Curriculum Vitæ Date of Revision: July, 2019

A. BIOGRAPHICAL INFORMATION

PERSONAL

RICHARD ZEMEL

Department of Computer Science
University of Toronto
6 King's College Road
Toronto, ON M5S 3H5

phone: (416) 978-7497
fax: (416) 978-1455
e-mail: zemel@cs.toronto.edu
url: www.cs.toronto.edu/~zemel

DEGREES

Ph.D. 1994 Department of Computer Science, University of Toronto
Dissertation: *A Minimum Description Length Framework for Unsupervised Learning.*

M.Sc. 1989 Department of Computer Science, University of Toronto.
Dissertation: *TRAFFIC: A Connectionist Model of Object Recognition.*

B.A. 1984 Department of History and Science, Harvard University.
Dissertation: *Old Ideas and New Viewpoints in Artificial Intelligence.*

EMPLOYMENT

Dept. of Computer Science, University of Toronto	Professor	2009-
Vector Institute for Artificial Intelligence	Founder, Research Director	2017-
Dept. of Computer Science, University of Toronto	Associate Professor	2001-2009
Tenure		July 1, 2001
School of Graduate Studies appointment		June 23, 2002
Dept. of Elec. Eng. & Comp. Sci., U. of Calif., Berkeley	Visiting Scholar	Jul-Jun 2008-09
Dept. of Computer Science, Stanford University	Visiting Professor	Jul-Jun 2008-09
Dept. of Biophysics, Princeton University	Visiting Faculty	Jan-Jun 2004
Dept. of Computer Science, University of Toronto	Assistant Professor	2000-01
Dept. of Psychology, University of Arizona Depts. of Computer Science and Cognitive Science	Assistant Professor	1996-2000
Dept. of Psychology, Carnegie Mellon University	Postdoctoral Fellow	1994-96
Computational Neurobiology Laboratory, Salk Institute	Postdoctoral Fellow	1993-94

HONORS

NSERC Industrial Research Chair in Machine Learning, 2018-present

Pioneers of AI, NVIDIA, 2017

Senior Fellow, Canadian Institute for Advanced Research, 2007-present

Executive Board, Neural Information Processing Society, 2012-present

Google Focused Research Award, 2012

Discovery Accelerator Awards, National Sciences and Engineering Research Council, 2009, 2014

Dean's Excellence Award, University of Toronto, 2005, 2006, 2007, 2008, 2011, 2013, 2014

Young Investigator Award, Office of Naval Research, 1998

National Sciences and Engineering Research Council Postgraduate Scholarship, 1989-1991

Ontario Graduate Scholarship, University of Toronto, 1989-1990

U.S. Alumni Fellowship, University of Toronto, 1987-1988

B.A. degree *magna cum laude*, Harvard University, 1984

John Harvard Scholarship, 1980-1984

Richard King Mellon National Merit Scholarship, 1980-1984

Presidential Scholar, 1980

B. ACADEMIC HISTORY

RESEARCH AWARDS

NSERC Industrial Rewsearch Chair in Machine Learning
"Making Deep Networks Expressive, Controllable and Fair"
2018-2023. PI, award: $300,000

DARPA Learning with Little Labels
"Few-Shot Learning via Embeddings"
2019-2023. Co-PI, total award: $1,650,000

IARPA Microns Program Grant
"Machine Intelligence from Cortical Networks"
2015-2020. Co-PI, total award: $1,250,000

DARPA Explainable AI Grant
"Deep Attention-Based Respresentations for Explanation"
2017-2021. Co-PI, award: $183,000

National Sciences and Engineering Research Council Discovery Grant
"Machine learning to understand images and text"
2014-2019. PI, total award: $310,000

National Sciences and Engineering Research Council Discovery Accelerator Supplement
2014-2017. PI, total award: $120,000

Deep learning for Samsung.
2014-2016. Co-PI, total award: $125,000

Mitacs Accelerate Graduate Research Internship
"Neural network model for predicting NBA shot outcomes"
2014-2016. PI, total award: $60,000

Amazon Web Services in Education
2014-2015. PI, total award: 10,000 credits

Mitacs Accelerate Graduate Research Internship
"Book recommendation: Improving collaborative filtering with content information"
2013. PI, total award: $30,000

Google Focused Research Award
2013-2015. Co-PI, total award: $1M

National Sciences and Engineering Research Council Research Tools and Instruments Grant
"Large scale machine learning using GPUs"
2011-2013. Co-PI, total award: $59,951

National Sciences and Engineering Research Council Discovery Grant
"Machine learning for image undertanding and user modelling"
2009-2014. PI, total award: $250,000

National Sciences and Engineering Research Council Discovery Accelerator Supplement
2009-2012. PI, total award: $120,000

RESEARCH AWARDS (cont.)

National Sciences and Engineering Research Council Strategic Partnerships Grant
"New machine learning methods for collaborative filtering, document retrieval and image retrieval"
2008-2010. Co-PI, total award: $194,000

Microsoft Research/Live Labs Research Grant,
"Digital photo organization"
2007-2009. PI, total award: $100,000

James S. McDonnell Foundation Collaborative Research Program,
"Understanding recovery from brain injury: Putting network models to work"
2005-2008. Co-PI, total award: $940,000

Communications and Information Technology Ontario (CITO) Partnerships Grant
"Semantic retrieval of image and video data"
2005-2007. Co-PI, total award: $250,000

National Sciences and Engineering Research Council Discovery Grant
"Learning intermediate representations"
2004-2009. $150,000

Canada Foundation for Innovation (CFI) New Opportunities Grant
"Machine learning and neural networks"
2002-2005. Co-PI, total award: $413,000

Institute for Robotics and Intelligent Systems (IRIS) Research Grant
"Learning algorithms"
2002-2005. Co-PI, total award: $495,000

Canadian Institutes for Health Research (CIHR) New Emerging Teams Grant
"Development of an integrative computational neuroscience program"
2002-2007. Co-PI, total award:$968,000

Office of Naval Research, Young Investigator Award
"Computational models of population codes"
2000-2002. $144,000

National Sciences and Engineering Research Council Research Grant
"Population code belief networks"
2000-2004. $100,000

Bell University Laboratories Research Grant
2000-2003. $60,000

National Institutes of Health Program Project Grant
1997-2002. Co-PI, total award:$1.1M

James McDonnell Foundation Research Grant
"Mechanisms of object attention"
1994-1998. $200,000

C. SCHOLARLY AND PROFESSIONAL WORK

Career Publication Count	
Scholarly Books (authored)	0
Scholarly Books (edited)	0
Chapters in Books	6
Papers in refereed journals	29
Papers in refereed conferences	101
Major invited contributions	5
Other conference abstracts/posters/contributions	12
Other publications	11

7. Refereed Publications

A. Articles in Refereed Journals

Kiros, R, Salakhtudinov, R, & Zemel, RS (2015). Unifying visual-semantic embeddings with multimodal neural language models. *Transactions of the Association for Computational Linguistics*, To appear.

Zheng, Y, Zhang, Y, Zemel, RS, & Larochelle, H. (2015). A neural autoregressive approach to attention-based recognition. *International Journal of Computer Vision*, To appear.

Volkovs, M, & Zemel, RS (2013). New learning methods for supervised and unsupervised preference aggregation. *Journal of Machine Learning Research*, 15(Mar):11351176.

Schmah, T, Zemel, RS, Hinton, GE, Small, S, & Strother, S. (2010). Comparing classification methods for longitudinal fMRI studies. *Neural Computation*. 22(1): 2729-2762.

Ross, D, Tarlow, D, & Zemel, RS. (2010). Learning articulated structure from motion. *International Journal of Computer Vision*. 88(2): 214-242.

Snoek, J, Hoey, J, Stewart, L, Zemel, RS, & Mihalidis, A. (2009). Automated detection of unusual events on stairs. *Image and Vision Computing Journal*. 27(1): 153-166.

Natarajan, R, Huys, Q, Dayan, P, & Zemel, RS (2008). Encoding and decoding dynamic population codes. *Neural Computation*. 20(9): 2325-2360.

Stewart, L, He, X, & Zemel, RS (2008). Learning flexible features for conditional random fields. *IEEE Transactions: Pattern Analysis and Machine Intelligence*. 30(8): 1415-1426.

Klam, F, Zemel, RS, & Pouget, A (2008). Population coding with motion energy filters: The impact of correlations. *Neural Computation*, 20(1): 146 - 175.

Huys, Q, Zemel, RS, Natarajan, R, & Dayan, P (2007). Fast population coding. *Neural Computation*, 19(2): 460-497.

He, X, Zemel, RS., & Mnih, V. (2006). Topological map learning from outdoor image sequences. *Journal of Field Robotics*, 23(11/12): 1091-1104.

Ross, D & Zemel, RS (2006). Learning parts-based representations of data. *Journal of Machine Learning Research*: 2369-2397.

Welling, M, Zemel, RS, & Hinton, GE (2004). Probabilistic sequential independent components analysis. *Transactions in Neural Networks*, Special Issue on Information Theory, 15(4): 838-849.

Pouget, A, Dayan, P, & Zemel, RS (2003). Inference and computation with population codes. *Annual Reviews of Neuroscience: 26*, pp. 381-410.

Zemel, RS & Mozer, MC (2002). Localist attractor networks. *Neural Computation*, 13:5, pp. 1045-1064.

Zemel, RS, Behrmann, M, Mozer, MC, & Bevalier, D (2002). Experience-dependent perceptual grouping and object-based attention. *Journal of Experimental Psychology: Human Perception and Performance*, 28:1, pp. 202-217.

Zemel, RS & Pillow, J (2000). Encoding multiple orientations in a recurrent network. *Neurocomputing*, 32-33, pp. 609-616. Preliminary version appeared in: *Proceedings of the 1999 Computational Neuroscience Meeting*.

Behrmann, M, Zemel, RS, & Mozer, MC (2000). Occlusion, symmetry, and object-based attention: Reply to Saiki (2000). *Journal of Experimental Psychology: Human Perception and Performance*, 26:4, pp. 1497-1505.

Pouget, A, Dayan, P, & Zemel, RS (2000). Information processing with population codes. *Nature Reviews Neuroscience*, 1, pp. 125-132.

Zemel, RS, Dayan, P, & Pouget, A (1998). Probabilistic interpretation of population codes. *Neural Computation*, 10:2, pp. 403-430.

Zemel, RS & Sejnowski, TJ (1998). A model for encoding multiple object motions and self-motion in area MST of primate visual cortex. *The Journal of Neuroscience*, 18:1, pp. 531-547.

Behrmann, M, Zemel, RS, & Mozer, MC (1998). Object-based attention and occlusion: Evidence from normal subjects and a computational model. *Journal of Experimental Psychology: Human Perception and Performance*, 24:4, pp. 1011-1036.

Gray, M, Pouget, A, Zemel, RS, Nowlan, SJ, & Sejnowski, TJ (1998). Reliable disparity estimation through selective integration. *Visual Neuroscience*, 15, pp. 511-528.

Zemel, RS & Hinton, GE (1995). Developing population codes by minimizing description length. *Neural Computation*, 7:3, pp. 549-564.

Zemel, RS, Williams, CKI, & Mozer, MC (1995). Lending direction to neural networks. *Neural Networks*, 8:4, pp. 503-512.

Dayan, P & Zemel, RS (1995). Competition and multiple cause models. *Neural Computation*, 7:3, pp. 565-579.

Dayan, P, Hinton, GE, Neal, R, & Zemel, RS (1995). The Helmholtz Machine. *Neural Computation*, 7:5, pp. 889-904.

Mozer, MC, Zemel, RS, Behrmann, M, & Williams, CKI (1992). Learning to segment images using dynamic feature binding. *Neural Computation*, 4:5, pp. 647-662.

B. Refereed Conference Papers

Liao, R, Zhao, Z, Urtasun, R & Zemel, RS (2019). LanzcosNet: Multiscale deep graph convolutional networks. In: *ICLR 2019: International Conference on Learning Representations.* [29% acceptance rate].

Law, M, Snell, J, Farahmand, A, Urtasun, R & Zemel, RS (2019). Dimensionality reduction for representing the knowledge of probabilistic models. In: *ICLR 2019: International Conference on Learning Representations.* [29% acceptance rate].

Jacobsen, J, Behrmann, J, Zemel, RS & Bethge, M (2019). Excessive invariance causes adversarial vulnerability. In: *ICLR 2019: International Conference on Learning Representations.* [29% acceptance rate].

Lucas, J, Sun, S, Zemel, RS & Grosse, R (2019). Aggregated momentum: Stability through passive damping. In: *ICLR 2019: International Conference on Learning Representations.* [29% acceptance rate].

Madras, D, Creager, E, Pitassi, T & Zemel, RS (2019). Fairness through causal awareness: Learning causal latent-variable models for biased data. In *ACM Conference on Fairness, Accountability and Transparencey (FAT*).* [19% acceptance rate].

Klys, J, Snell, J & Zemel, RS (2018). Learning latent subspaces in variational autoencoders. In: *NIPS 2018: Advances in Neural Information Processing Systems.*

Madras, D, Pitassi, T & Zemel, RS (2018). Predict responsibly: improving fairness and accuracy by learning to defer. In: *NIPS 2018: Advances in Neural Information Processing Systems.* [21% acceptance rate].

Zhang, L, Rosenblatt, G, Fetaya, E, Liao, R, Byrd, W, Might, M, Urtasun, R & Zemel, RS (2018). Neural guided constraint logic programming for program synthesis. In: *NIPS 2018: Advances in Neural Information Processing Systems.* [21% acceptance rate].

Liao, R, Xiong, Y, Fetaya, E, Zhang, L, Yoon, K, Pitkow, Z, Urtasun, R & Zemel, RS (2018). Reviving and improving recurrent back-propagation. In: *NIPS 2018: Advances in Neural Information Processing Systems.* [21% acceptance rate].

Madras, D, Creager, E, Pitassi, T & Zemel, RS (2018). Learning adversarially fair and transferable

representations. In: *ICML 2018: International Conference on Machine Learning.* [22% acceptance rate].

Wang, K, Vicol, P, Lucas, J, Gu, L, Grosse, R & Zemel, RS (2018). Adversarial distillation of Bayesian neural network posteriors. In: *ICML 2018: International Conference on Machine Learning.* [22% acceptance rate].

Kipf, T, Fetaya, E, Wang, K, Welling, M & Zemel, RS (2018). Neural relational inference for interacting systems. In: *ICML 2018: International Conference on Machine Learning.* [22% acceptance rate].

Ren, M, Triantafillou, E, Ravin, S, Swersky, K, Tenenbaum, J, Larochelle, H& Zemel, RS (2018). Meta-learning for semi-supervised few-shot classification. In: *ICLR 2018: International Conference on Learning Representations.* [24% acceptance rate].

Snell, J, Swersky, K & Zemel, RS (2017). Prototypical networks for few-shot learning. In: *NIPS 2017: Advances in Neural Information Processing Systems.* [23% acceptance rate].

Li, Y, Schwing, A, Wang, J, & Zemel, RS (2017). Dualing GANs. In: *NIPS 2017: Advances in Neural Information Processing Systems.* [23% acceptance rate].

Triantafillou, E, Zemel, RS & Urtasun, R (2017). Few-shot learning through an information retrieval lens. In: *NIPS 2017: Advances in Neural Information Processing Systems.* [23% acceptance rate].

Louizos, C, Shalit, Mooij, J, Sontag, D, Zemel, RS & Welling , M (2017). Causal effect inference with deep latent-variable models. In: *NIPS 2017: Advances in Neural Information Processing Systems.* [23% acceptance rate].

Ren, M & Zemel, RS (2017). End-to-end instance segmentation with recurrent attention. In: *CVPR 2017: 30th IEEE Conference on Computer Vision and Pattern Recognition.* [21% acceptance rate].

Snell, J & Zemel, RS (2017). Stochastic segmentation trees. In: *UAI 2017: Uncertainty in Artificial Intelligence.* [25% acceptance rate].

Ren, M, Liao, R, Urtasun, R, Sinz, F & Zemel, RS (2017). Normalizing the normalizers: Comparing and extending network normalization schemes. In: *ICLR 2017: International Conference on Learning Representations.* [24% acceptance rate].

Triantafillou, E, Kiros, R, Urtasun, R, & Zemel, RS (2016). Towards generalizable sentence embeddings. In: *ACL Workshop on Representation Learning for NLP.* [21% acceptance rate].

Luo, W, Li, Y, Urtasun, R & Zemel, RS (2016). Understanding the effective receptive field in deep convolutional neural networks. In: *NIPS 2016: Advances in Neural Information Processing Systems.* [22% acceptance rate].

Liao, R, Schwing, A, Zemel, RS & Urtasun, R (2016). Learning deep parsimonious representations. In: *NIPS 2016: Advances in Neural Information Processing Systems.* [22% acceptance rate].

Song, Y, Schwing, A, Zemel, RS & Urtasun, R (2016). Training deep neural networks via direct loss minimization In: *ICML 2016: 33rd International Conference on Machine Learning.* [24% acceptance rate].

Louizos, C, Li, Y, Swersky, K, Welling, M & Zemel, RS (2016). The variational fair autoencoder. In: *ICLR 2016: International Conference on Learning Representations.* [26% acceptance rate].

Li, Y, Tarlow, D, Brockschmidt, M & Zemel, RS (2016). Gated graph sequence neural networks. In: *ICLR 2016: International Conference on Learning Representations.* [26% acceptance rate].

Ren, M, Kiros, R & Zemel, RS (2015). Image question answering: A visual semantic embedding model and a new dataset. In: *NIPS 2015: Advances in Neural Information Processing Systems.* [23% acceptance rate].

Kiros, R, Zhu, Y, Salakhutdinov, R, Zemel, RS, Torralba, A, Urtasun, R, & Fidler, S (2015). Skip-thought vectors. In: *NIPS 2015: Advances in Neural Information Processing Systems.* [23% acceptance rate].

Zhu, Y, Kiros, R, Zemel, RS, Salakhutdinov, R, Urtasun, R, Torralba, A & Fidler, S (2015). Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In: *ICCV 2015: International Conference on Computer Vision.* [22% acceptance rate].

Li, Y, Swersky, K & Zemel, RS (2015). Generative moment matching networks. In: *ICML 2015: 32nd International Conference on Machine Learning.* [23% acceptance rate].

Xu, K, Ba, J, Kiros, R, Cho, K, Courville, A, Salakhutdinov, R, Zemel, RS, & Bengio, Y (2015). Show, attend and tell: Neural image caption generation with visual attention. In: *ICML 2015: 32nd International Conference on Machine Learning.* [23% acceptance rate].

Kiros, R, Zemel, RS, & Salakhtudinov, R (2014). A multiplicative model for learning distributed text-based attribute representations. In: *NIPS 2014: Advances in Neural Information Processing Systems.* [24% acceptance rate].

Charlin, L, Larochelle, H, & Zemel, RS (2014). Leveraging user libraries to bootstrap collaborative filtering. In: *KDD 2014: 20th ACM Conference on Knowledge Discovery and Data Mining.* [14% acceptance rate].

Li, Y & Zemel, RS (2014). High order regularization for semi-supervised learning of structured output problems. In: *ICML 2014: 31st International Conference on Machine Learning.* [22% acceptance rate].

Kiros, R, Zemel, RS & Salakhtudinov, R (2014). Multimodal neural models. In: *ICML 2014: 31st International Conference on Machine Learning.* [22% acceptance rate].

Snoek, J, Swersky, K, Zemel, RS & Adams, R (2014). Input warping for Bayesian optimization of non-stationary functions. In: *ICML 2014: 31st International Conference on Machine Learning.* [22% acceptance rate].

Snoek, J, Adams, R, & Zemel, RS (2013). A Determinantal Point Process latent variable model for inhibition in neural spiking data. In: *NIPS 2013: Advances in Neural Information Processing Systems.* [25% acceptance rate].

Martens, J, Pitassi, T, Chattopadhyay, A, & Zemel, RS (2013). On the expressive power of Restricted Boltzmann Machines. In: *NIPS 2013: Advances in Neural Information Processing Systems.* [25% acceptance rate].

Volkovs, M & Zemel, RS (2013). Supervised CRF framework for preference aggregation. In: *CIKM-2013: International Conference on Information and Knowledge Management.* [17% acceptance rate].

Zemel, RS, Pitassi, T, Yu, W, Swersky, K, & Dwork, C (2013). Learning fair representations. In: *ICML-2013: 30th International Conference on Machine Learning.* [24% acceptance rate].

Swersky, K, Tarlow, D, Charlin, L, Sutskever, I, & Zemel, R (2013). Stochastic k-neighborhood selection for supervised and unsupervised learning. In: *ICML-2013: 30th International Conference on Machine Learning.* [24% acceptance rate].

Li, Y, Tarlow, D, & Zemel, R (2013). Exploring compositional high order pattern potentials for structured output learning. In: *CVPR-2013: 26th IEEE Conference on Computer Vision and Pattern Recognition.* [25% acceptance rate].

Denton, E, & Zemel, R (2013). Conditional random fields for spiking populations of neurons. In: *COSYNE-2013: 10th Conference on Computational and Systems Neuroscience.* [65% acceptance rate].

Snoek, J, Adams, R, & Zemel, R (2013). Modeling inter-neuron inhibition with determinantal point processes In: *COSYNE-2013: 10th Conference on Computational and Systems Neuroscience.* [65% acceptance rate].

Adams, R, Hinton, GE, & Zemel, R (2013). Unsupervised learning of latent spiking representations In: *COSYNE-2013: 10th Conference on Computational and Systems Neuroscience.* [65% acceptance rate].

Volkovs, M & Zemel, RS (2012). Collaborative ranking with 17 parameters. In: *NIPS-2012: Advances in Neural Information Processing Systems.* [26% acceptance rate].

Swersky, K, Tarlow, D, Zemel, RS & Frey, B (2012). Probabilistic n-choose-k models for classification and ranking. In: *NIPS-2012: Advances in Neural Information Processing Systems.* [26% acceptance rate].

Volkovs, M & Zemel, RS (2012). Efficient sampling for bipartite matching problems. In: *NIPS-2012: Advances in Neural Information Processing Systems.* [26% acceptance rate].

Swersky, K, Tarlow, D, Sustkever, I, Zemel, RS & Adams, R (2012). Cardinality restricted Boltzmann machines. In: *NIPS-2012: Advances in Neural Information Processing Systems.* [26% acceptance rate].

Volkovs, M, Larochelle, H, & Zemel, RS (2012). Learning to rank by aggregating expert preferences. In *CIKM-2012: International Conference on Information and Knowledge Management.* [22% acceptance rate].

Tarlow, D, Swersky, K, Zemel, RS, & Adams, R (2012). Fast exact inference for recursive cardinality models. In *The 28th Conference on Uncertainty in Artificial Intelligence (UAI 2012).* [29% acceptance rate].

Charlin, L, Zemel, RS, & Boutilier, C. (2012). Active learning for matching problems. In: *29th International Conference on Machine Learning (ICML 2012).* [27% acceptance rate].

Volkovs, M & Zemel, RS (2012). A flexible generative model for preference aggregation. In: *WWW 2012: 21st International World Wide Web Conference.* [12% acceptance rate].

Tarlow, D, Adams, R & Zemel, RS (2012). Randomized optimum models for structured prediction. In: *The 15th International Conference on Artificial Intelligence and Statistics (AISTATS 2012).* [28% acceptance rate].

Tarlow, D & Zemel, RS (2012). Learning with high order loss functions. In: *The 15th International Conference on Artificial Intelligence and Statistics (AISTATS 2012).* [28% acceptance rate].

Dwork, C, Hardt, M, Pitassi, T, Reingold, O, & Zemel, RS (2012). Fairness through awareness. In: *The 3rd International Conference on Innovations in Theoretical Computer Science (ITCS 2012).*

Tarlow, D, Givoni, I, Zemel, RS, & Frey, B (2011). Graph-cuts is a max-product algorithm. In: *UAI 2011: 27th Conference on Uncertainty in Artificial Intelligence.* [32% acceptance rate]. *Runner-up, Best student paper award.*

Charlin, L, Zemel, RS, & Boutilier, C (2011). A framework for optimizing paper matching. In: *UAI 2011: 27th Conference on Uncertainty in Artificial Intelligence.* [32% acceptance rate].

Marlin, B, Zemel, RS, Roweis, S, & Slaney, M (2011). Recommender systems: Missing data and statistical model estimation In: *The 22nd International Joint Conference on Artificial Intelligence (IJCAI 2011).*

Tarlow, T, Givoni, I & Zemel, RS (2010). HOP-MAP: Efficient message passing with high order potentials. In: *The 13th International Conference on Artificial Intelligence and Statistics (AISTATS 2010).* [22% acceptance rate].

Marlin, B & Zemel, RS (2009). Collaborative prediction and ranking with non-random missing data. In: *RecSys-2009: 3rd ACM Conference on Recommender Systems.* [26% acceptance rate]. *Best paper award.*

Volkovs, M & Zemel, RS (2009). BoltzRank: Learning to maximize expected ranking gain. In: *ICML-2009: Proceedings of the 24th International Conference on Machine Learning.* [21% acceptance rate]. *Best student paper award.*

Natarajan, R, Murray, I, Shams, L, & Zemel, RS (2009). Validating a Bayesian model of conflicting sensory inputs. In: *COSYNE-09: Computational and Systems Neuroscience Conference* [6% oral acceptance rate].

He, X & Zemel, RS (2008). Learning hybrid models for image annotation with partially labeled data. In: *NIPS-2008: Advances in Neural Information Processing Systems.* [24% acceptance rate].

Natarajan, R, Murray, I, Shams, L, & Zemel, RS (2008). Comparing model predictions of response bias and variance in cue combination. In: *NIPS-2008: Advances in Neural Information Processing Systems.* [24% acceptance rate].

Schmah, T, Hinton, G, & Zemel, RS (2008). Competing RBM density models for classification of fMRI images. In: *NIPS-2008: Advances in Neural Information Processing Systems.* [24% acceptance rate].

Ross, D, Tarlow, D, & Zemel, RS (2008). Unsupervised learning of skeletons from motion. In: *ECCV-2008: European Conference on Computer Vision.* [26% acceptance rate].

Tarlow, D, Zemel, RS, & Frey, B (2008). Flexible priors for exemplar-based clustering. In: *UAI-2008: 24th Conference on Uncertainty in Artificial Intelligence.* [29% acceptance rate].

He, X & Zemel, RS (2008). Latent topic random fields: Learning using a taxonomy of labels. In: *CVPR-2008: IEEE Conference on Computer Vision and Pattern Recognition* [31% acceptance rate]. Los Alamitos, CA, IEEE Computer Society.

Meeds, E, Ross, D, Roweis, S. & Zemel, RS (2008). Learning stick-figure models using nonparametric Bayesian priors over trees. In: *CVPR-2008: IEEE Conference on Computer Vision and Pattern Recognition* [31% acceptance rate]. Los Alamitos, CA, IEEE Computer Society.

Marlin, B, Zemel, RS, Roweis, S & Slaney, M (2007). Collaborative filtering and the missing at random assumption. In: *UAI-2007: 23rd Conference on Uncertainty in Artificial Intelligence.* [31% acceptance rate].

Huys Q, & Zemel, RS, Natarajan, R, & Dayan, P (2007). Gaussian priors for population coding. In: *CNS-07: Computational Neuroscience Conference* [19% acceptance rate].

He, X, Zemel, RS, & Ray, D (2006). Learning and incorporating top-down cues in image segmentation. In: *ECCV-2006: European Conference on Computer Vision*, LNCS 2951, (pp. 338-351). [22% acceptance rate].

Snoek, J, Hoey, J, Stewart, L & Zemel, RS. (2006). Automated detection of unusual events on stairs. In: *Canadian Conference on Computer and Robot Vision* (pp. 5-13). [24% acceptance rate].

Natarajan, R, Huys, Q, Dayan, P & Zemel, RS (2006). Population codes for dynamic cue combination. In: *COSYNE-06: Computational and Systems Neuroscience Conference* [4% oral acceptance rate].

Ross, D, Osindero, S, & Zemel, RS (2006). Combining discriminative features to infer complex trajectories. In: *ICML-2006: Proceedings of the 21st International Conference on Machine Learning* (pp.

761-769). [19% acceptance rate].

Natarajan, R, Huys, Q, Dayan, P & Zemel, RS (2006). Population codes for natural dynamic stimuli. In: *CNS-06: Computational Neuroscience Conference.* [34% acceptance rate].

Marlin, B, Roweis, S, & Zemel, RS (2005). Unsupervised learning with non-ignorable missing data. In: *Proceedings of the Tenth International Workshop on Artificial Intelligence and Statistics*, C. M. Bishop & B. J. Frey (Eds.). [36% acceptance rate].

Zemel, RS, Huys, Q, Natarajan, R, & Dayan, P (2004). Probabilistic computation in spiking populations. In: *NIPS-17: Advances in Neural Information Processing Systems 17.* Cambridge, MA, MIT Press. (pp. 1609-1616). [25% acceptance rate].

Carreira-Perpinan, M & Zemel, RS (2004). Proximity graphs for clustering and manifold learning. In: *NIPS-17: Advances in Neural Information Processing Systems 17.* Cambridge, MA, MIT Press. (pp. 225-232). [25% acceptance rate].

Marlin, B & Zemel, RS (2004). The multiple multiplicative factor model for collaborative filtering. In: *ICML-2004: Proceedings of the 21st International Conference on Machine Learning.*

He, X, Zemel, RS, & Carreira-Perpinan, M (2004). Multiscale conditional random fields for image labelling. In: *CVPR-2004: IEEE Conference on Computer Vision and Pattern Recognition.* (pp. 695-703). Los Alamitos, CA, IEEE Computer Society.

Boutilier, C, Zemel, RS, & Marlin, B (2003). Active collaborative filtering. In: *UAI-2003: 19th Conference on Uncertainty in Artificial Intelligence.* (pp. 98-106). San Francisco, CA, Morgan Kaufmann.

Welling, M, Zemel, RS, & Hinton, GE (2003). Efficient parametric projection pursuit density estimation. In: *UAI-2003: 19th Conference on Uncertainty in Artificial Intelligence.* (pp. 575-583). San Francisco, CA, Morgan Kaufmann.

Boutilier, C & Zemel, RS (2003). Online queries for collaborative filtering. In: *Proceedings of the Ninth International Workshop on Artificial Intelligence and Statistics*, C. M. Bishop & B. J. Frey (Eds.).

Ross, DA & Zemel, RS (2002). Multiple-cause vector quantization. In: *NIPS-15: Advances in Neural Information Processing Systems 15.* (pp. 521-528). Cambridge, MA, MIT Press.

Welling, M, Zemel, RS, & Hinton, GE (2002). Self supervised boosting. In: *NIPS-15: Advances in Neural Information Processing Systems 15.* (pp. 878-885). Cambridge, MA, MIT Press.

Zemel, RS, Hungerford, M, & Mozer, MC. (2002). Experience effects in neural coding. In: *Proceedings of 2002 Neural Information and Coding Conference.*

Zemel, RS & Pitassi, T. (2001). A gradient-based boosting algorithm for regression problems. In: *NIPS-13: Advances in Neural Information Processing Systems 13.* (pp. 696-702). T. K. Leen, T. G. Dietterich, & V. Tresp (Eds.). Cambridge, MA, MIT Press.

Cowen, S, Zemel, RS, Kudrimoti, H, Gerrard, J, & McNaughton, B (2000). Spike train segmentation and recognition: Identifying the neural correlate of decisions. *Proceedings of the 2000 Computational Neuroscience Meeting.*

Zemel, RS & Mozer, MC (2000). A generative model for attractor dynamics. *NIPS-12: Advances in Neural Information Processing Systems 12.* (pp. 80-86). S. A. Solla, T. K. Leen & K.-R. Muller (Eds.). Cambridge, MA: MIT Press.

Yang, Z & Zemel, RS (2000). Managing uncertainty in cue combination. *NIPS-12: Advances in Neural Information Processing Systems 12* (pp. 869-875). S. A. Solla, T. K. Leen & K.-R. Muller (Eds.). Cambridge, MA: MIT Press.

Zemel, RS & Dayan, P (1999). Distributional population codes and multiple motion models. *NIPS-11: Advances in Neural Information Processing Systems 11.* (pp. 174-180). Cambridge, MA: MIT Press.

Zemel, RS & Dayan, P (1997). Combining probabilistic population codes. *IJCAI-97: Fifteenth International Joint Conference on Artificial Intelligence* (pp. 1114-1119). San Francisco, CA: Morgan Kaufmann.

Zemel, RS, Dayan, P, & Pouget, A (1997). Probabilistic interpretation of population codes. *NIPS-9: Advances in Neural Information Processing Systems 9* (pp. 676-683). M. C. Mozer, M. I. Jordan, & T. Petsche (Eds.). Cambridge, MA: MIT Press.

Zemel, RS, & Sejnowski, TJ (1995). Grouping components of three-dimensional moving objects in area MST of visual cortex. In G. Tesauro, D. S. Touretzky, & T. K. Leen (Eds.), *NIPS-7: Advances in Neural Information Processing Systems 7* (pp. 165-172). Cambridge, MA: MIT Press.

Zemel, RS, & Hinton, GE (1994). Developing population codes by minimizing description length. In *NIPS-6: Advances in Neural Information Processing Systems 6* (pp. 11–18). J. D. Cowan, G. Tesauro, J. Alspector (Eds.). San Mateo, CA: Morgan Kaufmann.

Hinton, GE, & Zemel, RS (1994). Autoencoders, Minimum Description Length, and Helmholtz free energy. In *NIPS-6: Advances in Neural Information Processing Systems 6* (pp. 3–10). J. D. Cowan, G. Tesauro, J. Alspector (Eds.). San Mateo, CA: Morgan Kaufmann.

Zemel, RS, Williams, CKI, & Mozer, MC (1993). Directional-unit Boltzmann machines. In C. L. Giles, S. J. Hanson, & J. D. Cowan (Eds.), *NIPS-5: Advances in Neural Information Processing Systems 5* (pp. 172-179). San Mateo, CA: Morgan Kaufmann.

Mozer, MC, Zemel, RS & Behrmann, M (1992). Learning to segment images using dynamic feature binding. In J. E. Moody, S. J. Hanson, & R. P. Lippmann (Eds.), *NIPS-4: Advances in Neural Information Processing Systems 4* (pp. 436-443). San Mateo, CA: Morgan Kaufmann.

Mozer, MC, Zemel, RS & Behrmann, M (1992). Discovering and using perceptual grouping principles in visual information processing. *Proceedings of the 14th Annual Conference of the Cognitive Society.* Bloomington, IN: Lawrence Erlbaum.

Zemel, RS & Hinton, GE (1991). Discovering viewpoint-invariant relationships that characterize objects. In R. P. Lippmann, J. E. Moody, & D. S. Touretzky (Eds.), *NIPS-3: Advances in Neural Information Processing Systems 3* (pp. 299-305). San Mateo, CA: Morgan Kaufmann.

Zemel, RS, Mozer, MC, & Hinton, GE (1990). TRAFFIC: Object recognition using hierarchical reference frame transformations. In D. S. Touretzky, (Ed.), *NIPS-2: Advances in Neural Information Processing Systems 2* (pp. 266-273). San Mateo, CA: Morgan Kaufmann.

C. Book Chapters

Natarajan, R & Zemel, RS. (2011). Dynamic cue combination in distributional population code networks. In: M Landy (Ed.), *Sensory Cue Integration* (pp. 368-392). Oxford University Press.

Pouget, A & Zemel, RS. (2007). In: K Doya & K Ishii (Eds.), *Bayesian Brain: Probabilistic Approaches to Neural Coding* (pp. 115-130). MIT Press.

Zemel, RS (2003). Cortical belief networks. In: R. Hecht-Neilsen & T. McKenna (Eds.), *Computational Models for Neuroscience* (pp. 267-287). Springer-Verlag.

Zemel, RS & Pillow, J (2002). An information-theoretic objective for population codes. In: R. Rao & B. Olshausen (Eds.) *Statistical Theories of Cortical Function* (pp. 223-242). MIT Press,

Zemel, RS (2002). Minimizing description length and neural networks. To appear in: M. A. Arbib (Ed.), *The Handbook of Brain Theory and Neural Networks, Volume II*. MIT Press.

Becker, S & Zemel, RS (2002). Unsupervised learning with global objective functions. To appear in: M. A. Arbib (Ed.), *The Handbook of Brain Theory and Neural Networks, Volume II*. MIT Press.

D. LIST OF COURSES

A. Undergraduate Courses Taught

CSC 412: Probabilistic Graphical Models

CSC 411: Machine Learning and Data Mining

CSC 384: Introduction to Artificial Intelligence

B. Graduate Courses Taught

CSC 2515: Introduction to Machine Learning

CSC 2535: Advanced Machine Learning

CSC 2541: Natural Scene Statistics

Career Student Numbers		
	In Progress	Completed
MSc	2	21
PhD	9	12
Postdoctoral Fellows	3	6

C. PhD Student Supervision

Elliot Creager, Dept. of Computer Science, Univ. of Toronto, PhD Advisor. *Learning adversarially fair and transferrable representations.* 2017-present (ongoing).

Mengye Ren, Dept. of Computer Science, Univ. of Toronto, PhD Advisor. *Learning with few labels.* 2017-present (ongoing).

James Lucas, Dept. of Computer Science, Univ. of Toronto, PhD Co-Advisor. *Novel optimization methods for deep learning.* 2017-present (ongoing).

Will Grathwohl, Dept. of Computer Science, Univ. of Toronto, PhD Co-Advisor. *Learning discrete latent variable models.* 2017-present (ongoing).

David Madras, Dept. of Computer Science, Univ. of Toronto, PhD Advisor. *Predicting responsibly.* 2016-present (ongoing).

Jackson Wang, Dept. of Computer Science, Univ. of Toronto, PhD Advisor. *Learning to represent uncertainty in deep networks.* 2016-present (ongoing).

Renjie Liao, Dept. of Computer Science, Univ. of Toronto, PhD Co-Advisor. *Graph Neural Networks.* 2015-present (ongoing).

Eleni Triantafillou, Dept. of Computer Science, Univ. of Toronto, PhD Co-Advisor. *Meta-learning of new categories.* 2015-present (ongoing).

Jake Snell, Dept. of Computer Science, Univ. of Toronto, PhD Advisor. *Learning to Produce Multiple Valid Responses.* 2014-present (ongoing).

Yujia Li, Dept. of Computer Science, Univ. of Toronto, PhD Advisor. *Learning Shape Models for Image Labeling.* Defended 2017.

Jamie Ryan Kiros, Dept. of Computer Science, Univ. of Toronto, PhD Co-Advisor. *Multimodal Neural Language Models.* Defended 2017.

Kevin Swersky, Dept. of Computer Science, Univ. of Toronto, PhD Advisor. *Efficient Learning of Sparse Models.* Defended 2017.

James Martens, Dept. of Computer Science, Univ. of Toronto, PhD Advisor. *Second-order optimization methods for neural networks.* Defended 2015.

Laurent Charlin, Dept. of Computer Science, Univ. of Toronto, PhD Co-Advisor. *Supervised and Active Learning for Recommender Systems.* Defended 2014.

Maksims Volkovs, Dept. of Computer Science, Univ. of Toronto, PhD Advisor. *Machine Learning Methods and Models for Ranking.* Defended 2013.

Daniel Tarlow, Dept. of Computer Science, Univ. of Toronto, PhD Advisor. *Efficient Machine Learning with High Order and Combinatorial Structures.* Defended 2013.

Rama Natarajan, Dept. of Computer Science, Univ. of Toronto, PhD Advisor. *Neural representation, learning and manipulation of uncertainty.* Defended 2009.

Benjamin Marlin, Dept. of Computer Science, Univ. of Toronto, PhD Advisor. *Missing Data Problems in Machine Learning.* Defended 2008.

David Ross, Dept. of Computer Science, Univ. of Toronto, PhD Advisor. *Learning Probabilistic Models for Visual Motion.* Defended 2008.

Xuming He, Dept. of Computer Science, Univ. of Toronto, PhD Advisor. *Learning Structured Models for Image Labeling.* Defended 2007.

Stephen Cowen, Dept. of Psychology, Univ. of Arizona, PhD Co-Advisor. *Operant Conditioning of Neural Population Activity.* Defended 2006.

MSc Student Supervision

Seyed Seyed Ghasemipour, Dept. of Computer Science, Univ. of Toronto, MSc Advisor. *Option-based reinforcement learning.* 2017-present (ongoing).

Marc-Etienne Brunet, Dept. of Computer Science, Univ. of Toronto, MSc Co-Advisor. *Learning to unbias word representations.* 2017-present (ongoing).

Lisa Zhang, Dept. of Computer Science, Univ. of Toronto, MSc Co-Advisor. *Neural guided constraint logic programming for program synthesis.* Defended, May, 2018.

David Madras, Dept. of Computer Science, Univ. of Toronto, MSc Advisor. *Learning fair and transferable representations.* Defended, January, 2018.

Eleni Triantafillou, Dept. of Computer Science, Univ. of Toronto, MSc Co-Advisor. *Transfer learning across natural language tasks.* Defended, January, 2017.

Mengye Ren, Dept. of Computer Science, Univ. of Toronto, MSc Advisor. *Learning to count objects in images*, Defended, January, 2017.

Jackson Wang, Dept. of Computer Science, Univ. of Toronto, MSc Advisor. *Learning Basketball Plays.* Defended, January, 2016.

Gregory Koch, Dept. of Computer Science, Univ. of Toronto, MSc Co-Advisor. *Zero-shot learning of complex categories.* Defended, April, 2014.

Guang Yu, Dept. of Computer Science, Univ. of Toronto, MSc Applied Computing. *Collaborative ranking.* Defended, December, 2014.

Jake Snell, Dept. of Computer Science, Univ. of Toronto, MSc Advisor. *Learning with Multiple Ground Truths*. Defended, January, 2014.

Sagan Bolliger, Dept. of Computer Science, Univ. of Toronto, MSc Co-Advisor. *Learning to Represent Novel Stimuli with Spiking Population Codes*. Defended, January, 2014.

Qingwei Ge, Dept. of Computer Science, Univ. of Toronto, MSc Applied Computing. *Book recommendation: Improving collaborative filtering with content information*. Defended, December, 2013.

Yujia Li, Dept. of Computer Science, Univ. of Toronto, MSc Advisor. *Exploring Compositional High Order Pattern Potentials for Structured Output Learning*. Defended, January, 2013.

Nikola Karamanov, Dept. of Computer Science, Univ. of Toronto, MSc Advisor. *Inferring Depth for Image Labelling*. Defended, April, 2011.

James Martens, Dept. of Computer Science, Univ. of Toronto, MSc Co-Advisor. *Efficient Approximations to Dynamical Systems*. Defended April, 2009.

Maksims Volkovs, Dept. of Computer Science, Univ. of Toronto, MSc Advisor. *Learning Ranking Metrics*. Defended January, 2009.

Daniel Tarlow, Dept. of Computer Science, Univ. of Toronto, MSc Advisor. *Flexible Priors for Exemplar-Based Clustering*. Defended January, 2008.

Stephen Fung, Dept. of Computer Science, Univ. of Toronto, MSc Advisor. *Learning Intrinsic Image Decomposition*. Defended September, 2007.

Liam Stewart, Dept. of Computer Science, Univ. of Toronto, MSc Advisor. *Learning Flexible Features for Sequence Labeling*. Defended September, 2005.

Benjamin Marlin, Dept. of Computer Science, Univ. of Toronto, MSc Advisor, *Machine Learning Methods for Collaborative Filtering*. Defended April, 2004.

David Ross, Dept. of Computer Science, Univ. of Toronto, MSc Advisor. *Learning Parts-Based Representations of Data*. Defended April, 2003.

David Towers, Dept. of Psychology, Univ. of Arizona, MSc Advisor. *Learning in Object-Based Attention*. Defended March, 2002.

Stephen Cowen, Dept. of Psychology, Univ. of Arizona, MSc Advisor. *Decoding Hippocampal Population Codes for Navigation*. Defended August, 2000.

Postdoctoral Fellow Supervision

Joern-Henrik Jcobsen, Dept. of Computer Science, University of Toronto, 2018-present.

Ethan Fetaya, Dept. of Computer Science, University of Toronto, 2017-present.

Marc Law, Dept. of Computer Science, University of Toronto, 2016-present.

Tanya Schmah, Dept. of Computer Science, University of Toronto, 2007-2014.

Amit Gruber, Dept. of Computer Science, University of Toronto, 2009-2011.

Miguel Carreira-Perpinan, Dept. of Computer Science, University of Toronto, 2003-2005.

Daniel Roobaert, Dept. of Computer Science, University of Toronto, 2001-2002.

Zhiyong Yang, Dept. of Psychology, University of Arizona, 1998-2000.

PhD Committees

Nikhil Bhagwat, Dept. of Computer Science, Univ. of Toronto, 2014-present (ongoing).

Elias AbouZeid, Dept. of Computer Science, Univ. of Toronto, 2014-present (ongoing).

Jian Yao, Dept. of Computer Science, Univ. of Toronto, 2014-present (ongoing).

Wenjie Luo, Dept. of Computer Science, Univ. of Toronto, 2014-present (ongoing).

Kaustav Kundu, Dept. of Computer Science, Univ. of Toronto, 2014-present (ongoing).

Libby Barak, Dept. of Computer Science, Univ. of Toronto, 2012-present (ongoing).

Amirali Salehi-Abari, Dept. of Computer Science, Univ. of Toronto, 2012-present (ongoing).

Tom Lee, Dept. of Computer Science, Univ. of Toronto, 2011-2016.

Charlie Tang, Dept. of Computer Science, Univ. of Toronto, 2011-2015.

Aida Nematzadeh, Dept. of Computer Science, Univ. of Toronto, 2011-2015.

Alex Krizhevsky, Dept. of Computer Science, Univ. of Toronto, 2009-present (ongoing).

Stefan Mathe, Dept. of Computer Science, Univ. of Toronto, 2010-2015.

George Dahl, Dept. of Computer Science, Univ. of Toronto, 2010-2015.

Tyler Lu, Dept. of Computer Science, Univ. of Toronto, 2008-2014.

Tijmen Tieleman, Dept. of Computer Science, Univ. of Toronto, 2007-2014.

Kevin Regan, Dept. of Computer Science, Univ. of Toronto, 2006-2013.

Grigori Yourganov, Dept. of Physiology, Rotman Medical Institute, 2006-2013.

Vlad Mnih, Dept. of Computer Science, Univ. of Toronto, 2008-2013.

Darius Braziunas, Dept. of Computer Science, Univ. of Toronto, 2006-2012.

Gerald Quon, Dept. of Computer Science, Univ. of Toronto, 2006-2012.

Chris Parisien, Dept. of Computer Science, Univ. of Toronto, 2006-2011.

Bowen Hui, Dept. of Computer Science, Univ. of Toronto, 2006-2011.

Russ Salaktudinov, Dept. of Computer Science, Univ. of Toronto, 2005-2009.

Dustin Lang, Dept. of Computer Science, Univ. of Toronto, 2004-2009.

Graham Taylor, Dept. of Computer Science, Univ. of Toronto, 2004-2010.

Roland Memisevic, Dept. of Computer Science, Univ. of Toronto, 2002-2008.

Afra Alishahi, Dept. of Computer Science, Univ. of Toronto, 2002-2008.

Anitha Kannan, Dept. of Computer Science, Univ. of Toronto, 2000-2006.

Kannan Achan, Dept. of Computer Science, Univ. of Toronto, 2000-2006.

Vincent Cheung, Dept. of Electrical & Computer Engineering, Univ. of Toronto, 2003-2009.

Jesse Gillis, Dept. of Neurophysiology, Univ. of Toronto, 2002-2007.

Francisco Estrada, Dept. of Computer Science, Univ. of Toronto, 2000-2005.

Gustavo Carneiro, Dept. of Computer Science, Univ. of Toronto, 2000-2005.

Pascal Poupart, Dept. of Computer Science, Univ. of Toronto, 2000-2007.

Hai Wang, Dept. of Computer Science, Univ. of Toronto, 2000-2003 (defended October, 2003).

Yee Whye Teh, Dept. of Computer Science, Univ. of Toronto, 2000-2003 (defended April, 2003).

Richard Murray, Dept. of Psychology, Univ. of Toronto, 2000-2002 (defended September, 2002).

Jason Gold, Dept. of Psychology, Univ. of Toronto, 2000-2002 (defended April, 2002).

Alberto Paccanaro, Dept. of Computer Science, Univ. of Toronto, 2000-present (defended March, 2002).

Karen Reinke, Dept. of Psychology, Univ. of Arizona Member, 1996-1998 (defended October, 1998).

D. Other Teaching

Computational Neuroscience Course

Okinawa Institute of Technology; guest faculty member for two week course

Summer School Course

Neural Computation (Complex Systems Summer School, Santa Fe Institute;
guest faculty member for week-long course)

E. ADMINISTRATIVE POSITIONS

University Service

Research Director & Co-Founder, Vector Institute for Artificial Intelligence, 2017-present.

Chief Scientist, Creative Destruction Lab, Rotman School of Management, Univ. of Toronto, 2015-present.

Chair, Graduate Recruiting Committee, Univ. of Toronto Computer Science Department, 2011-2012.

Awards Committee, Univ. of Toronto Computer Science Department, 2011-2012.

Graduate Admissions Committee, Univ. of Toronto Computer Science Department, 2010-2011.

Professional Masters Committee, Univ. of Toronto Computer Science Department, 2010-2011.

Chair, Graduate Affairs Committee, Univ. of Toronto Computer Science Department, 2009-2010.

Associate Chair, Graduate Studies, Univ. of Toronto Computer Science Department, 2005-2008.

School of Graduate Studies Awards Committee, Univ. of Toronto, 2006-2008.

Graduate Affairs Committee, Univ. of Toronto Computer Science Department, 2007-2008.

Graduate Admissions Committee, Univ. of Toronto Computer Science Department, 2006-2008.

Chair, Graduate Admissions Committee, Univ. of Toronto Computer Science Department, 2004-2006.

Chair, Graduate Affairs Committee, Univ. of Toronto Computer Science Department, 2004-2007.

Undergraduate Committee, Univ. of Toronto Computer Science Department, 2001-2003.

Graduate Committee, Univ. of Toronto Computer Science Department, 2000-2001.

Graduate Admissions Committee, Univ. of Arizona Psychology Department, Cognition & Neural Systems, 1998-2000.

Computing Resources Committee, Univ. of Arizona Psychology Department, 1998-2000.

Referee

Grant proposals: National Science Foundation, National Science and Engineering Research Council, Mathematics of Information Technology and Complex Systems Network: Centers of Excellence, National Institutes of Health, National Science Foundation

Journal papers: *Nature; Science; Machine Learning; International Journal of Computer Vision; Neural Computation; Vision Research; Journal of Experimental Psychology; Cerebral Cortex; Cognitive Science; Graphics and Image Processing: Image Understanding; Neural Networks; IEEE Transactions on Neural Networks; Journal of Neuroscience*

Conference papers: International Joint Conference on Artificial Intelligence (IJCAI); National Conference on Artificial Intelligence (AAAI); Neural Information Processing Society (NIPS); International Conference on Computer Vision (ICCV); International Conference on Artificial Neural Networks (ICANN)