and biology at various times between 1975 and 1980.

Academic History

Research Endeavours

I have active research interests in the following areas:

Bayesian inference: I am interested in fundamental issues concerning the Bayesian approach to statistical inference, in methods for solving the computational problems it gives rise to, and in the development of Bayesian models for practical applications.

Markov chain Monte Carlo: The principal way of solving many computational problems in Bayesian inference and statistical physics is to use Monte Carlo methods based on the simulation of a Markov chain. I am interested in improving these methods, and in applying them to problems of Bayesian inference, particularly for complex models such as those based on neural networks.

Flexible models for regression and classification: I am interested in methods for modeling regression and classification functions that allow a wide range of functional forms, including those based on neural networks and on Gaussian process priors.

Latent variable models: Probabilistic models with unobservable latent variables are widely used in some areas of science, and represent one possible approach to “concept learning” in artificial intelligence. I am interested in the development of these models, such as mixture models and models based on Dirichlet diffusion trees, and in their computational implementation, using Markov chain Monte Carlo and other methods.

Reinforcement learning: Organisms and artificial agents often need to learn how to act based only on the reward obtained following their actions. I am interested in methods for learning in this context, particularly when the environment is only partially observed, and memory is limited.

The R language for statistical programming: I am working on ways of speeding up the implementation of the widely-used R language for statistical programming, and on language extensions to make it faster and more useful.

Computer arithmetic: I am interested in fast ways of performing arithmetic with exactly-correct rounding, and in issues regarding reproducible computations.

Application areas: I am interested in biological applications such as haplotype inference, analysis of gene expression data, and analysis of spectroscopic data, in physics applications such as inference for high-energy physics experiments and calibration of models for the history of glaciation, and in public health applications such as modeling disease spread and the health effects of air pollution.

I have done work in the past on error correcting codes, data compression, empirical evaluation of statistical methods, programming environments, operating systems, compilers, and computer graphics.

Research Prize

In May 2004, I received the Lindley Prize, awarded every two years by the International Society for Bayesian Analysis for “innovative research in Bayesian Statistics” that was presented at the preceding ISBA or Valencia meeting. This award was for my paper on “Density modeling and clustering using Dirichlet diffusion trees”.

Research Awards

- 2011 \$105,000 (\$21,000 per year) for the period 1 April 2011 to 31 March 2016, individual research grant from the Natural Sciences and Engineering Research Council (NSERC), for research on Bayesian statistical computation, methodology, and applications.
- 2011 Renewal of Canada Research Chair, which comes with a \$10,000 per year research grant.
- 2005 \$170,000 (\$34,000 per year) for the period 1 April 2005 to 31 March 2010, individual research grant from the Natural Sciences and Engineering Research Council (NSERC), for research on Bayesian analysis using flexible models.
- 2003 Awarded a Canada Research Chair, which comes with a \$10,000 per year research grant, commencing 1 July 2003.
- 2001 \$96,000 (\$24,000 per year) for the period 1 April 2001 to 31 March 2005, individual research grant from the Natural Sciences and Engineering Research Council (NSERC), for research on theory, computation, and applications of Bayesian inference.
- 1999 \$100,000 over the period from 1999 to 2003, Premier’s Research Excellence Award.

- 1999 \$30,000 (\$15,000 per year) for the period 1 April 1999 to 31 March 2001, individual research grant from the Natural Sciences and Engineering Research Council (NSERC), for research on theory and computation for realistic Bayesian models.
- 1998 \$58,914 (about \$15,000 per year) for the period 1 April 1998 to 31 March 2002, my portion of a grant from the Institute for Robotics and Intelligent Systems for “Basic research and continuing education in learning algorithms”, with Y. Bengio (lead PI), R. Tibshirani, and F. Major.
- 1995 \$48,000 (\$12,000 per year) for the period 1 April 1995 to 31 March 1999, individual research grant from the Natural Sciences and Engineering Research Council (NSERC), for research on Markov chain Monte Carlo methods.
- Became a Principal Investigator with G. Hinton (lead PI), R. Tibshirani, D. Lowe, and M. LeBlanc on a project to research applications of neural networks and other statistical methods for medical diagnosis and financial forecasting, financed by a grant of \$669,736 (over four years) from the Institute for Robotics and Intelligent Systems (IRIS).
- 1994 \$6,200 from the Connaught Fund, as a new faculty start-up grant.

Scholarly and Professional Work

Book

- Neal, R. M. (1996) *Bayesian Learning for Neural Networks*, Lecture Notes in Statistics No. 118, New York: Springer-Verlag, 195 pages.

Refereed Research Papers

- Shestopaloff, A. Y. and Neal, R. M. (2018) “Sampling latent states for high-dimensional non-linear state space models with the embedded HMM method”, *Bayesian Analysis*, 26 pages, advance publication at projecteuclid.org/euclid.ba/1508551720. [Joint Author]
- Shahbaba, B., Lan, S., Johnson, W. O., and Neal, R. M. (2014) “Split Hamiltonian Monte Carlo”, *Statistics and Computing*, vol. 24, pp. 339-349. [Joint Author]
- Tarasov, L., Dyke, A. S., Neal, R. M., and Peltier, W. R. (2012) “A data-calibrated distribution of deglacial chronologies for the North American ice complex from glaciological modeling”, *Earth and Planetary Science Letters*, vol. 315-316, pp. 30-40. [Secondary Author]
- Shahbaba, B. and Neal, R. M. (2009) “Nonlinear models using Dirichlet process mixtures”, *Journal of Machine Learning Research*, vol. 10, pp. 1829-1850. [Secondary Author]
- Li, L. and Neal, R. M. (2008) “Compressing parameters in Bayesian high-order models with application to logistic sequence models”, *Bayesian Analysis*, vol. 3, pp. 793-822. [Joint Author]
- Li, L., Zhang, J., and Neal, R. M. (2008) “A method for avoiding bias from feature selection with application to naive Bayes classification models”, *Bayesian Analysis*, vol. 3, pp. 171-196. [Joint Author]

- Sun, S., Greenwood, C. M. T., and Neal, R. M. (2007) “Haplotype inference using a Bayesian hidden Markov model”, *Genetic Epidemiology*, vol. 31, pp. 937-948. [Joint Author]
- Jain, S. and Neal, R. M. (2007) “Splitting and merging components of a nonconjugate Dirichlet process mixture model” (with discussion), *Bayesian Analysis*, vol. 2, pp. 445-472. [Secondary Author]
- Listgarten, J., Neal, R. M., Roweis, S. T., Wong, P., and Emili, A. (2007) “Difference detection in LC-MS data for protein biomarker discovery”, *Bioinformatics*, vol. 23, pp. e198-e204. [Joint Author]
- Shahbaba, B. and Neal, R. M. (2007) “Improving classification when a class hierarchy is available using a hierarchy-based prior”, *Bayesian Analysis*, vol. 2, pp. 221-238, posted online 2006-11-16. [Joint Author]
- Shahbaba, B. and Neal, R. M. (2006) “Gene function classification using Bayesian models with hierarchy-based priors”, *BMC Bioinformatics*, 7:448, 9 pages. [Joint Author]
- Jain, S. and Neal, R. M. (2004) “A split-merge Markov chain Monte Carlo procedure for the Dirichlet process mixture model”, *Journal of Computational and Graphical Statistics*, vol. 13, pp. 158-182. [Joint Author]
- Neal, R. M. (2003) “Slice sampling” (with discussion), *Annals of Statistics*, vol. 31, pp. 705-767.
- Neal, R. M. (2001) “Annealed importance sampling”, *Statistics and Computing*, vol. 11, pp. 125-139.
- Diaconis, P., Holmes, S., and Neal, R. M. (2000) “Analysis of a non-reversible Markov chain sampler”, *Annals of Applied Probability*, vol. 10, pp. 726-752. [Joint Author]
- Neal, R. M. (2000) “Markov chain sampling methods for Dirichlet process mixture models”, *Journal of Computational and Graphical Statistics*, vol. 9, pp. 249-265.
- Moffat, A., Neal, R. M., and Witten, I. H. (1998) “Arithmetic coding revisited”, *ACM Transactions on Information Systems*, vol. 16, pp. 256-294. [Joint author]
- Neal, R. M. and Dayan, P. (1997) “Factor analysis using delta-rule wake-sleep learning”, *Neural Computation*, vol. 9, pp. 1781-1803. [Primary author]
- MacKay, D. J. C. and Neal, R. M. (1996) “Near Shannon limit performance of low density parity check codes”, *Electronics Letters*, vol. 32, pp. 1645-1646. Reprinted with printing errors corrected in vol. 33, pp. 457-458. [Secondary author]
- Neal, R. M. (1996) “Sampling from multimodal distributions using tempered transitions”, *Statistics and Computing*, vol. 6, pp. 353-366.
- Dayan, P., Hinton, G. E., Neal, R. M., and Zemel, R. S. (1995) “The Helmholtz machine”, *Neural Computation*, vol. 7, pp. 1022-1037. [Secondary author]
- Hinton, G. E., Dayan, P., Frey, B. J., and Neal, R. M. (1995) “The wake-sleep algorithm for unsupervised neural networks”, *Science*, vol. 268, pp. 1158-1161. [Secondary author]

Neal, R. M. (1994) “An improved acceptance procedure for the hybrid Monte Carlo algorithm”, *Journal of Computational Physics*, vol. 111, pp. 194-203.

Neal, R. M. (1992) “Connectionist learning of belief networks”, *Artificial Intelligence*, vol. 56, pp. 71-113.

Witten, I. H., Birtwistle, G. M., Cleary, J. G., Hill, D. R., Levinson, D., Lomow, G. A., Neal, R., Peterson, M., Unger, B. W., and Wyvill, B. L. M. (1983) “JADE: A distributed software prototyping environment”, *ACM Operating Systems Review*, vol. 7, no. 3. [Secondary author]

Witten, I. H. and Neal, R. M. (1982) “Using Peano curves for bilevel display of continuous-tone images”, *IEEE Computer Graphics and Applications*, vol. 2, no. 3. [Joint author]

Refereed Research Notes

Neal, R. M. (2000) “On deducing conditional independence from d -separation in causal graphs with feedback:”, *Journal of Artificial Intelligence Research*, vol. 12, pp. 87-91.

Gilks, W. R., Neal, R. M., Best, N. G., and Tan, K. K. C. (1997) “Corrigendum: Adaptive rejection Metropolis sampling”, *Applied Statistics*, vol. 46, pp. 541-542. [Joint author, not an author of the original paper]

Neal, R. M. (1993) “Comments on ‘A theoretical analysis of Monte Carlo algorithms for the simulation of Gibbs random field images’”, *IEEE Transactions on Information Theory*, vol. 39, p. 310.

Neal, R. M. (1992) “Asymmetric parallel Boltzmann machines are belief networks”, *Neural Computation*, vol. 4, pp. 832-834.

Refereed Tutorial/Review Papers

Kass, R. E., Carlin, B. P., Gelman, A., and Neal, R. M. (1998) “Markov Chain Monte Carlo in Practice: A Roundtable Discussion”, *The American Statistician*, vol. 52, pp. 93-100. [Joint author]

Witten, I. H., Neal, R. M., and Cleary, J. G. (1987) “Arithmetic coding for data compression”, *Communications of the ACM*, vol. 30, pp. 520-540. [Joint author]

Refereed Conference Papers

Neal, R. M. (2008) “Computing likelihood functions for high-energy physics experiments when distributions are defined by simulators with nuisance parameters”, in the proceedings of the PHYSTAT-LHC Workshop on Statistical Issues for LHC Physics, June 2007, CERN 2008-001, pp. 119-126.

Listgarten, J., Neal, R. M., Roweis, S. T., Puckrin, R., and Cutler, S. (2007) “Bayesian detection of infrequent differences in sets of time series with shared structure”, in B. Schölkopf, *et al.* (editors), *Advances in Neural Information Processing Systems 19* (aka NIPS*2006), 8 pages. [Joint Author]

- Meeds, E., Ghahramani, Z., Neal, R., and Roweis, S. (2007) “Modeling dyadic data with binary latent features”, in B. Schölkopf, *et al.* (editors), *Advances in Neural Information Processing Systems 19*, (aka NIPS*2006), 8 pages. [Secondary Author]
- Listgarten, J., Neal, R. M., Roweis, S. T., and Emili, A. (2005) “Multiple alignment of continuous time series”, in L. K. Saul, *et al.* (editors), *Advances in Neural Information Processing Systems 17*, 8 pages. [Joint Author]
- Neal, R. M., Beal, M. J., and Roweis, S. T. (2004) “Inferring state sequences for non-linear systems with embedded hidden Markov models”, in S. Thrun, *et al.* (editors) *Advances in Neural Information Processing Systems 16*, 8 pages. [Joint author]
- Neal, R. M. (2003) “Density modeling and clustering using Dirichlet diffusion trees”, in J. M. Bernardo, *et al.* (editors) *Bayesian Statistics 7*, Oxford University Press, pp. 619-629.
- Harvey, M. and Neal, R. M. (2000) “Inference for belief networks using coupling from the past”, in C. Boutilier and M. Goldszmidt (editors), *Uncertainty in Artificial Intelligence: Proceedings of the Sixteenth Conference (2000)*, pp. 256-263. [Joint Author]
- Neal, R. M. (1999) “Regression and classification using Gaussian process priors” (with discussion), in J. M. Bernardo, *et al.* (editors) *Bayesian Statistics 6*, Oxford University Press, pp. 475-501.
- MacKay, D. J. C. and Neal, R. M. (1995) “Good codes based on very sparse matrices”, in C. Boyd (editor) *Cryptography and Coding: 5th IAM Conference*, Lecture Notes in Computer Science No. 1025, pp. 100-111. Springer-Verlag. [Secondary author]
- Moffat, A., Neal, R., and Witten, I. H. (1995) “Arithmetic coding revisited”, in J. A. Storer and M. Cohn (editors) *Proceedings of the Fifth IEEE Data Compression Conference (Snowbird, Utah, 1995)*, pp. 202-211, Los Alamitos, California: IEEE Computer Society Press. [Joint author]
- Neal, R. M. (1993) “Bayesian learning via stochastic dynamics”, in C. L. Giles, S. J. Hanson, and J. D. Cowan (editors) *Advances in Neural Information Processing Systems 5*, pp. 475-482, San Mateo, California: Morgan Kaufmann.

Unrefereed Conference Papers

- Thompson, M. B. and Neal, R. M. (2010) “Slice sampling with adaptive multivariate steps: The shrinking-rank method”, *JSM 2010, Section on Statistical Computing*, pp. 3890-3896. [Secondary author]
- Hinton, G. E., Dayan, P., To, A., and Neal, R. (1995) “The Helmholtz machine through time”, in F. Fogelman-Soulié and R. Gallinari (editors) *ICANN-95*, pp. 483-490. [Secondary author]
- Neal, R. M. (1992) “Bayesian mixture modeling”, in C. R. Smith, G. J. Erickson, and P. O. Neudorfer (editors) *Maximum Entropy and Bayesian Methods: Proceedings of the 11th International Workshop on Maximum Entropy and Bayesian Methods of Statistical Analysis* (Seattle, 1991), pp. 197-211, Dordrecht: Kluwer Academic Publishers.

- Neal, R. M., Lomow, G. A., Peterson, M., Unger, B. W., and Witten, I. H. (1984) “Experience with an inter-process communication protocol in a distributed programming environment”, *CIPS Session 84 Conference* (Calgary). [Primary author]
- Wyvill, B. L. M., Neal, R., Levinson, D., and Bramwell, B. (1984) “JAGGIES: A distributed hierarchical graphics system”, *CIPS Session 84 Conference* (Calgary). [Secondary author]

Book Chapters

- Neal, R. M. (2011) “MCMC using Hamiltonian dynamics”, in the *Handbook of Markov Chain Monte Carlo*, S. Brooks, A. Gelman, G. L. Jones, and X.-L. Meng (editors), Chapman & Hall / CRC Press, pp. 113-162.
- Neal, R. M. and Zhang, J. (2006) “High dimensional classification with Bayesian neural networks and Dirichlet diffusion trees”, in I. Guyon, S. Gunn, M. Nikravesh, and L. A. Zadeh (editors) *Feature Extraction: Foundations and Applications*, Studies in Fuzziness and Soft Computing, Volume 207, Springer, pp. 265-2956. [Primary Author]
- Neal, R. M. (2006) “Classification with Bayesian neural networks”, in J. Quiñero-Candela, B. Magnini, I. Dagan, and F. D’Alche-Buc (editors) *Machine Learning Challenges. Evaluating Predictive Uncertainty, Visual Object Classification, and Recognising Textual Entailment*, Lecture Notes in Computer Science no. 3944, Springer-Verlag, pp.28-32.
- Neal, R. M. (1998) “Assessing relevance determination methods using DELVE”, in C. M. Bishop (editor), *Neural Networks and Machine Learning*, Berlin: Springer-Verlag, pp. 97-129.
- Neal, R. M. (1998) “Suppressing random walks in Markov chain Monte Carlo using ordered overrelaxation”, in M. I. Jordan (editor), *Learning in Graphical Models*, pp. 205-228, Dordrecht: Kluwer Academic Publishers.
- Neal, R. M. and Hinton, G. E. (1998) “A view of the EM algorithm that justifies incremental, sparse, and other variants”, in M. I. Jordan (editor), *Learning in Graphical Models*, Kluwer Academic Press, pp. 355-368, Dordrecht: Kluwer Academic Publishers. [Primary author]

Other Publications

- Neal, R. M. (2007) Book review: Pattern Recognition and Machine Learning, by Christopher M. Bishop, *Technometrics*, vol. 49, p. 366.

Technical Reports

Technical Reports not duplicating other papers listed:

- Neal, R. M. (2020) “Non-reversibly updating a uniform $[0, 1]$ value for Metropolis accept/reject decisions”, Technical Report, 14 pages. <http://arxiv.org/abs/2001.11950>
- Neal, R. M. (2015) “Fast exact summation using small and large superaccumulators”, Technical Report, 22 pages. <http://arxiv.org/abs/1505.05571>

- Neal, R. M. (2015) “Representing numeric data in 32 bits while preserving 64-bit precision”, Technical Report, 16 pages. <http://arxiv.org/abs/1504.02914>
- Shestopaloff, A. Y. and Neal, R. M. (2014) “Efficient Bayesian inference for stochastic volatility models with ensemble MCMC methods”, Technical Report, 17 pages. <http://arxiv.org/abs/1412.3013> [Joint Author]
- Shestopaloff, A. Y. and Neal, R. M. (2013) “On Bayesian inference for the M/G/1 queue with efficient MCMC sampling”, Technical Report, 18 pages. <http://arxiv.org/abs/1401.5548> [Joint Author]
- Wang, C. and Neal, R. M. (2013) “MCMC methods for Gaussian process models using fast approximations for the likelihood”, Technical Report, 21 pages. <http://arxiv.org/abs/1305.2235> [Joint Author]
- Shestopaloff, A. Y. and Neal, R. M. (2013) “MCMC for non-linear state space models using ensembles of latent sequences”, Technical Report, 18 pages. <http://arxiv.org/abs/1305.0320> [Joint Author]
- Wang, C. and Neal, R. M. (2012) “Gaussian process regression with heteroscedastic or non-Gaussian residuals”, Technical Report, 19 pages. <http://arxiv.org/abs/1212.6246> [Joint Author]
- Neal, R. M. (2012) “How to view an MCMC simulation as a permutation, with applications to parallel simulation and improved importance sampling”, Technical Report No. 1201, Dept. of Statistics, University of Toronto, 42 pages.
- Neal, R. M. (2010) “MCMC using ensembles of states for problems with fast and slow variables such as Gaussian process regression”, Technical Report No. 1011, Dept. of Statistics, University of Toronto, 24 pages.
- Thompson, M. and Neal, R. M. (2010) “Covariance-adaptive slice sampling”, Technical Report No. 1002, Dept. of Statistics, University of Toronto, 17 pages. [Secondary Author]
- Neal, R. M. (2006) “Puzzles of anthropic reasoning resolved using full non-indexical conditioning”, Technical Report No. 0607, Dept. of Statistics, University of Toronto, 53 pages.
- Neal, R. M. (2005) “Estimating ratios of normalizing constants using Linked Importance Sampling”, Technical Report No. 0511, Dept. of Statistics, University of Toronto, 37 pages.
- Neal, R. M. (2005) “The short-cut Metropolis method”, Technical Report No. 0506, Dept. of Statistics, University of Toronto, 28 pages.
- Neal, R. M. (2004) “Taking bigger Metropolis steps by dragging fast variables”, Technical Report No. 0411, Dept. of Statistics, University of Toronto, 9 pages.
- Neal, R. M. (2004) “Improving asymptotic variance of MCMC estimators: Non-reversible chains are better”, Technical Report No. 0406, Dept. of Statistics, University of Toronto, 25 pages.

- Neal, R. M. (2003) “Markov chain sampling for non-linear state space models using embedded hidden Markov models”, Technical Report No. 0304, Dept. of Statistics, University of Toronto, 9 pages.
- Neal, R. M. (2002) “Circularly-coupled Markov chain sampling”, Technical Report No. 9910 (original version November 1999, revised version February 2002), Dept. of Statistics, University of Toronto, 49 pages.
- Neal, R. M. (2001) “Transferring prior information between models using imaginary data”, Technical Report No. 0108, Dept. of Statistics, University of Toronto, 29 pages.
- Neal, R. M. (2001) “Defining priors for distributions using Dirichlet diffusion trees”, Technical Report No. 0104, Dept. of Statistics, University of Toronto, 25 pages.
- Pinto, R. L. and Neal, R. M. (2001) “Improving Markov Chain Monte Carlo Estimators by Coupling to an Approximating Chain”, Technical Report No. 0101, Dept. of Statistics, University of Toronto, 13 pages. [Joint Author]
- Neal, R. M. (1997) “Monte Carlo implementation of Gaussian process models for Bayesian regression and classification”, Technical Report No. 9702, Dept. of Statistics, University of Toronto, 24 pages.
- Neal, R. M. (1993) *Probabilistic Inference Using Markov Chain Monte Carlo Methods*, Technical Report CRG-TR-93-1, Dept. of Computer Science, University of Toronto, 144 pages.
- Neal, R. M. (1992) “Bayesian training of backpropagation networks by the hybrid Monte Carlo method”, Technical Report CRG-TR-92-1, Dept. of Computer Science, University of Toronto, 21 pages.

Technical Reports largely similar to other papers listed:

- Shestopaloff, A. Y. and Neal, R. M. (2016) “Sampling latent states for high-dimensional non-linear state space models with the embedded HMM method”, Technical Report, 23 pages (revised version). <http://arxiv.org/abs/1602.06030> [Joint Author] [A revised version was published in *Bayesian Analysis*]
- Shahbaba, B., Lan, S., Johnson, W. O., and Neal, R. M. (2011) “Split Hamiltonian Monte Carlo”, <http://arxiv.org/abs/1106.5941>. [Joint Author] [A revised version was published in *Statistics and Computing*]
- Shahbaba, B. and Neal, R. M. (2007) “Nonlinear models using Dirichlet process mixtures”, Technical Report No. 0707, Dept. of Statistics, 16 pages. [Secondary Author] [A revised version was published in the *Journal of Machine Learning Research*]
- Li, L., Zhang, J., and Neal, R. M. (2007) “A method for avoiding bias from feature selection with application to naive Bayes classification models”, Technical Report No. 0705, Dept. of Statistics, University of Toronto, 21 pages. [Joint Author] [A revised version was published in *Bayesian Analysis*]
- Shahbaba, B. and Neal, R. M. (2006) “Gene function classification using Bayesian models with hierarchy-based priors”, Technical Report No. 0606, Dept. of Statistics, University of Toronto, 14 pages. [Joint Author] [A revised version was published in *BMC Bioinformatics*]

- Shahbaba, B. and Neal, R. M. (2005) “Improving classification when a class hierarchy is available using a hierarchy-based prior”, Technical Report No. 0510, Dept. of Statistics, University of Toronto, 11 pages. [Joint Author] [A revised version was published in *Bayesian Analysis*]
- Jain, S. and Neal, R. M. (2005) “Splitting and merging components of a nonconjugate Dirichlet process mixture model”, Technical Report No. 0507, Dept. of Statistics, University of Toronto, 37 pages. [Secondary Author] [A revised version was published in *Bayesian Analysis*]
- Neal, R. M. (2000) “Slice sampling”, Technical Report No. 2005, Dept. of Statistics, University of Toronto, 40 pages. [A revised version was published in *Annals of Statistics*]
- Jain, S. and Neal, R. M. (2000) “A split-merge Markov chain Monte Carlo procedure for the Dirichlet process mixture model”, Technical Report No. 2003, Dept. of Statistics, University of Toronto, 32 pages. [Joint Author] [A revised version was published in the *Journal of Computational and Graphical Statistics*]
- Neal, R. M. (1998) “Annealed importance sampling”, Technical Report No. 9805, Dept. of Statistics, University of Toronto, 25 pages (first version 18 February 1998, revised 1 September 1998). [A revised version was published in *Statistics and Computing*]
- Neal, R. M. (1998) “Markov chain sampling methods for Dirichlet process mixture models”, Technical Report No. 9815, Dept. of Statistics, University of Toronto, 17 pages. [A revised version was published in the *Journal of Computational and Graphical Statistics*]
- Neal, R. M. (1997) “Markov chain Monte Carlo methods based on ‘slicing’ the density function”, Technical Report No. 9722, Dept. of Statistics, University of Toronto, 27 pages. [Was revised and extended, and re-titled as “Slice sampling”]
- Diaconis, P., Holmes, S., and Neal, R. M. (1997) “Analysis of a non-reversible Markov chain sampler”, Technical Report BU-1385-M, Biometrics Unit, Cornell University, 26 pages. [Joint Author] [A revised version was published in *Annals of Applied Probability*]
- Neal, R. M. and Dayan, P. (1996) “Factor analysis using delta-rule wake-sleep learning”, Technical Report No. 9607, Dept. of Statistics, University of Toronto, 23 pages. [Primary Author] [A revised version was published in *Neural Computation*]
- Neal, R. M. (1995) “Suppressing random walks in Markov chain Monte Carlo using ordered overrelaxation”, Technical Report No. 9508, Dept. of Statistics, University of Toronto, 22 pages. [A revised version became a chapter in *Learning in Graphical Models*]
- Neal, R. M. (1994) “Sampling from multimodal distributions using tempered transitions”, Technical Report No. 9421, Dept. of Statistics, University of Toronto, 22 pages. [A revised version was published in *Statistics and Computing*]
- Neal, R. M. (1994) “Priors for infinite networks” Technical Report CRG-TR-94-1, Dept. of Computer Science, University of Toronto, 22 pages. [A revised version is part of *Bayesian Learning for Neural Networks*]

Neal, R. M. (1991) “Bayesian mixture modeling by Monte Carlo simulation”, Technical Report CRG-TR-91-2, Dept. of Computer Science, University of Toronto, 23 pages. [A condensed version became “Bayesian mixture modeling”, in *Maximum Entropy and Bayesian Methods: Proceedings of the 11th International Workshop*.]

Neal, R. M. (1990) “Learning stochastic feedforward networks”, Technical Report CRG-TR-90-7, Dept. of Computer Science, University of Toronto, 34 pages. [A revised version became “Connectionist learning of belief networks”, published in *Artificial Intelligence*.]

Publicly-Released Software

pqR – A faster and extended version of the open-source R interpreter: I first released this improved version of the R interpreter on 2013-06-20. The most recent version was released on 2020-07-23; it includes a preliminary implementation of automatic differentiation. See pqR-project.org for up-to-date news of the project.

Software for Flexible Bayesian Modeling: Implements the models and algorithms described in *Bayesian Learning for Neural Networks*, in “Monte Carlo implementation of Gaussian process models for Bayesian regression and classification”, in “Markov chain sampling for Dirichlet process mixture models”, and in “Density modeling and clustering using Dirichlet diffusion trees”. Also allows Markov chain Monte Carlo methods to be applied to distributions specified by simple formulas. Approximately 250 pages of documentation and 53,000 lines of C program code. First public release August 1995 (for neural network models only), most recent release January 2020.

Software for Low Density Parity Check codes: Supports research into Low Density Parity Check (LDPC) codes, algorithms for encoding them, and algorithms for decoding them using probability propagation. Also includes modules implementing operations on dense and sparse modulo-2 matrices. Approximately 50 pages of documentation and 7000 lines of C program code. First public release March 2000, most recent release April 2012.

Plotutils+: I have modified and extended the GNU plotutils package, aiming especially to make it work better in conjunction with a command-line interface, as for my software for Flexible Bayesian Modeling. First release was 2020-12-21.

Xsum: Routines for performing float-point summation with exactly-rounded results. First release 2015-05-20, most recent update 2018-12-02.

SGGC - A Segmented Generational Garbage Collector: I have used this garbage collection module in releases of pqR since 2017-06-09, but it is designed for use in other projects as well. It is available at gitlab.com/radfordneal/sggc. Approximately 25 pages of API and implementation documentation and 4000 lines of C code, plus test programs.

The ‘matprod’ library for fast matrix multiplication that preserves naive rounding: Used in pqR since 2013, but potentially useful for other purposes. The most recent release has major performance improvements that exploit SIMD instructions and multiple cores. Available at github.com/radfordneal/matprod. Approximately 15 pages of API and implementation documentation and 7500 lines of C code, plus test programs.

The ‘helpers’ library for parallel computations: Used in pqR since 2013, but potentially useful for other purposes. Available at github.com/radfordneal/helpers. Approximately 25 pages of API and implementation documentation and 3500 lines of C code, plus test programs.

The ‘quotedargs’ package: An R package on CRAN for writing functions that quote their arguments. Released in 2017. Approximately 4 pages of documentation and 500 lines of C and R code.

GRIMS – General R Interface for Markov Sampling: I released a very preliminary version of this software for MCMC simulation in R in June 2011, and have released some updated (but still preliminary) versions since then. Approximately 15 pages of documentation and 2500 lines of R code.

R function for performing univariate slice sampling: A small R function released March 2008.

The DELVE software for assessing learning methods: I contributed substantially to the design of the DELVE software, with C. E. Rasmussen, G. E. Hinton, R. Tibshirani, and others, and wrote about half of the 102 page user’s manual. First public release December 1996.

Low-precision arithmetic coding software: This consists of about 3 pages of documentation and 1300 lines of code. First public release July 1991; last updated May 1993. I also contributed to the design of the arithmetic coding software maintained by Alistair Moffat, Dept. of Computer Science, University of Melbourne, which implements the algorithms in “Arithmetic coding revisited”.

Talks at Conferences Without Proceedings

- Dec 2019 Neal, R. M. “Non-reversibly updating a uniform $[0, 1]$ value for accept/reject decisions” (poster), 2nd Symposium on Advances in Approximate Bayesian Inference, Vancouver.
- Oct 2019 Neal, R. M. “Using deterministic maps when sampling from complex distributions”, Evolution of Deep Learning Symposium, Toronto.
- July 2019 Neal, R. M. “An automatic differentiation extension for R, and its implementation in pqR”, RIOT2019, Toulouse.
- July 2017 Neal, R. M. “Advances in Memory Management and Symbol Lookup in pqR”, 3rd Workshop on R Implementation, Optimization and Tooling (RIOT 2017), Brussels, Belgium.
- July 2017 Neal, R. M. “Recent and Planned Language Extensions in pqR”, 3rd Workshop on R Implementation, Optimization and Tooling (RIOT 2017), Brussels, Belgium.
- June 2015 Neal, R. M. “Can interpreting be as fast as byte compiling? + Other developments in pqR”, R Summit Conference, Copenhagen, Denmark.
- Jan 2015 Neal, R. M. “Learning to randomize and remember in partially-observed Environments”, Fields Institute Workshop on Big Data and Statistical Machine Learning, Toronto.

- June 2014 Neal, R. M. “Proposals for extending the R language”, Directions in Statistical Computing (DSC 2014), Brixen / Bressanone, Italy.
- June 2014 Neal, R. M. “Speed improvements in pqR: Current status and future plans”, Directions in Statistical Computing (DSC 2014), Brixen / Bressanone, Italy.
- Jul 2011 Neal, R. M. “Probability and anthropic reasoning in small, large, and infinite universes”, invited talk at the Perimeter Institute conference on Challenges for Early Universe Cosmology, Waterloo, Ontario.
- Jul 2011 Neal, R. M. “New Monte Carlo methods based on Hamiltonian dynamics”, invited talk at MaxEnt 2011, Waterloo, Ontario.
- Apr 2010 Neal, R. M. “Can playing billiards lead to over-relaxation?”, invited talk at 2010 Statistics Graduate Student Research Day, Fields Institute, Toronto.
- Oct 2004 Sun, S., Greenwood, C. T. M., and Neal, R. M. “Haplotype inference using a Hidden Markov Model with efficient Markov Chain sampling”, poster presentation, American Society of Human Genetics annual meeting, Toronto. [Joint Author]
- Aug 2004 Neal, R. M. “Discussion of a talk by M. Jordan on “Variational inference in exponential families: Alternatives to MCMC”, Joint Statistical Meetings, Toronto.
- May 2004 Neal, R. M. “A new proof of Peskun’s theorem regarding the asymptotic variance of MCMC estimators”, ISBA 2004, Viña del Mar, Chile.
- Apr 2004 Tarasov, L., Neal, R., and Peltier, W. R. “A Bayesian calibration methodology applied to ice-sheet modeling”, European Geosciences Union general assembly, Nice. [Joint author, presented by L. Tarasov]
- Aug 2001 Neal, R. M. “Survival Analysis Using a Bayesian Neural Network”, Joint Statistical Meetings, Atlanta.
- Aug 1999 Neal, R. M. “Markov Chain Sampling Using Hamiltonian Dynamics”, Joint Statistical Meetings, Baltimore.
- June 1996 Neal, R. M. “Faster Markov chain Sampling Using Overrelaxation Methods”, Statistical Society of Canada annual meeting, Waterloo.
- May 1994 Neal, R. M. “Hierarchical Bayesian inference for Hidden Markov Models”, Statistical Society of Canada annual meeting, Banff.

Invited Tutorial Presentation

- Dec 2004 Tutorial on “Bayesian Methods in Machine Learning”, NIPS 2004.

Invited Workshop Participation

- June 2007 PHYSTAT-LHC Workshop on Statistical Issues for LHC Physics, Organizers: Louis Lyons, Albert De Roeck, and others.

- May 2007 Third Workshop on Monte Carlo Methods, Harvard University. Organizer: Jun Liu.
- July 2006 Workshop on “Statistical Inference Problems in High Energy Physics and Astronomy”, Banff International Research Station. Organizers: James Linnemann, Louis Lyons, and Nancy Reid.
- April 2006 Workshop on “Nonlinearity and Randomness in Complex Systems”, Dept. of Mathematics, State University of New York at Buffalo. Organizers: Gino Biondini, Bruce Pitman, and William Kath.
- June 2005 Workshop on “Mathematical Issues in Molecular Dynamics”, Banff International Research Station. Organizers: Robert Skeel and Paul Tupper.
- Aug 2004 Second Workshop on Monte Carlo Methods, Harvard University. Organizer: Jun Liu.
- Dec 2003 NIPS 2003 workshop on feature extraction, Whistler. Organizer: Isabelle Guyon.
- Sep 2002 First Cape Cod Workshop on Monte Carlo Methods, Hyannis, Massachusetts. Sponsored by the NSF, IMS, and ISBA.
- Dec 2001 Workshop on “New Trends in Optimization and Computational Algorithms”, Kyoto, Japan. Sponsored by the The Institute of Statistical Mathematics.
- May 2001 Workshop on “Statistical Physics and Capacity-Approaching Codes”, Trieste, Italy. Sponsored by the Abdus Salam International Centre for Theoretical Physics.
- Mar 2001 Workshop on “Statistical Mixtures and Latent Structure Models”, ICMS, Edinburgh, UK. Organizers: C. P. Robert and D. M. Titterton.
- Oct 1999 Workshop on “Relating Causal Structure to Conditional Independence Structure”, The Fields Institute, Toronto. Organizers: Thomas Richardson and Peter Spirtes.
- Aug 1999 Workshop on “Codes, Systems and Graphical Models”, Institute for Mathematics and its Applications, Minneapolis. Organizers: G. David Forney, *et al.*
- Oct 1998 Workshop on Monte Carlo Methods, part of the programme on “Probability and Its Applications” at the Fields Institute, Toronto. I was a co-organizer (with N. Madras and J. Rosenthal) of this workshop.
- Aug 1997 Programme on “Neural Networks and Machine Learning” at the Newton Institute, University of Cambridge. Organizers: Chris Bishop, *et al.*
- June 1996 AMS-IMS-SIAM Joint Workshop on “Stochastic Inference, Monte Carlo, and Empirical Methods”, Mt. Holyoke, Massachusetts, Organizers: Ian McKeague and Priscilla Greenwood.

- Apr 1995 Workshop on “Statistics and Neural Networks”, International Centre for Mathematical Sciences, Edinburgh. Organizers: Jim Kay and Mike Titterton.
- June 1994 AMS-IMS-SIAM Joint Workshop on “Markov Chain Monte Carlo Methods”, Mt. Holyoke, Massachusetts. Organizers: A. Gelfand, C. Geyer, P. Green, and A. F. M. Smith.
- Nov 1993 The “Statistical basis of learning and classification” section of the Computer Vision programme at the Newton Institute, University of Cambridge. Organizer: Brian Ripley.
- Apr 1992 Workshop on “Neural networks and nonparametric regression”, Math Sciences Research Institute, Berkeley. Organizers: Trevor Hastie and Rob Tibshirani.
- Dec 1991 Workshop on Bayesian neural network learning following the 1991 conference on Neural Information Processing Systems. Organizers: David MacKay and Steve Nowlan.

Invited Lectures

- 14 May 2020 University of Oxford, Statistics, “MCMC for Hierarchical Bayesian Models Using Non-reversible Langevin Methods” (virtual seminar).
- 18 Mar 2020 School of Computational Science & Engineering, McMaster University, “Non-reversible Langevin Methods for Sampling Complex Distributions” (virtual seminar).
- 27 Apr 2016 Greater Toronto R User’s Group, “Performance improvements and future language extensions in the pqR implementation of R”.
- 22 Apr 2016 CRM – University of Ottawa Distinguished Lecture, “Reinforcement learning with randomization, memory, and prediction”.
- 29 Nov 2013 Dept. of Mathematics and Statistics, University of Guelf, “Speeding up R with Multithreading, Task Merging, and Other Techniques”.
- 15 Nov 2010 Dept. of Computer Science, University of Toronto, “MCMC Using Ensembles of States with Application to Gaussian Process Regression”.
- 22 Oct 2010 Dept. of Economics, University of Toronto, “MCMC Using Ensembles of States with Application to Gaussian Process Regression”.
- 17 Jan 2008 Dept. of Statistics, University of Toronto, “Constructing Efficient MCMC Methods Using Temporary Mapping and Caching”.
- 29 Nov 2007 Dept. of Statistics, University of Western Ontario, “Constructing Efficient MCMC Methods Using Temporary Mapping and Caching”.
- 27 Sep 2007 Canadian Institute for Theoretical Astrophysics, “A Critique of Anthropic Reasoning, with Applications to Cosmology and the Fermi Paradox”.

- 12 Feb 2007 Dept. of Computer Science, University of Toronto, “Constructing Efficient MCMC Methods Using Temporary Mapping and Caching”.
- 11 Dec 2006 Dept. of Statistics, Columbia University, “Constructing Efficient MCMC Methods Using Temporary Mapping and Caching”.
- 21 Feb 2003 Dept. of Mathematics and Statistics, York University, “Density Modeling and Clustering Using Dirichlet Diffusion Trees”.
- 12 Jun 2002 Cavendish Laboratory, Cambridge University, “Lossy Data Compression”.
- 10 Jun 2002 Dept. of Mathematics, University of Bristol, “Circularly-Coupled Markov Chain Sampling”.
- 22 May 2002 The Applied Biostatistics Association / Clinical Biostatistics Rounds, Toronto, “Survival analysis using a Bayesian neural network”.
- 12 Apr 2002 Dept. of Chemistry (Statistical Mechanics Group), University of Toronto: “Defining an Importance Sampling Distribution Using Hamiltonian Dynamics with Annealing”.
- 29 Nov 2001 Dept. of Statistics, University of Toronto: “Density Estimation and Clustering Using Dirichlet Diffusion Trees”.
- 1 Feb 2001 Dept. of Statistics, University of Toronto: “Circularly-Coupled Markov Chain Sampling”.
- 30 Jun 2000 Dept. of Statistics, Stanford University: “Circularly-Coupled Markov Chain Sampling”.
- 19 Apr 2000 Graduate School of Business, University of Chicago: “Circularly-Coupled Markov Chain Sampling”.
- 11 Mar 1999 Dept. of Physiology, University of Toronto: “Viewing Learning as Statistical Estimation for Models with Latent Variables”.
- 3 Dec 1998 Dept. of Statistics, University of Toronto: “Markov Chain Sampling Methods for Dirichlet Process Mixture Models”.
- 25 Nov 1998 Dept. of Mathematics and Statistics, York University: “Markov Chain Sampling Methods for Dirichlet Process Mixture Models”.
- 29 May 1998 Dept. of Engineering, University of Cambridge: “Annealed Importance Sampling”.
- 28 May 1998 Computer Science and Applied Mathematics, Aston University, Birmingham: “Annealed Importance Sampling”.
- 30 Nov 1997 Industrial Engineering and Management, Technion: “Bayesian Learning for Neural Networks”.
- 19 Nov 1997 Dept. of Statistics, Carnegie Mellon University: “Markov chain Monte Carlo methods based on ‘slicing’ the density function”.

- 20 Aug 1997 Dept. of Pure Mathematics and Mathematical Statistics, Cambridge University: “Markov chain Monte Carlo methods based on ‘slicing’ the density function”.
- 12 May 1997 Dept. of Statistics, University of Chicago: “Regression and classification using Gaussian process priors”.
- 4 April 1996 Dept. of Statistics, University of Western Ontario: “Markov chain Monte Carlo methods based on ‘slicing’ the density function”.
- 23 Feb 1995 Dept. of Statistics, University of Toronto: “Suppressing random walks in Markov chain Monte Carlo methods”.
- 14 Feb 1995 Dept. of Mathematics and Statistics, York University: “Sampling from multimodal distributions using tempered transitions”.
- 1 Dec 1994 Dept. of Statistics, University of Waterloo: “Sampling from multimodal distributions using tempered transitions”.
- 3 Feb 1994 Dept. of Statistics, University of Toronto: “Bayesian neural networks”.
- 11 Feb 1993 Dept. of Computer Science, University of Toronto: “Learning concepts with hierarchical latent variable models”.

Teaching Activities

Except as noted, I was solely responsible for teaching these courses.

Undergraduate Courses Taught Recently (Fall 2012 – Spring 2017)

- Spring 2017 CSC 121 — Computer Science for Statistics.
- Spring 2016 CSC 120 — Computer Science for the Sciences (R Section).
- Fall 2015 STA 410 — Computational Statistics.
- Spring 2015 CSC 120 — Computer Science for the Sciences (R Section).
- Fall 2014 STA 410 — Computational Statistics.
- Spring 2014 STA 414 — Statistical Methods for Machine Learning and Data Mining. (cross listing of STA 2104)
- Spring 2013 STA 414 — Statistical Methods for Machine Learning and Data Mining. (cross listing of STA 2104)

Graduate Courses Taught Recently (Fall 2012 – Spring 2017)

- Fall 2015 STA 2102 — Computational Techniques in Statistics. (cross listing of STA 410)
- Fall 2014 STA 2102 — Computational Techniques in Statistics. (cross listing of STA 410)

- Fall 2014 STA 3000 — Advanced Theory of Statistics.
I taught the Fall term of this course; Mike Evans taught the Spring term.
- Spring 2014 STA 3000 — Advanced Theory of Statistics.
I taught the Spring term of this course; Nancy Reid taught the Fall term.
- Spring 2014 STA 2104 — Statistical Methods for Machine Learning and Data Mining.
- Spring 2013 STA 4503 — Advanced Monte Carlo Methods and Applications.
(six week course)
- Spring 2013 STA 2104 — Statistical Methods for Machine Learning and Data Mining.

Graduate Students and Postdoctoral Fellows

I supervised the following graduate students, from the time listed to completion:

- Alexander Shestopaloff, PhD, Statistics, 2016 (from 2009)
Thesis title: *MCMC Methods for Non-linear State Space Models*
- Chunyi Wang, PhD, Statistics, 2014 (from 2009)
Thesis title: *Gaussian Process Regression with Heteroscedastic Residuals and Fast MCMC Methods*
- Meng Du, PhD, Statistics, 2013 (from 2008)
Thesis title: *Adjusting for Selection Bias Using Gaussian Process Models*
- Madeleine Thompson, PhD, Statistics, 2011 (from 2008)
Thesis title: *Slice Sampling with Multivariate Steps*
- Longhai Li, PhD, Statistics, 2007 (from 2003)
Thesis title: *Bayesian Classification and Regression with High-Dimensional Features*
- Peter Liu, MSc, Computer Science, 2007 (co-supervisor, from 2005)
Thesis title: *Using Gaussian Process Regression to Denoise Images and Remove Artefacts from Microarray Data*
- Babak Shahbaba, PhD, Public Health Sciences (Biostatistics), 2007 (from 2003)
Thesis title: *Improving Classification Models When a Class Hierarchy is Available*
- Jennifer Listgarten, PhD, Computer Science, 2006 (co-supervisor, from 2003)
Thesis title: *Analysis of Sibling Time Series Data: Alignment and Difference Detection*
- Shuying Sun, PhD, Statistics, 2006 (co-supervisor, from 2003)
Thesis title: *Haplotype Inference Using a Hidden Markov Model with Efficient Markov Chain Sampling*
- Ruxandra Pinto, PhD, Statistics, 2002 (from 1997)
Thesis title: *Improving Markov Chain Monte Carlo Estimators Using Overrelaxation and Coupling Techniques*

Sonia Jain, PhD, Statistics, 2002 (from 1998)

Thesis title: *Split-Merge Techniques for Bayesian Mixture Models*

Krunoslav Kovac, MSc, Computer Science, 2005 (from 2002)

Thesis title: *Multitask Learning for Bayesian Neural Networks*

Kiam Choo, MSc, Computer Science, 2000 (from 1997)

Thesis title: *Learning Hyperparameters for Neural Network Models Using Hamiltonian Dynamics*

Michael Harvey, MSc, Computer Science, 1999 (part-time, from 1996)

Thesis title: *Monte Carlo Inference for Belief Networks Using Coupling From the Past*

I supervised the following postdoctoral fellows:

Matthew Beal, Computer Science (July 2003 – July 2004)

Iain Murray, Computer Science (August 2007 – August 2009, co-supervisor)

I am currently on the advisory committees of the following PhD students:

Chris Cremer, Computer Science

Eric Corlett, Computer Science

Will Grathwohl, Computer Science

Mohammad Kaviul Khan, Biostatistics

Shengyang Sun, Computer Science

Yuhuai Wu, Computer Science

I played substantial roles on the advisory committees of the following PhD students (completion date given):

Victor Veitch, Statistics, 2017

Yanshuai Cao, Computer Science, 2017

Mohammad Norouzi, Computer Science, 2015

James Martens, Computer Science, 2015

Tijmen Tieleman, Computer Science, 2014

Navdeep Jaitly, Computer Science, 2014

Ilya Sutskever, Computer Science, 2012

Andriy Mnih, Computer Science, 2009

Ruslan Salakhutdinov, Computer Science, 2009

Ofer Shai, Electrical and Computer Engineering, 2009

Jim Huang, Electrical and Computer Engineering, 2009

Edward Meeds, Computer Science, 2008

Zi Jin, Statistics, 2009

Vladimir Jojic, Computer Science, 2007

Matthew Kowgier, Biostatistics, 2011

Marcus Lee, Management, 2003

Yee Whye Teh, Computer Science, 2003

Greg Kondrak, Computer Science, 2002

Andrew Brown, Computer Science, 2002

Brian Sallans, Computer Science, 2001

Sageev Oore, Computer Science, 2001
Alison Gibbs, Statistics, 2000
Xiang Sun, Statistics, 1999
Brendan Frey, Electrical and Computer Engineering, 1997
Carl Edward Rasmussen, Computer Science, 1996

I served on the SGS examining committees for the following PhD students (in addition to serving on the examining committees of my own students):

Mohammad Norouzi, Computer Science, 2015
James Martens, Computer Science, 2015
Tijmen Tieleman, Computer Science, 2014
Navdeep Jaitly, Computer Science, 2014
Billy Chang, Biostatistics, 2014
Chi Hay Tong, Aerospace Science and Engineering, 2013
Ilya Sutskever, Computer Science, 2012
Andriy Mnih, Computer Science, 2009
Ruslan Salakhutdinov, Computer Science, 2009
Ofer Shai, Electrical and Computer Engineering, 2009
Jim Huang, Electrical and Computer Engineering, 2009
Vladimir Jovic, Computer Science, 2007
Sigfrido Iglesias-Gonzalez, Statistics, 2007
Matthew Kowgier, Biostatistics, 2011
Marcus Lee, Management, 2003
Yee Whye Teh, Computer Science, 2003
Andrew Brown, Computer Science, 2002
Brian Sallans, Computer Science, 2001
Sageev Oore, Computer Science, 2001
Alison Gibbs, Statistics, 2000
Xiang Sun, Statistics, 1999
George Tomlinson, Biostatistics, 1998
Tomoyuki Yamakami, Computer Science, 1997
Brendan Frey, Electrical and Computer Engineering, 1997
Carl Edward Rasmussen, Computer Science, 1996
Margeurite Ennis, Statistics, 1996
Kin Chung Lo, Economics, 1995

I served as an external examiner for the following PhD students:

Xiaoqiang Li, Statistics, University of Western Ontario, 2002.
Paul Vanetti, Statistics, University of Oxford, 2020.

I served as a reader/examiner for the following MSc students:

James Martens, Computer Science, 2009
Ilya Sutskever, Computer Science, 2007
Brian Sallans, Computer Science, 1998
Rafal Kustra, Biostatistics, 1997

Administrative Positions

Major Administrative Positions

I was Associate Chair for Undergraduate Studies (Statistics) from July 2015 through June 2016.

I was Associate Chair for Graduate Studies in the Dept. of Statistics from January through August 2007.

Departmental Committee Work

Listed by academic year, commencing in July.

Computing Committee, Dept. of Statistics: Member in 1995, 1996; Chair in 1997, 1998, 1999, 2000 (1st half); Member in 2000 (2nd half), 2001; Chair in 2002, 2003, 2004; Member in 2005 (2nd half); Chair in 2006, 2008, 2009, 2010; Member in 2011, 2012 (2nd half), Chair in 2014.

Graduate Committee, Dept. of Statistics: Member in 1995, 1996, 1998 (2nd half), 2002, 2003, 2005 (2nd half); Chair in 2006; Member in 2009.

Hiring Committee, Dept. of Statistics: Member in 1998 (1st half), 1999, 2000 (1st half), 2001, 2002, 2003, 2010, 2013 (UTSC), 2014 (UTM), 2015 (two positions).

Curriculum Committee, Dept. of Statistics: Member in 2010, 2011, 2012.

Undergraduate Committee, Dept. of Statistics: Member in 2014, Chair 2015.

Seminar Coordinator, Dept. of Statistics: 1999.

PTR committee, Dept. of Statistics: Member in 2011, 2015.

Promotions Committee, Dept. of Statistics: Member in 2010, 2015.

Graduate admissions committee, Dept. of Computer Science, 2009, 2015.

Faculty Committee Work

Introductory Mathematics and Science Committee, Faculty of Arts and Science, 2012.

Other Professional Activity

Professional Committee Work

I was a member (in 2016) of a working group on standardizing and documenting the R API, sponsored by the R Consortium.

Organizational Activities

In 1998, I co-organized (with N. Madras and J. Rosenthal) a workshop on Monte Carlo methods held at the Fields Institute in October 1998, as part of the programme on Probability and Its Applications.

In 1995/1996, I organized a series of local Toronto meetings involving researchers on Markov chain Monte Carlo methods from Statistics, Computer Science, and Chemistry. In 1997/1998, I was co-organizer for a similar series of meetings (with J. Rosenthal and S. Whittington).

I organized and chaired an invited session on “Bayesian inference for high dimensional problems: Models and computation” for the Statistical Computing Section of the ASA, as part of the August 1996 Joint Statistical Meetings in Chicago.

Refereeing

I have refereed papers for the following journals:

- Annals of Applied Probability (1)
- Annals of Statistics (2)
- Annals of the Institute of Statistical Mathematics (1)
- Artificial Intelligence (2)
- Bayesian Analysis (5)
- Bernoulli (1)
- Biometrics (1)
- Biometrika (2)
- Canadian Journal of Forestry Research (1)
- Canadian Journal of Statistics (1)
- Communications in Statistics (1)
- Computational Statistics and Data Analysis (1)
- Electronic Journal of Probability (1)
- Erkenntnis (1)
- IEEE Communications Letters (1)
- IEEE Transactions on Computers (1)
- IEEE Transactions on Communications (1)
- IEEE Transactions on Information Theory (2)
- IEEE Transactions on Knowledge and Data Engineering (1)
- IEEE Transactions on Neural Networks (6)
- IEEE Transactions on Pattern Analysis and Machine Intelligence (3)
- IEEE Transactions on Signal Processing (1)
- Journal of Artificial Intelligence Research (1)
- Journal of Chemical Theory and Computation (1)
- Journal of Computational and Graphical Statistics (3)
- Journal of Lightwave Technology (1)
- Journal of Machine Learning Research (2)
- Journal of the American Statistical Association (7)
- Journal of the Royal Statistical Society B (5)
- Journal of Physics A (1)
- Machine Learning (2)
- Neural Computation (22)
- Neural Networks (2)
- Proceedings of the National Academy of Sciences (USA) (2)
- Random Structures and Algorithms (1)
- Statistical Science (1)
- Statistics and Computing (6)

Statistics and Probability Letters (2)
Stochastic Processes and their Applications (1)

I have refereed submissions to the following conferences:

Neural Information Processing Systems, 1993, 1997, 1999, 2003, 2004, 2005, 2018, 2019. I was also a referee and mentor for the NeurIPS 2019 “New in ML” workshop.
International Workshop on Artificial Intelligence and Statistics, 1997.
SIGGRAPH Asia, 2008.

I have also refereed 17 grant proposals for NSERC (Canada), 1 grant proposal for NERC (United Kingdom), and 1 grant proposal for FWF (Austria).

I also comment on research articles, books, and software on my blog, which is located at <http://radfordneal.wordpress.com/>.