

GRAHAM TAYLOR

Thornbrough Building, Room 2413
School of Engineering
University of Guelph
Guelph, Ontario, N1G 2W1 Canada

T +1 519 824 4120 x53644
F +1 519 836 0227
[account name in url below]@uoguelph.ca
<http://www.uoguelph.ca/~gwtaylor>

RESEARCH INTERESTS

Statistical machine learning with an emphasis on unsupervised learning and time series analysis.

Application areas

Motion capture and analysis • Video and image understanding • Computer graphics
Climate modeling • Speech and natural language processing • Finance • Crowd gaming

EDUCATION

- 2009 **PhD in Computer Science**
University of Toronto
Composable, distributed-state models for high-dimensional time series
Advisors: Geoffrey Hinton and Sam Roweis
- 2004 **MASc in Systems Design Engineering**
University of Waterloo
Advisor: Hamid Tizhoosh
- 2003 **BASc in Systems Design Engineering (with distinction)**
University of Waterloo

RESEARCH EXPERIENCE

- 2012-present **School of Engineering, University of Guelph** Guelph, Canada
Assistant Professor (current appointment)
- 2009-2012 **Courant Institute of Mathematical Sciences, New York University** New York, USA
Research Scientist
Research with Chris Bregler, Rob Fergus, and Yann LeCun of the Vision, Learning, and Graphics group. Co-advising graduate students.
- 2007 **Speech Technology Group, Microsoft Research** Redmond, USA
Research Intern
Research with Michael Seltzer, Li Deng, and Alex Acero. Developed a novel algorithm for the separation of acoustic sources.
- 2003-2004 **INSA de Lyon Technical and Scientific University** Lyon, France
Visiting Scientist
Research with Christian Wolf and Jean-Michel Jolion of the Lyon Research Centre for Images and Intelligent Information Systems.

INDUSTRY EXPERIENCE

- 2014-present **Investment Industry Regulatory Organization of Canada (IIROC)** Toronto, Canada
Consultant
- Applied machine learning to gain insight from Canada's largest dataset of equities trading
- 2011-2012 **OANDA** Toronto, Canada
Quantitative Analyst
- Applied machine learning to build statistical models of financial transactions and web click-stream data
- 2008-2009 **Mixamo** San Francisco, USA
Consultant (Part of founding team)
- Developed generative models of human motion for the animation community
 - Transferred research into a commercially successful product

Industry experience prior to 2008 available upon request

TEACHING EXPERIENCE

Courses taught, University of Guelph

- 2013-2014 **ENGG 6500: Machine Learning** (20 students, graduate)
Taught introductory applied machine learning course.
- 2012-2014 **ENGG 4450: Large-Scale Software Architecture Engineering** (25 students, 4th year)
Emphasized agile software processes, open-source, and distributed development.
- 2013-2015 **ENGG 1210: Engineering Mechanics** (240 students, 1st year)
Taught core course to engineering students from all disciplines.

Teaching assistantships, University of Toronto

- 2008 **CSC 411: Machine Learning and Data Mining** (30 students, 4th year)
Assisted instructor with preparation of course notes, prepared and delivered lectures, conducted tutorials, marked exams/assignments and offered individual help.
- 2005,2008 **CSC 321: Machine Learning and Neural Networks** (60 students, 3rd year)
Prepared and delivered lectures, conducted tutorials, marked exams/assignments and offered individual help.
- 2008 **CSC 108: Introduction to Computer Programming** (300 students, 1st year)
Conducted tutorials, marked exams/assignments and offered individual help.
- 2005-2007 **CSC 104: The How and Why of Computing** (100 students, 1st year)
Conducted tutorials, marked exams/assignments and offered individual help. Taught computing concepts to students from diverse academic backgrounds.
- 2004 **CSC 180: Introduction to Computer Programming** (100 students, 1st year)
Conducted tutorials, marked exams/assignments and offered individual help. Assisted students who had no programming experience understand difficult and abstract concepts.

Teaching experience prior to 2004 available upon request

Formal training

- 2008 **THE500: Teaching in Higher Education** (Certificate)
Office of Teaching Advancement, University of Toronto

PUBLICATIONS

Under Review

Available upon request

Preprints

1. Jan Rudy, Weiguang Ding, Jiwoong Im, and Graham Taylor. Neural network regularization via robust weight factorization. *arXiv preprint arXiv:1412.6630*, 2014.
2. Jiwoong Im, Ethan Buchman, and Graham Taylor. Understanding minimum probability flow for RBMs under various kinds of dynamics. *arXiv preprint arXiv:1412.6617*, 2014.
3. Jan Rudy and Graham Taylor. Generative class-conditional autoencoders. *arXiv preprint arXiv:1412.7009*, 2014.

Refereed Conference Proceedings

4. Natalia Neverova, Christian Wolf, and Graham Taylor. Hand segmentation into parts with convolutional learning. In *Asian Conference on Computer Vision (ACCV)*, 2014.
5. Arjun Jain, Jonathan Tompson, Mykhaylo Andriluka, Graham Taylor, and Christoph Bregler. Learning human pose estimation features with convolutional networks. In *International Conference on Learning Representations (ICLR)*, 2014.
6. Terrance Devries, Kumar Biswaranjan, and Graham Taylor. Multi-task learning of facial landmarks and expression. In *14th Canadian Conference on Computer and Robot Vision (CRV)*, 2014.
7. Matthew Zeiler, Graham Taylor, Leonid Sigal, Iain Matthews, and Rob Fergus. Input-output temporal restricted Boltzmann machines for facial expression transfer. In *Advances in Neural Information Processing Systems 24 (NIPS)*, 2011. (Accepted: 305/1400, 22%)
8. Matthew Zeiler, Graham Taylor, and Rob Fergus. Adaptive deconvolutional networks for mid and high level feature learning. In *Proc. of the 13th International Conference on Computer Vision (ICCV)*, 2011. (Accepted: 340/1285, 26%)
9. Graham Taylor, Ian Spiro, Christoph Bregler, and Rob Fergus. Learning invariance through imitation. In *Proc. of the 24th IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011. (Accepted: 440/1677, 26%)
10. Isabelle Guyon, Gideon Dror, Vincent Lemaire, Daniel Silver, Graham Taylor, and David Aha. Unsupervised and transfer learning challenge. In *Proc. of the International Joint Conference on Neural Networks (IJCNN)*, 2011.
11. Graham Taylor, Rob Fergus, Ian Spiro, George Williams, and Christoph Bregler. Pose-sensitive embedding by nonlinear NCA regression. In *Advances in Neural Information Processing Systems 23 (NIPS)*, pages 2280–2288, 2010. (Accepted: 293/1219, 24%)
12. Graham Taylor, Rob Fergus, Yann LeCun, and Christoph Bregler. Convolutional learning of spatio-temporal features. In *Proc. of the 11th European Conference on Computer Vision (ECCV)*, pages 140–153, 2010. (Accepted: 325/1174, 28%)
13. George Williams, Graham Taylor, Kirill Smolskiy, and Christoph Bregler. Body motion analysis for multi-modal identity verification. In *Proc. of the 20th International Conference on Pattern Recognition (ICPR)*, 2010. (Accepted for plenary oral: 385/2140, 18%; Overall: 1147/2140, 54%)
14. Graham Taylor, Leonid Sigal, David Fleet, and Geoffrey Hinton. Dynamical binary latent variable models for 3d human pose tracking. In *Proc. of the 23rd IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010. (Accepted for plenary oral: 78/1724, 5%; Overall: 461/1724 27%)

15. Matthew Zeiler, Dilip Krishnan, Graham Taylor, and Rob Fergus. Deconvolutional networks. In *Proc. of the 23rd IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010. (Accepted: 461/1724, 27%)
16. Graham Taylor and Geoffrey Hinton. Products of Hidden Markov Models: It takes N>1 to tango. In *Proc. of the 25th Conference on Uncertainty in Artificial Intelligence (UAI)*, 2009. (Accepted for plenary oral: 30/243, 12%; Overall: 76/243, 31%)
17. Graham Taylor and Geoffrey Hinton. Factored conditional restricted Boltzmann machines for modeling motion style. In *Proc. of the 26th International Conference on Machine Learning (ICML)*, pages 1025–1032, 2009. (Accepted: 160/595, 27%)
18. Matthew Zeiler, Graham Taylor, Niko Troje, and Geoffrey Hinton. Modeling pigeon behaviour using a conditional restricted Boltzmann machine. In *Proc. of the 17th European Symposium on Artificial Neural Networks (ESANN)*, 2009.
19. Ilya Sutskever, Geoffrey Hinton, and Graham Taylor. The recurrent temporal restricted Boltzmann machine. In *Advances in Neural Information Processing Systems 21 (NIPS)*, 2009. (Accepted: 250/1022, 25%)
20. Graham Taylor, Michael Seltzer, and Alex Acero. Maximum a posteriori ICA: Applying prior knowledge to the separation of acoustic sources. In *Proc. of the IEEE International Conference on Audio, Speech and Signal Processing (ICASSP)*, pages 1821–1824, 2008.
21. Graham Taylor, Geoffrey Hinton, and Sam Roweis. Modeling human motion using binary latent variables. In *Advances in Neural Information Processing Systems 19 (NIPS)*, pages 1345–1352, 2007. (Accepted for poster spotlight: 63/833, 8%; Overall 204/833, 24%)
22. Graham Taylor. A reinforcement learning framework for parameter control in computer vision applications. In *Proc. of the 1st Canadian Conference on Computer and Robot Vision (CRV)*, pages 496–503, 2004.
23. Graham Taylor and Christian Wolf. Reinforcement learning for parameter control of text detection in images from video sequences. In *Proc. of the 1st International Conference on Information and Communication Technologies*, pages 517–518, 2004.

Journal articles

24. Mingyuan Jiu, Christian Wolf, Graham Taylor, and Atilla Baskurt. Human body part estimation from depth images via spatially-constrained deep learning. *Pattern Recognition Letters*, 2013. In press.
25. Isabelle Guyon, Gideon Dror, Vincent Lemaire, Daniel Silver, Graham Taylor, and David Aha. Analysis of the IJCNN 2011 unsupervised and transfer learning challenge. *Neural Networks*, 32:174 – 178, 2012.
26. Graham Taylor, Geoffrey Hinton, and Sam Roweis. Two distributed-state models for generating high-dimensional time series. *Journal of Machine Learning Research*, 12(Mar):1025–1068, 2011.
27. Hamid Tizhoosh and Graham Taylor. Reinforced contrast adaptation. *International Journal of Image and Graphics*, 6(3):377–392, 2006.
28. Doris Miller, Alexandra Zecevic, and Graham Taylor. Hurdle preflight in springboard diving: A case of diminishing returns. *Research Quarterly for Exercise and Sport*, 73:134–145, 2002.

Refereed workshop papers

29. Weiguang Ding and Graham Taylor. Mental rotation by optimizing transforming distance. In *Neural Information Processing Systems 27 (NIPS) Workshop on Deep Learning and Representation Learning*, 2014.

30. Jiwoong Im and Graham Taylor. Analyzing the dynamics of gated auto-encoders. In *Neural Information Processing Systems 27 (NIPS) Workshop on Deep Learning and Representation Learning*, 2014.
31. Natalia Neverova, Christian Wolf, Graham Taylor, and Florian Nebout. Multi-scale deep learning for gesture detection and localization. In *ECCV ChaLearn Workshop on Looking at People*, 2014.
32. Natalia Neverova, Christian Wolf, Giulio Paci, Giacomo Sommavilla, Graham Taylor, and Florian Nebout. A multi-scale approach to gesture detection and recognition. In *ICCV Workshop on Understanding Human Activities: Context and Interactions*, 2013.
33. Mayank Rana, Graham Taylor, Ian Spiro, and Chris Bregler. 3d skeletal reconstruction from low-resolution multi-view images. In *IEEE CVPR Workshop on Human Activity Understanding from 3D Data (HAU3D)*, 2012.
34. Atul Kanaujia, Niels Haering, Graham Taylor, and Chris Bregler. 3d human pose and shape estimation from multi-view imagery. In *IEEE CVPR Workshop on Human Activity Understanding from 3D Data (HAU3D)*, 2011.
35. Ian Spiro, Graham Taylor, George Williams, and Christoph Bregler. Hands by hand: Crowd-sourced motion tracking for gesture annotation. In *IEEE CVPR Workshop on Advancing Computer Vision with Humans in the Loop (ACVHL)*, 2010.

Refereed extended abstracts

36. Jiwoong Im and Graham Taylor. Improving semi-supervised neural networks for scene understanding by learning the neighbourhood graph. In *IEEE CVPR Workshop on Scene Understanding*, 2014.
37. Graham Taylor, Ian Spiro, Chris Bregler, and Rob Fergus. Metric learning by active crowd-sourcing. In *The Learning Workshop (Snowbird)*, 2011.
38. Graham Taylor, Rob Fergus, Yann LeCun, and Christoph Bregler. Convolutional learning of spatio-temporal features. In *New York Academy of Sciences Machine Learning Symposium*, 2010.
39. Graham Taylor and Christoph Bregler. Learning local spatio-temporal features for activity recognition. In *The Learning Workshop (Snowbird)*, 2010.
40. Matthew Zeiler, Dilip Krishnan, Graham Taylor, and Rob Fergus. Deconvolutional networks for feature learning. In *The Learning Workshop (Snowbird)*, 2010.
41. George Williams, Graham Taylor, Kirill Smolskiy, and Christoph Bregler. Identifying people based on their motion signature. In *The Learning Workshop (Snowbird)*, 2010.
42. Graham Taylor, Geoffrey Hinton, and Sam Roweis. Deep generative models for modeling animate motion. In *Proc. 4th International Symposium on Adaptive Motion of Animals and Machines (AMAM)*, 2008.

Unpublished works

43. Graham Taylor. *Composable, distributed-state models for high-dimensional time series*. PhD thesis, University of Toronto, 2009.
44. Graham Taylor. Reinforcement learning for parameter control of image-based applications. Master's thesis, University of Waterloo, 2004.

GRANTS

- | | |
|-----------|--|
| 2015 | NSERC Engage: <i>Optimization Techniques for the Design of the Archimedes Screw Hydro Generator Model</i> (\$24,868) |
| 2014-2015 | Google ATAP: <i>Continuous Authentication Program</i> (\$111,750 USD) |
| 2014-2015 | NSERC Engage Plus: <i>Machine Learning for the remote monitoring of insect pests in agriculture</i> |

(\$25,000)

2014-2015 NSERC Engage Plus: *Deep Scene Parsing from Hyperspectral Imagery* (\$22,263)

2013-2018 CFI Leader's Opportunity Fund: *Massively parallel hardware accelerators for large-scale machine learning* (\$299,929)

2013-2014 NSERC Engage: *Learning to detect insects in the field: towards a fully automated intelligent system for remote pest monitoring* (\$25,000)

2013-2014 NSERC Engage: *Developing a deep scene parser for UAV-acquired images* (\$23,500)

2013-2018 NSERC Discovery Grant: *Deep Learning and Representation Learning for Sequential Data* (\$25,000/yr for 5 years)

2013 NSERC Research Tools and Instruments (\$98,258)

2013-2014 MITACS Globalink (3 @ ~\$10,000 ea.)

ADVISING

2014-present Carolyn Augusta, University of Guelph
Current PhD student (co-supervised with R. Deardon)

2014-present Griffin Lacey, University of Guelph
Current MSc student (co-supervised with S. Areibi)

2014-present Matthew Veres, University of Guelph
Current MSc student (co-supervised with M. Moussa)

2014 Ruoyan (Ryan) Wang, Dalian University of Technology
Summer research student through MITACS Globalink.

2014 Weiguang (Gavin) Ding, University of Guelph
Current research scientist.

2013-present Jiwoong (Daniel) Im, University of Guelph
Current MSc student.

2013-present Jan Rudy, University of Guelph
Current MSc student.

2013-present Ethan Buchman, University of Guelph
Current MSc student.

2013-present Ammar Abu Leil, University of Guelph
Current MSc student (co-supervised with M. Moussa).

2013-2014 Oana Burlacu, University of Guelph
MEng student: Classification of spastic patients based on clinical trial (co-supervised with H. Abdullah).

2013 Terrance Devries, University of Guelph
Summer undergraduate research assistant: Multi-Task Learning of Facial Landmarks and Expression.

2013 Kumar Biswaranjan, IIT Guwahati
Summer research student through MITACS Globalink.

2013 Samresh Satapathy, NIT Rourkela
Summer research student through MITACS Globalink (co-advised with Dave Lubitz).

2013 Mingyuan Jiu, INSA de Lyon
Visiting Researcher: Deep Learning for Activity Recognition

2012-2013 Shamir Alavi, University of Guelph

- Research scientist: Learning to detect insects in the field
- 2012 Malte Probst, University of Mainz
Visiting Researcher: Energy-based models for reinforcement learning
- 2012 Eric Noland, University of Waterloo
Undergraduate Co-op: Machine learning applied to foreign exchange
- 2008-2009 Matthew Zeiler, University of Toronto
BASc Thesis: Learning Pigeon Behaviour Using Binary Latent Variables (2009)
- 2009-2012 Matthew Zeiler, New York University (faculty advisor: Rob Fergus)
PhD Thesis: Hierarchical Convolutional Deep Learning in Computer Vision
- 2009-2012 Ian Spiro, New York University (faculty advisor: Chris Bregler)
PhD Thesis: Augmenting Information Flow for Visual Privacy

SELECTED MEDIA COVERAGE

- Dana Liebelson. Do Androids Dream of Electric Lolcats?. *Mother Jones Magazine*. September-October 2014.
- Marco Hochgemuth. Computer Vision meets Dutch Music Video. *Radio Netherlands Worldwide*, Jun 8, 2011.
- Clark Boyd. Dutch band C-Mon & Kypski's Crowdsourcing Video. *Public Radio International: The World*, Jun 1, 2011.
- Discovery Channel. Motion Capture. *Innovation Nation with Miles O'Brien*, Dec 29, 2010.
- Christoph Weiss. Crowd2cloud. *FM4 ORF Austrian radio*, Sep 14, 2010. (in German).
- Ars Electronica Festival (TV segment). *Servus TV*, Sep 11, 2010. (in German and English).
- Ars Electronica Festival (TV segment). *Repair TV*, Sep 5, 2010. (in German and English).
- Hugh Hart. Hallucinatory Art Snags Attention at Ars Electronica Festival. *Wired Blog*, Aug 31, 2010.

HONORS & AWARDS

- 2008-2009 Monica Ryckman Bursary
- 2005-2007 Walter C. Sumner Fellowship
- 2005-2006 Ontario Graduate Scholarship
- 2003-2005 NSERC Post Graduate Scholarship (PGS-A)

SERVICE TO THE PROFESSION

Program Committee

- 2012-2015 International Conference on Machine Learning (ICML)
- 2012-2014 Uncertainty in Artificial Intelligence (UAI)
- 2012-2014 Face and Gesture (FG)
- 2009-2014 Neural Information Processing Systems (NIPS)
- 2011-2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)
- 2013-2015 International Conference on Learning Representations (ICLR)
- 2011,2013 Artificial Intelligence and Statistics (AISTATS)
- 2011 International Conference on Computer Vision (ICCV)

Other Committees and Organization

- 2012 CVPR Tutorial on Deep Learning Methods for Vision
- 2011-2012 NIPS Workshop on Big Learning: Algorithms, Systems, and Tools for Learning at Scale
- 2011 NIPS Workshop on Learning Semantics
- 2011 ICML Workshop on Unsupervised and Transfer Learning
- 2011 CVPR Workshop on Gesture Recognition
- 2010-2011 NIPS Deep and Unsupervised Feature Learning Workshop
- 2010-2011 Unsupervised and Transfer Learning Challenge

Journal reviewing

International Journal of Computer Vision (IJCV)
IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)
Journal of Machine Learning Research (JMLR)
IEEE Signal Processing Letters
Machine Learning (MLJ)
Neural Networks

SERVICE TO THE COMMUNITY

- 2007-2012 **Sky's The Limit Youth Organization**
Member, Board of Directors
 - Placed over 1025 computers to date into the homes of economically disadvantaged youth
 - Worked directly with families: installation, support, and education
- 2005-2008 **Pathways to Education: Regent Park**
Volunteer and Program Facilitator
 - Tutored and mentored under-resourced high-school students
 - Trained and supervised volunteer tutors
- 2005-2007 **89 Chestnut Residence, University of Toronto**
Don (Residence Assistant)
 - Acted as a "first contact" to support first and upper-year university students
 - Promoted diversity, spirit, cooperation, an academic atmosphere and mutual respect