

Curriculum Vitae

Greg Mori

Professor

School of Computing Science

8888 University Drive

Burnaby, BC V5A 1S6, Canada

Email: mori@cs.sfu.ca

Web: <http://www.cs.sfu.ca/~mori>

Tel.: +1 (778) 782 7111

Fax: +1 (778) 782 3045

Contents

1	Background	2
1.1	Education	2
1.2	Employment History	2
1.3	Awards	2
2	Research	3
2.1	Awards	3
2.2	Research Program	4
2.3	Most Significant Research Contributions	4
2.4	Publications	6
2.5	Selected Invited Talks	17
2.6	Research Funding	20
3	Teaching	24
3.1	Award	24
3.2	Graduate Student Supervision	24
3.3	Postdoctoral Fellow Supervision	30
3.4	Undergraduate Student Supervision	30
3.5	Course Teaching at Simon Fraser University	31
3.6	Summary of Student Course Evaluations	33
3.7	Other Teaching	33
4	Service	35
4.1	Service To Simon Fraser University	35
4.2	Service to the Academic Community	35

1 Background

1.1 Education

- 2004 Ph.D. in Computer Science
Department of Electrical Engineering and Computer Sciences
University of California at Berkeley
Thesis: Detecting and Localizing Human Figures
Advisor: Jitendra Malik
- 1999 Hon. B.Sc. in Computer Science and Mathematics with High Distinction
University of Toronto

1.2 Employment History

- 06/2018 - current Research Director
RBC Borealis AI Vancouver
- 09/2015 - current Professor
School of Computing Science, Simon Fraser University
- 05/2015 - 04/2018 Director
School of Computing Science, Simon Fraser University
- 12/2014 - 05/2015 Visiting Scientist
Google Inc., Mountain View, CA
- 09/2010 - 08/2015 Associate Professor
School of Computing Science, Simon Fraser University
- 08/2004 - 08/2010 Assistant Professor
School of Computing Science, Simon Fraser University
- 08/1999 - 07/2004 Graduate Student Researcher
Department of Electrical Engineering and Computer Sciences, UC Berkeley
Advisor: Jitendra Malik
- 06/2000 - 08/2000 Student Intern, Computer Vision Group
Intel Corporation, Santa Clara, CA
Host: Gary Bradski
- 06/1999 - 08/1999 Student Intern, graphics research
Electronic Arts, Burnaby, BC
Host: John Buchanan
- 10/1997 - 08/1998 Student Intern, Media Integration and Communications Research Laboratories
Advanced Telecommunications Research (ATR), Kyoto, Japan
Host: Seiki Inoue
- 05/1997 - 08/1997 Student Intern, Consortium for Software Engineering Research (CSER)
University of Toronto / IBM Centre for Advanced Studies, Toronto, ON
Advisors: Kostas Kontogiannis and John Mylopoulos

1.3 Awards

- 2017 ICCV Helmholtz Prize
Type: Research
Organization: IEEE Technical Committee on Pattern Analysis and Machine Intelligence

- 2017 IAPR MVA Most Influential Paper over the Decade Award
Type: Research
Organization: IAPR International Conference on Machine Vision Applications
- 2016 Discovery Accelerator Supplement
Type: Research
Organization: Natural Sciences and Engineering Research Council of Canada (NSERC)
- 2014 Award for Service to the Research Community
Type: Service
Organization: Canadian Image Processing and Pattern Recognition Society (CIPPRS)
- 2011 Award for Excellence in Research (early career)
Type: Research
Organization: SFU Faculty of Applied Science
- 2010 Outstanding Reviewer
Type: Service
Organization: IEEE Conference on Computer Vision and Pattern Recognition (CVPR)
- 2009 Outstanding Reviewer
Type: Service
Organization: IEEE International Conference on Computer Vision (ICCV)
- 2008 Discovery Accelerator Supplement
Type: Research
Organization: Natural Sciences and Engineering Research Council of Canada (NSERC)
- 2008 Award for Research Excellence and Service
Type: Research / Service
Organization: Canadian Image Processing and Pattern Recognition Society (CIPPRS)
- 2006 Excellence in Undergraduate Teaching Award
Type: Teaching
Organization: SFU Undergraduate Computing Science Student Society (CSSS)
- 1999 Regents' Fellowship
Type: Fellowship
Organization: University of California at Berkeley
- 1995-1997 Regents' Scholarship
Type: Fellowship
Organization: Victoria College (University of Toronto)

2 Research

2.1 Awards

ICCV Helmholtz Prize

The Helmholtz Prize is an award given biyearly by the TCPAMI at the International Conference on Computer Vision (ICCV) for fundamental contributions in Computer Vision. The award recognizes ICCV papers from ten years ago with significant impact on computer vision research. I received this award in 2017 for the paper: A.A. Efros, A.C. Berg, G. Mori, and J. Malik, *Recognizing Action at A Distance*.

IAPR MVA Most Influential Paper over the Decade Award

This award recognizes IAPR MVA papers from 10 years ago with significant influence on machine vision technologies over the decade. I received this award in 2017 for the paper: A. Rova, G. Mori, and L.M. Dill, *One Fish, Two Fish, Butterfish, Trumpeter: Recognizing Fish in Underwater Video*.

NSERC Discovery Accelerator Supplement 2008-2011, 2016-2019

The NSERC Discovery Accelerator Supplement (DAS) provides substantial and timely resources to a small group of outstanding researchers who have a well-established research program, and who show strong potential to become international leaders in their respective area of research¹. The DAS provides a research grant of \$120,000 over 3 years. NSERC recognizes 125 Canadian researchers, across all fields of science and engineering, with DAS awards annually. I was a recipient twice, in 2008 and again in 2016.

SFU FAS Award for Excellence in Research (early career) 2011

The Faculty of Applied Sciences (FAS) at SFU has established five awards to recognize faculty members who have demonstrated superior performance in teaching, research or service, and a staff member who has made excellent contributions to service. I received the early career research award in 2011.

CIPPRS Award for Research Excellence and Service 2008

The Canadian Image Processing and Pattern Recognition Society (CIPPRS) annually recognizes a researcher for research excellence and contributions to the research community. I received this award in 2008 for my research on human pose estimation and activity recognition, and service to the Canadian computer vision community.

2.2 Research Program

My research is in computer vision, and is concerned with developing algorithms that automatically interpret images and videos, particularly those containing people. I have made significant contributions towards solving the problems of human pose estimation and human action recognition. At a broad level, the methodology followed is to construct features and representations that capture our intuition regarding these vision problems. We operationalize these via machine learning algorithms, adapting them to suit our purposes.

Examples of this methodology are described in detail below. Specific examples of features and representations include work on *superpixels* for representing images, motion features for human action recognition, and our structured models for video sequences and group activities. We have developed variants of machine learning algorithms such as hidden Conditional Random Fields (hCRF), Latent Dirichlet Allocation (LDA), latent SVMs, and deep networks to implement these ideas.

2.3 Most Significant Research Contributions

Conference publications are extremely important in computer vision. Top conferences utilize rigorous double-blind peer-review processes and are very selective. Acceptance rates for CVPR, ECCV, ICCV, and NIPS, the major computer vision and machine learning conferences, are typically in the 20-30% range. I regularly publish in these conferences, and in T-PAMI, a top journal in computer vision. My work has been cited 8000 times, h-index 37, according to Google Scholar.

Publication numbers below refer to those in Sec. 2.4, which includes conference acceptance rate data.

¹http://www.nserc-crsng.gc.ca/Professors-Professeurs/Grants-Subs/DGAS-SGSA_eng.asp

Human Action Recognition

We have developed a body of work on human action recognition. This work has spanned many sub-topics, and includes models for spatio-temporal structure [J6, C17, C20], relationships between multiple people in a scene [J9, C37], semantic tag representations for internet video [C53], and recognition of actions using motion cues [J5, C14].

Incorporating action sequence information into human action recognition is a standard challenge. Previous research takes one of two approaches. Either temporal sequence information is ignored, with video frames classified independently, or sequence models (e.g. HMMs, DBNs) are constructed. We developed a novel *bag-of-words* sequence model [J5] that falls in between – capturing temporal information via co-occurrence statistics among actions, without the complexity of a full temporal model. This method proved effective on a variety of standard benchmark datasets.

We have also developed part-based representations for human actions, learning motion part features [C14] and models combining global-scale and local part features [J6, C17, C20]. The mid-level motion feature learning method [C14] obtained high accuracy on benchmark KTH and Weizmann human action datasets, and is very efficient. It formed the basis for recent work on real-time gesture recognition used in robotics applications [C18, C23, C50]. Our work on global-scale and local part features [J6, C17, C20] has developed a principled method for combining these two sources of information, and shown the effectiveness of using the *max-margin* learning criterion for finding the parameters of this type of model.

Our work on structured models has explored the relationship between the actions of individuals, their interactions, and labelings of actions at multiple levels of detail. We have examined the use of person-person context and automatically inferring connections between individuals in a scene [J9, C30]. Further work built representations combining low-level actions, high-level events, and social roles of people [C37].

Human Pose Estimation

I have developed approaches for human body pose estimation. We did pioneering work in the use of *exemplar* methods for localizing human figures in still images. These methods are based on matching input images to a set of stored example 2D images of human figures with labelled joint locations. The work [J2, C2] is well cited (517 citations as of Apr. 2014 according to Google Scholar).

We have also developed methods for combining segmentation and recognition in the context of human pose estimation. Our work used a segmentation-as-preprocessing paradigm in which an input image is first over-segmented into small regions called superpixels [C6, C5] and is also widely cited (588 citations as of Apr. 2014 according to Google Scholar). This strategy is useful for general object recognition problems, and the problem of not only recognizing, but also segmenting objects has received a large amount of attention since our work in 2004. Numerous other papers have made use of the superpixel approach for object recognition.

Breaking Visual CAPTCHAs

A CAPTCHA is a program that can generate and grade tests that humans can pass but current computer programs cannot pass. CAPTCHA stands for “Completely Automated Public Turing test to Tell Computers and Humans Apart”. CAPTCHAs are employed by internet companies such as Yahoo and TicketMaster to prevent *bots* from signing up for free email accounts or purchasing tickets. The most commonly deployed CAPTCHAs are the word-puzzle type, where a distorted word placed in a cluttered background is presented to a user who is asked to read the word.

We developed methods for breaking these word-based CAPTCHAs ([C3], [J1]) which were based on our general-purpose object recognition algorithms. This domain provided evidence of the effectiveness of our computer vision algorithms. I have also done consulting work with companies to develop CAPTCHAs which are resistant to automated attacks. In addition, this work has received a large amount of attention from the popular press. This work was featured in the New York Times Science Section. The webpage describing our approach has received over 600,000 visitors.

Applications of Vision-based Activity Monitoring

In addition to fundamental algorithm development, I regularly collaborate with researchers in various application areas. We bring to bear state-of-the-art vision algorithms on data collection problems and generate novel algorithms and models for these problems.

For scientists and engineers, gathering data can be a labour-intensive and expensive process. I have collaborated with researchers in civil engineering studying the design of road systems (T. Sayed, UBC), kinesiologists attempting to understand the causes and circumstances of falls by nursing home residents (S. Robinovitch, SFU), and biologists examining the behaviours of animals in the wild (L. Dill, K. Rothley, SFU).

In each of these instances, the collaborators desired real-world data on their subjects, and collected video footage. However, cameras generate large amounts of data, which are typically sorted manually to collect the required information. As a computer vision researcher, there is a great opportunity to aid scientists by automating parts of the video analysis process.

We have applied our algorithms which we developed for monitoring the activities of, and detecting the presence of humans to the problem of animal activity monitoring. In collaborations with our biologist partners, Prof. Larry Dill (FRSC) and Prof. Kristina Rothley, we have developed systems for counting fish viewed from underwater cameras [C11], analyzing the activities of grasshoppers in cages [C10], and detecting grizzly bears at a remote site in the Yukon [J4]. We developed novel contextual models for group activity recognition that can be used to find nursing home falls in an off-line setting for data collection [J9]. Novel tracking methods [J8] were deployed to detect and track road users to collect data about pedestrians and vehicle-pedestrian interactions in intersections [J15, J11, J10].

2.4 Publications

Legend:

Names in **bold** face are my students.

The acceptance rate is mentioned where available (mainly for the top-tier conferences).

Refereed Journal Papers

- [J21] **N. Nauata, H. Hu, G.-T. Zhou, Z. Deng, Z. Liao**, and G. Mori. Structured Label Inference for Visual Understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, T-PAMI, 2018.
- [J20] **F. Tung** and G. Mori. Deep Neural Network Compression by In-Parallel Pruning-Quantization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, T-PAMI, 2018.
- [J19] S. Yeung, O. Russakovsky, N. Jin, M. Andriluka, G. Mori and L. Fei-Fei. Every Moment Counts: Dense Detailed Labeling of Actions in Complex Videos. *International Journal of Computer Vision*, IJCV 2017.
- [J18] **H. Hajimirsadeghi** and G. Mori. Multi-Instance Classification by Max-Margin Training of Cardinality-Based Markov Networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, T-PAMI 39(9) pp.1839-1852, 2017.
- [J17] **Y. Sefidgar, A. Vahdat**, S. Se, and G. Mori. Discriminative Key-Component Models for Interaction Detection and Recognition. *Computer Vision and Image Understanding*, CVIU 135 pp.16-30, 2015.
- [J16] **J. Li, H. Hajimirsadeghi**, M. Zaki, G. Mori, and T. Sayed. Computer Vision Techniques to Collect Helmet-Wearing Data on Cyclists. *Transportation Research Record: Journal of the Transportation Research Board*, 2468, pp.1-10, 2014.
- [J15] H. Hediye, T. Sayed, M. Zaki, and G. Mori. Pedestrian Gait Analysis Using Automated Computer Vision Techniques. *Transportmetrica A: Transport Science*, 10(3), pp.214-232, 2014.
- [J14] O. Aziz, E. Park, G. Mori, and S. Robinovitch. Distinguishing the Causes of Falls in Humans Using an Array of Wearable Tri-Axial Accelerometers. *Gait & Posture*, 39(1), pp.506-512, 2014.

- [J13] S. Oh, S. McCloskey, I. Kim, **A. Vahdat**, **K. Cannons**, **H. Hajimirsadeghi**, G. Mori, A. G. Perera, M. Pandey, J. J. Corso. Multimedia Event Detection and Recounting with Multimodal Feature Fusion and Temporal Concept Localization. *Machine Vision and Applications*, MVA 25(1) pp.49-69, 2014.
- [J12] **M. Ranjbar**, **T. Lan**, **Y. Wang**, S. Robinovitch, Z. Li, and G. Mori. Optimizing Non-Decomposable Loss Functions in Structured Prediction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, T-PAMI 35(4) pp.911-924, 2013.
- [J11] M. Zaki, T. Sayed, and G. Mori. Classifying Road-Users in Urban Scenes Using Movement Patterns. *ASCE Journal of Computing in Civil Engineering*, 27(4) pp.395-406, 2013.
- [J10] S. Li, T. Sayed, M. Zaki, G. Mori, **F. Stefanus**, **B. Khanloo**, N. Saunier. Automating Collection of Pedestrian Data Through Computer Vision Techniques. *Transportation Research Record: Journal of the Transportation Research Board*, 2299, pp.121-127, 2012.
- [J9] **T. Lan**, **Y. Wang**, **W. Yang**, S. Robinovitch and G. Mori. Discriminative Latent Models for Recognizing Contextual Group Activities. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, T-PAMI 34(8) pp.1549-1562, 2012.
- [J8] **B. Y. S. Khanloo**, **F. Stefanus**, **M. Ranjbar**, Z.-N. Li, N. Saunier, T. Sayed, and G. Mori. A Large Margin Framework for Single Camera Offline Tracking with Hybrid Cues. *Computer Vision and Image Understanding*, CVIU 116 pp.676-689, 2012.
- [J7] **P. Wighton**, T. Lee, G. Mori, H. Lui, D. I. McLean and M. S. Atkins. Conditional Random Fields and Supervised Learning in Automated Skin Lesion Diagnosis. *International Journal of Biomedical Imaging, Special Issue on Machine Learning in Medical Imaging*, 2011.
- [J6] **Y. Wang** and G. Mori. Hidden Part Models for Human Action Recognition: Probabilistic vs. Max-Margin. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, T-PAMI 33(7) pp.1310-1323, 2011.
- [J5] **Y. Wang** and G. Mori. Human Action Recognition by Semi-Latent Topic Models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, T-PAMI 31(10) pp.1762-1774, 2009.
- [J4] **J. Wawerla**, **S. Marshall**, G. Mori, K. Rothley, and **P. Sabzmeydani**. BearCam: Automated Wildlife Monitoring At The Arctic Circle. *Machine Vision and Applications*, MVA 20(5) pp.303-317, June 2009.
- [J3] **R. Botchen**, **S. Bachthaler**, **F. Schick**, M. Chen, G. Mori, D. Weiskopf, and T. Ertl. Action-based Multi-field Video Visualization. *IEEE Transactions on Visualization & Computer Graphics*, T-VCG 14(4) pp.885-899, July/August 2008.
- [J2] G. Mori and J. Malik. Recovering 3d Human Body Configurations Using Shape Contexts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, T-PAMI 28(7) pp.1052-1062, July 2006.
- [J1] G. Mori, S. Belongie, and J. Malik. Efficient Shape Matching Using Shape Contexts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, T-PAMI 27(11) pp.1832-1837, Nov. 2005.

Refereed Book Chapters

- [B2] **W. Yang**, **Y. Wang**, and G. Mori. Learning Transferable Distance Functions for Human Action Recognition *Machine Learning for Vision-based Motion Analysis*, Springer, 2010.
- [B1] S. Belongie, G. Mori, and J. Malik. Matching with Shape Contexts. *Analysis and Statistics of Shapes*, eds. T. Yezzi and H. Krim, Birkhäuser, 2006.

Refereed Conference Papers

- [C102] **X. Chang**, F. Tung, G. Mori. Learning Discriminative Prototypes with Dynamic Time Warping. *IEEE/CVF Computer Vision and Pattern Recognition, CVPR 2021*.
- [C101] **M. Zhai, L. Chen**, G. Mori. Hyper-LifelongGAN: Scalable Lifelong Learning for Image Conditioned Generation. *IEEE/CVF Computer Vision and Pattern Recognition, CVPR 2021*.
- [C100] **S. Hu**, Z. Yang, G. Mori. Neural fidelity warping for efficient robot morphology design. *IEEE International Conference on Robotics and Automation, ICRA 2021*. Xian, China, May 2021.
- [C99] **Y. Gong**, H. Hajimirsadeghi, J. He, T. Durand, G. Mori. Variational Selective Autoencoder: Learning from Partially-Observed Heterogeneous Data. *24th International Conference on Artificial Intelligence and Statistics, AISTATS 2021*. Virtual, April 2021.
- [C98] **R. Deng**, B. Chang, M. A. Brubaker, G. Mori, A. Lehrmann. Modeling Continuous Stochastic Processes with Dynamic Normalizing Flows. *Neural Information Processing Systems, NIPS 2020*. Vancouver, Canada, December 2020.
- [C97] **M. Zhai, L. Chen**, J. He, **M. Nawhal**, F. Tung, and G. Mori. Piggyback GAN: Efficient Lifelong Learning for Image Conditioned Generation. *European Conference on Computer Vision, ECCV 2020*. Glasgow, Scotland, August 2020.
- [C96] **M. Nawhal, M. Zhai**, A. Lehrmann, L. Sigal, and G. Mori. Generating Videos of Zero-Shot Compositions of Actions and Objects. *European Conference on Computer Vision, ECCV 2020*. Glasgow, Scotland, August 2020.
- [C95] N. Nauata, K.-H. Chang, C.-Y. Cheng, G. Mori, and Y. Furukawa. House-GAN: Relational Generative Adversarial Networks for Graph-constrained House Layout Generation. *European Conference on Computer Vision, ECCV 2020*. Glasgow, Scotland, August 2020.
- [C94] **C. Chen, S. Hu**, P. Nikdel, G. Mori, and M. Savva. Relational Graph Learning for Crowd Navigation. *IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2020*. Las Vegas, NV, October 2020.
- [C93] **L. Chen**, J. Chen, H. Hajimirsadeghi, and G. Mori. Adapting Grad-CAM for Embedding Networks. *IEEE Winter Conference on Applications of Computer Vision, WACV 2020*. Snowmass Village, CO, March 2020.
- [C92] **F. Tung** and G. Mori. Similarity-Preserving Knowledge Distillation. *IEEE International Conference on Computer Vision, ICCV 2019*. Seoul, Korea, October 2019.
- [C91] **M. Zhai, L. Chen, F. Tung, J. He, M. Nawhal**, and G. Mori. Lifelong GAN: Continual Learning for Conditional Image Generation. *IEEE International Conference on Computer Vision, ICCV 2019*. Seoul, Korea, October 2019.
- [C90] **A. A. Jyothi, T. Durand, J. He**, L. Sigal, and G. Mori. LayoutVAE: Stochastic Scene Layout Generation from a Label Set. *IEEE International Conference on Computer Vision, ICCV 2019*. Seoul, Korea, October 2019.
- [C89] **T. Durand**, N. Mehrasa and G. Mori. Learning a Deep ConvNet for Multi-label Classification with Partial Labels. *IEEE Computer Vision and Pattern Recognition, CVPR 2019*. Long Beach, CA, June 2019.
- [C88] **N. Mehrasa, A. Abdu Jyothi, T. Durand, J. He**, L. Sigal, and G. Mori. A Variational Auto-Encoder Model for Stochastic Point Processes. *IEEE Computer Vision and Pattern Recognition, CVPR 2019*. Long Beach, CA, June 2019.

- [C87] **J. He, Y. Gong**, J. Marino, G. Mori, and A. Lehrmann. Variational Autoencoders with Jointly Optimized Latent Dependency Structures. *International Conference on Learning Representations, ICLR 2019*. New Orleans, LA, May 2019.
- [C86] **Z. Deng, J. Chen, Y. Fu**, and G. Mori. Probabilistic Neural Programmed Networks for Scene Generation. *Neural Information Processing Systems, NIPS 2018*. Montreal, Canada, December 2018.
- [C85] **M. Ibrahim** and G. Mori. Hierarchical Relational Networks for Group Activity Recognition and Retrieval. *European Conference on Computer Vision, ECCV 2018*. Munich, Germany, September 2018.
- [C84] L. Zhu, **R. Deng**, M. Maire, **Z. Deng**, G. Mori, and P. Tan. Sparsely Aggregated Convolutional Networks. *European Conference on Computer Vision, ECCV 2018*. Munich, Germany, September 2018.
- [C83] **C. Chen, F. Tung**, N. Vedula, and G. Mori. Constraint-Aware Deep Neural Network Compression. *European Conference on Computer Vision, ECCV 2018*. Munich, Germany, September 2018.
- [C82] **J. He**, A. Lehrmann, J. Marino, G. Mori, and L. Sigal. Probabilistic Video Generation using Holistic Attribute Control. *European Conference on Computer Vision, ECCV 2018*. Munich, Germany, September 2018.
- [C81] **F. Baradel**, N. Neverova, C. Wolf, J. Mille, and G. Mori. Object Level Visual Reasoning in Videos. *European Conference on Computer Vision, ECCV 2018*. Munich, Germany, September 2018.
- [C80] **M. Zhai, R. Deng, J. Chen, L. Chen, Z. Deng**, and G. Mori. Adaptive Appearance Rendering. *British Machine Vision Conference, BMVC 2018*. Newcastle, UK September 2018. *accept rate: $\frac{258}{862} = 29.9\%$* .
- [C79] **F. Tung** and G. Mori. CLIP-Q: Deep Network Compression Learning by In-Parallel Pruning-Quantization. *IEEE Computer Vision and Pattern Recognition, CVPR 2018*. Salt Lake City, UT, June 2018.
- [C78] **J. He, Z. Deng, M. Ibrahim**, and G. Mori. Generic Tubelet Proposals for Action Localization. *IEEE Winter Conference on Applications of Computer Vision, WACV 2018*. Lake Tahoe, NV, March 2018.
- [C77] L. Shen, S. Yeung, J. Hoffman, G. Mori, and L. Fei-Fei. Scaling Human-Object Interaction Recognition through Zero-Shot Learning. *IEEE Winter Conference on Applications of Computer Vision, WACV 2018*. Lake Tahoe, NV, March 2018.
- [C76] **N. Mehrasa, Y. Zhong, F. Tung**, L. Bornn, and G. Mori. Deep Learning of Player Trajectory Representations for Team Activity Analysis. *Sloan Sports Analytics Conference, 2018*. Boston, MA, March, 2018.
- [C75] **F. Tung, S. Muralidharan**, and G. Mori. Fine-Pruning: Joint Fine-Tuning and Compression of a Convolutional Network with Bayesian Optimization. *British Machine Vision Conference, BMVC 2017*. London, UK, September 2017.
- [C74] **Z. Deng**, R. Navarathna, P. Carr, S. Mandt, Y. Yue, I. Matthews and G. Mori. Factorized Variational Autoencoders for Modeling Audience Reactions to Movies. *IEEE Computer Vision and Pattern Recognition, CVPR 2017*. Honolulu, HI, July 2017. *accept rate: $\frac{783}{2680} = 29.2\%$* .
- [C73] S. Yeung, V. Ramanathan, O. Russakovsky, L. Shen, G. Mori and L. Fei-Fei. Learning to Learn from Noisy Web Videos. *IEEE Computer Vision and Pattern Recognition, CVPR 2017*. Honolulu, HI, July 2017. *accept rate: $\frac{783}{2680} = 29.2\%$* .
- [C72] **K. Rashedi Nia** and G. Mori. Building Damage Assessment Using Deep Learning and Ground-Level Image Data. *Fourteenth Canadian Conference on Computer and Robot Vision, CRV 2017*. Edmonton, AB, May 2017.
- [C71] **M. Khodabandeh, S. Muralidharan**, A. Vahdat, **N. Mehrasa**, E. M. Pereira, S. Satoh, and G. Mori. Un-supervised Learning of Supervoxel Embeddings for Video Segmentation. *IAPR International Conference on Pattern Recognition, ICPR 2016*. Cancun, Mexico December 2016.

- [C70] **M. Ibrahim, S. Muralidharan, Z. Deng**, A. Vahdat, and G. Mori. A Hierarchical Deep Temporal Model for Group Activity Recognition. *IEEE Computer Vision and Pattern Recognition*, CVPR 2016. Las Vegas, NV, June 2016. *accept rate*: $\frac{643}{2145} = 29.9\%$.
- [C69] **H. Hu, G.-T. Zhou, Z. Deng**, Z. Liao, and G. Mori. Learning Structured Inference Neural Networks with Label Relations. *IEEE Computer Vision and Pattern Recognition*, CVPR 2016. Las Vegas, NV, June 2016. *accept rate*: $\frac{643}{2145} = 29.9\%$.
- [C68] **Z. Deng**, A. Vahdat, **H. Hu**, G. Mori. Structure Inference Machines: Recurrent Neural Networks for Analyzing Relations in Group Activity Recognition. *IEEE Computer Vision and Pattern Recognition*, CVPR 2016. Las Vegas, NV, June 2016. *accept rate*: $\frac{643}{2145} = 29.9\%$.
- [C67] S. Yeung, O. Russakovsky, G. Mori, L. Fei-Fei. End-to-end Learning of Action Detection from Frame Glimpses in Videos. *IEEE Computer Vision and Pattern Recognition*, CVPR 2016. Las Vegas, NV, June 2016. *accept rate*: $\frac{643}{2145} = 29.9\%$.
- [C66] V. Ramanathan, K. Tang, G. Mori, and L. Fei-Fei. Learning Temporal Embeddings for Complex Video Analysis. *IEEE International Conference on Computer Vision*, ICCV 2015. Santiago, Chile, December 2015. *accept rate*: $\frac{525}{1698} = 30.3\%$.
- [C65] **H. Hajimirsadeghi** and G. Mori. Learning Ensembles of Potential Functions for Structured Prediction with Latent Variables. *IEEE International Conference on Computer Vision*, ICCV 2015. Santiago, Chile, December 2015. *accept rate*: $\frac{525}{1698} = 30.3\%$.
- [C64] **Z. Deng, M. Zhai, L. Chen, Y. Liu, S. Muralidharan**, M. Roshtkhari, and G. Mori. Deep Structured Models For Group Activity Recognition. In *26th British Machine Vision Conference*, BMVC 2015. Swansea, UK, September 2015. *accept rate*: $\frac{185}{553} = 33\%$.
- [C63] **W. Yan, J. Yap**, and G. Mori. Multi-Task Transfer Methods to Improve One-Shot Learning for Multimedia Event Detection. In *26th British Machine Vision Conference*, BMVC 2015. Swansea, UK, September 2015. *accept rate*: $\frac{185}{553} = 33\%$.
- [C62] **N. Shapovalova** and G. Mori. Clustered Exemplar-SVM: Discovering Sub-Categories for Visual Recognition. In *IEEE International Conference on Image Processing*, ICIP 2015. Quebec City, PQ, September 2015.
- [C61] **H. Hajimirsadeghi**, W. Yan, **A. Vahdat**, and G. Mori. Visual Recognition by Counting Instances: A Multi-Instance Cardinality Potential Kernel. *IEEE Computer Vision and Pattern Recognition*, CVPR 2015. Boston, MA, June 2015. *accept rate*: $\frac{602}{2123} = 28.4\%$.
- [C60] **M. Zhai, L. Chen, M. Khodabandeh, J. Li**, and G. Mori. Object Detection in Surveillance Video from Dense Trajectories. In *IAPR Conference on Machine Vision Applications*, MVA 2015. Tokyo, Japan, May 2015.
- [C59] **J. Li, Y. Liu**, A. Tageldin, M. Zaki, G. Mori, and T. Sayed. Automated Region-Based Vehicle Conflict Detection Using Computer Vision Techniques. *Transportation Research Board 94th Annual Meeting*, 2015
- [C58] **A. Vahdat, G. Zhou**, and G. Mori. Discovering Video Clusters from Visual Features and Noisy Tags. *European Conference on Computer Vision*, ECCV 2014. Zurich, Switzerland, September 2014.
- [C57] **S. Pourmehr, V. Monajjemi, S. Sadat, F. Zhan**, J. Wawerla, G. Mori, and R. Vaughan. “You are green”: a touch-to-name interaction in an integrated multi-modal multi-robot HRI system. *9th ACM/IEEE International Conference on Human-Robot Interaction*, HRI 2014. Bielefeld, Germany, March 2014.

- [C56] **V. Monajjemi, S. Pourmehr, S. Sadat, F. Zhan, J. Wawerla, G. Mori, and R. Vaughan.** Integrating multi-modal interfaces to command UAVs. *9th ACM/IEEE International Conference on Human-Robot Interaction (Video Session)*, HRI 2014. Bielefeld, Germany, March 2014.
- [C55] **G. Zhou, T. Lan, A. Vahdat, and G. Mori.** Latent Maximum Margin Clustering. *Neural Information Processing Systems*, NIPS 2013. Lake Tahoe, NV, December 2013. *accept rate*: $\frac{360}{1420} = 25.4\%$.
- [C54] **N. Shapovalova, M. Raptis, L. Sigal, and G. Mori.** Action is in the Eye of the Beholder: Eye-gaze Driven Model for Spatio-Temporal Action Localization. *Neural Information Processing Systems*, NIPS 2013. Lake Tahoe, NV, December 2013. *accept rate*: $\frac{360}{1420} = 25.4\%$.
- [C53] **A. Vahdat and G. Mori.** Handling Uncertain Tags in Visual Recognition. *IEEE International Conference on Computer Vision*, ICCV 2013. Sydney, Australia, December 2013. *accept rate*: $\frac{455}{1505} = 30.2\%$.
- [C52] **A. Vahdat, K. Cannons, G. Mori, I. Kim, and S. Oh.** Compositional Models for Video Event Detection: A Multiple Kernel Learning Latent Variable Approach. *IEEE International Conference on Computer Vision*, ICCV 2013. Sydney, Australia, December 2013. *accept rate*: $\frac{455}{1505} = 30.2\%$.
- [C51] **T. Lan, M. Raptis, L. Sigal, and G. Mori.** From Subcategories to Visual Composites: A Multi-level Framework for Object Detection. *IEEE International Conference on Computer Vision*, ICCV 2013. Sydney, Australia, December 2013. *accept rate*: $\frac{455}{1505} = 30.2\%$.
- [C50] **V. Monajjemi, J. Wawerla, R. Vaughan, and G. Mori.** HRI in the Sky: Creating and Commanding Teams of UAVs with a Vision-mediated Gestural Interface. *IEEE/RSJ International Conference on Intelligent Robots and Systems*, IROS 2013. Tokyo, Japan, November 2013. *accept rate*: $\frac{903}{2089} = 43.2\%$.
- [C49] **S. Pourmehr, V. Monajjemi, R. Vaughan, and G. Mori.** “You Two! Take Off!”: Creating, Modifying and Commanding Groups of Robots Using Face Engagement and Indirect Speech in Voice Commands. *IEEE/RSJ International Conference on Intelligent Robots and Systems*, IROS 2013. Tokyo, Japan, November 2013. *accept rate*: $\frac{903}{2089} = 43.2\%$.
- [C48] **I. Kim, S. Oh, A. Vahdat, K. Cannons, A. G. Perera, G. Mori.** Segmental Multi-way Local Pooling for Video Recognition. *ACM Multimedia Conference*, ACM MM 2013. Barcelona, Spain, October 2013. (short paper)
- [C47] **H. Hajimirsadeghi, J. Li, G. Mori, M. Zaki, and T. Sayed.** Multiple Instance Learning by Discriminative Training of Markov Networks. *Conference on Uncertainty in Artificial Intelligence*, UAI 2013. Bellevue, WA, July 2013. *accept rate*: $\frac{73}{233} = 31.3\%$.
- [C46] **T. Lan and G. Mori.** A Max-Margin Riffled Independence Model for Image Tag Ranking. *IEEE Computer Vision and Pattern Recognition*, CVPR 2013. Portland, OR, June 2013. *accept rate*: $\frac{472}{1870} = 25.2\%$.
- [C45] **G. Zhou, T. Lan, W. Yang, and G. Mori.** Object Matching Based Distance Function Learning for Image Classification. *IEEE Computer Vision and Pattern Recognition*, CVPR 2013. Portland, OR, June 2013. *accept rate*: $\frac{472}{1870} = 25.2\%$.
- [C44] **Y. Zhu, T. Lan, Y. Yang, S. Robinovitch, and G. Mori.** Latent Spatio-temporal Models for Action Localization and Recognition in Nursing Home Surveillance Video. In *IAPR Conference on Machine Vision Applications*, MVA 2013. Kyoto, Japan, May 2013.
- [C43] **S. Pourmehr, V. Monajjemi, J. Wawerla, R. Vaughan, and G. Mori.** A Robust Integrated System for Selecting and Commanding Multiple Mobile Robots. *IEEE International Conference on Robotics and Automation*, ICRA 2013. Karlsruhe, Germany, May 2013.
- [C42] **W. Yang, Y. Wang, A. Vahdat, and G. Mori.** Kernel Latent SVM for Visual Recognition. *Neural Information Processing Systems*, NIPS 2012. Lake Tahoe, NV, December 2012. *accept rate*: $\frac{370}{1467} = 25.2\%$.

- [C41] **H. Hajimirsadeghi** and G. Mori. Multiple Instance Real Boosting with Aggregation Functions. *IAPR International Conference on Pattern Recognition*, ICPR 2012. Tsukuba, Japan, November 2012.
- [C40] **T. Lan, W. Yang, Y. Wang,** and G. Mori. Image Retrieval with Structured Object Queries Using Latent Ranking SVM. *European Conference on Computer Vision*, ECCV 2012. Florence, Italy, October 2012.
- [C39] **N. Shapovalova, A. Vahdat, K. Cannons, T. Lan,** and G. Mori. Similarity Constrained Latent Support Vector Machine: An Application to Weakly Supervised Action Classification. *European Conference on Computer Vision*, ECCV 2012. Florence, Italy, October 2012.
- [C38] **O. Aziz,** E. J. Park, G. Mori, S. Robinovitch. Distinguishing Near-Falls from Daily Activities with Wearable Accelerometers and Gyroscopes using Support Vector Machines. *34th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, EMBC 2012. San Diego, CA, Sept., 2012.
- [C37] **T. Lan,** L. Sigal, and G. Mori. Social Roles in Hierarchical Models for Human Activity Recognition. *IEEE Computer Vision and Pattern Recognition*, CVPR 2012. Providence, RI, June 2012. *accept rate:* $\frac{465}{1933} = 24\%$.
- [C36] **M. Ranjbar, A. Vahdat,** and G. Mori. Complex Loss Optimization via Dual Decomposition. *IEEE Computer Vision and Pattern Recognition*, CVPR 2012. Providence, RI, June 2012. *accept rate:* $\frac{465}{1933} = 24\%$.
- [C35] M. Zaki, T. Sayed, and G. Mori. Classifying Road-Users in Urban Scenes Using Movement Patterns, *Proceedings of the 91st Annual Meeting of the Transportation Research Board*, Washington, DC, January 2012.
- [C34] **T. Lan, Y. Wang,** and G. Mori. Discriminative Figure-Centric Models for Joint Action Localization and Recognition. *13th International Conference on Computer Vision*, ICCV 2011. Barcelona, Spain, November 2011. *accept rate:* $\frac{339}{1433} = 23.7\%$.
- [C33] **Z. F. Huang, W. Yang, Y. Wang,** and G. Mori. LatentBoost for Action Recognition. In *22nd British Machine Vision Conference*, BMVC 2011. Dundee, Scotland, August 2011. *accept rate:* $\frac{133}{418} = 31.8\%$.
- [C32] **Y. Wang** and G. Mori. Max-margin Latent Dirichlet Allocation for Image Classification and Annotation. In *22nd British Machine Vision Conference*, BMVC 2011. Dundee, Scotland, August 2011. *accept rate:* $\frac{133}{418} = 31.8\%$.
- [C31] **B. Milligan,** G. Mori, and R. Vaughan. Selecting and Commanding Groups of Robots in a Vision Based Multi-Robot System. In *6th ACM/IEEE International Conference on Human-Robot Interaction (Video Session)*, HRI 2011. Lausanne, Switzerland, March 2011. **Best Video Award Winner.**
- [C30] **T. Lan, Y. Wang, W. Yang,** and G. Mori. Beyond Actions: Discriminative Models for Contextual Group Activities. *Neural Information Processing Systems*, NIPS 2010. Vancouver, BC, Canada, December 2010. *accept rate:* $\frac{293}{1219} = 24.0\%$.
- [C29] **Y. Wang** and G. Mori. A Discriminative Latent Model of Image Region and Object Tag Correspondence. *Neural Information Processing Systems*, NIPS 2010. Vancouver, BC, Canada, December 2010. *accept rate:* $\frac{293}{1219} = 24.0\%$.
- [C28] **M. Ranjbar,** G. Mori, and **Y. Wang.** Optimizing Complex Loss Functions in Structured Prediction. *European Conference on Computer Vision*, ECCV 2010. Hersonissos, Greece, September 2010. *accept rate:* $\frac{322}{1174} = 27.4\%$.
- [C27] **Y. Wang** and G. Mori. A Discriminative Latent Model of Object Classes and Attributes. *European Conference on Computer Vision*, ECCV 2010. Hersonissos, Greece, September 2010. *accept rate:* $\frac{322}{1174} = 27.4\%$.
- [C26] **W. Yang, Y. Wang,** and G. Mori. Recognizing Human Actions from Still Images with Latent Poses. *IEEE Computer Vision and Pattern Recognition*, CVPR 2010. San Francisco, CA, June 2010. *accept rate:* $\frac{463}{1728} = 26.8\%$.

- [C25] **B. Y. S. Khanloo, F. Stefanus, M. Ranjbar**, Z.-N. Li, N. Saunier, T. Sayed, and G. Mori. Max-Margin Offline Pedestrian Tracking with Multiple Cues. *Seventh Canadian Conference on Computer and Robot Vision*, CRV 2010. Ottawa, ON, May 2010.
- [C24] **A. Couture-Beil**, R. Vaughan, and G. Mori. Selecting and Commanding Individual Robots in a Multi-Robot System. *Seventh Canadian Conference on Computer and Robot Vision*, CRV 2010. Ottawa, ON, May 2010.
- [C23] **A. Couture-Beil**, R. Vaughan, and G. Mori. Selecting and Commanding Individual Robots in a Vision-Based Multi-Robot System. In *5th ACM/IEEE International Conference on Human-Robot Interaction (Video Session)*, HRI 2010. Osaka, Japan, March 2010. *accept rate*: $\frac{12}{23} = 52.2\%$.
- [C22] **Y. Wang, G. Haffari, S. Wang**, and G. Mori. A Rate Distortion Approach for Semi-Supervised Conditional Random Fields. In *Neural Information Processing Systems*, NIPS 2009. Vancouver, BC, Canada, December 2009. *accept rate*: $\frac{263}{1105} = 23.8\%$.
- [C21] **W. Yang, Y. Wang**, and G. Mori. Efficient Human Action Detection using Transferable Distance Function. In *Ninth Asian Conference on Computer Vision*, ACCV 2009. Xi'an, China, Sept. 2009. *accept rate*: $\frac{175}{670} = 26.1\%$.
- [C20] **Y. Wang** and G. Mori. Max-Margin Hidden Conditional Random Fields for Human Action Recognition. In *IEEE Computer Vision and Pattern Recognition*, CVPR 2009. Miami, FL, June 2009. *accept rate*: $\frac{383}{1464} = 26.2\%$.
- [C19] **M. Norouzi, M. Ranjbar**, and G. Mori. Stacks of Convolutional Restricted Boltzmann Machines for Shift-Invariant Feature Learning. In *IEEE Computer Vision and Pattern Recognition*, CVPR 2009. Miami, FL, June 2009. *accept rate*: $\frac{383}{1464} = 26.2\%$.
- [C18] **M. Bayazit, A. Couture-Beil**, and G. Mori. Real-time Motion-based Gesture Recognition using the GPU. In *IAPR Conference on Machine Vision Applications*, MVA 2009. Yokohama, Japan, May 2009.
- [C17] **Y. Wang** and G. Mori. Learning a discriminative hidden part model for human action recognition. In *Neural Information Processing Systems*, NIPS 2008. Vancouver, BC, Canada, December 2008. *accept rate*: $\frac{250}{1022} = 24.5\%$.
- [C16] **Y. Wang** and G. Mori. Multiple Tree Models for Occlusion and Spatial Constraints in Human Pose Estimation. In *European Conference on Computer Vision*, ECCV 2008. Marseille, France, October 2008. *accept rate*: $\frac{243}{871} = 27.9\%$.
- [C15] **G. Haffari, Y. Wang, S. Wang, G. Mori**, and F. Jiao. Boosting with Incomplete Information. In *International Conference on Machine Learning*, ICML 2008. Helsinki, Finland, July 2008. *accept rate*: $\frac{155}{583} = 27\%$.
- [C14] **A. Fathi** and G. Mori. Action Recognition Using Mid-level Motion Features. In *IEEE Computer Vision and Pattern Recognition*, CVPR 2008. Anchorage, AK, June 2008. *accept rate*: $\frac{508}{1593} = 31.9\%$.
- [C13] **A. Fathi** and G. Mori. Human Pose Estimation using Motion Exemplars. In *IEEE International Conference on Computer Vision*, ICCV 2007. Rio de Janeiro, Brazil, October 2007. *accept rate*: $\frac{281}{1190} = 23.6\%$.
- [C12] **P. Sabzmejdani** and G. Mori. Detecting Pedestrians by Learning Shapelet Features. In *IEEE Computer Vision and Pattern Recognition*, CVPR 2007. Minneapolis, MN, June 2007. *accept rate*: $\frac{353}{1250} = 28.2\%$.
- [C11] **A. Rova, G. Mori**, and L. M. Dill. One Fish, Two Fish, Butterfish, Trumpeter: Recognizing Fish in Underwater Video. In *IAPR Conference on Machine Vision Applications*, MVA 2007. Tokyo, Japan, May 2007. *accept rate*: $\frac{137}{220} = 62.3\%$.
- [C10] **M. Moslemi Naeni, G. Dutton, K. Rothley**, and G. Mori. Action Recognition of Insects Using Spectral Clustering. In *IAPR Conference on Machine Vision Applications*, MVA 2007. Tokyo, Japan, May 2007. *accept rate*: $\frac{137}{220} = 62.3\%$.

- [C9] **H. Jiang, Y. Wang**, M. Drew, Z. Li, and G. Mori. Unsupervised Discovery of Action Classes. In *IEEE Conf. on Computer Vision and Pattern Recognition*, CVPR 2006. New York, NY, June 2006. *accept rate*: $\frac{318}{1131} = 28.1\%$.
- [C8] **X. Li**, G. Mori and H. Zhang. Expression-Invariant Face Recognition with Expression Classification. In *Third Canadian Conference on Computer and Robot Vision*, CRV 2006. Quebec City, PQ, June 2006. *accept rate*: $\frac{72}{113} = 62\%$.
- [C7] **O. van Kaick** and G. Mori. Automatic Classification of Outdoor Images by Region Matching. In *Third Canadian Conference on Computer and Robot Vision*, CRV 2006. Quebec City, PQ, June 2006. *accept rate*: $\frac{72}{113} = 62\%$.
- [C6] G. Mori. Guiding Model Search Using Segmentation. In *IEEE International Conference on Computer Vision*, ICCV 2005. Beijing, China, October 2005. *accept rate*: $\frac{244}{1230} = 19.8\%$.
- [C5] G. Mori, X. Ren, A.A. Efros, and J. Malik. Recovering Human Body Configurations: Combining Segmentation and Recognition. In *IEEE Conf. on Computer Vision and Pattern Recognition*, CVPR 2004. Washington, D.C., June 2004. *accept rate*: $\frac{260}{873} = 29.8\%$.
- [C4] A.A. Efros, A.C. Berg, G. Mori, and J. Malik. Recognizing Action at A Distance. In *IEEE International Conference on Computer Vision*, ICCV 2003. Nice, France, October 2003. *accept rate*: $\frac{199}{966} = 20.6\%$.
- [C3] G. Mori and J. Malik. Recognizing Objects in Adversarial Clutter: Breaking a Visual CAPTCHA. In *IEEE Conf. on Computer Vision and Pattern Recognition*, CVPR 2003. Madison, WI, June 2003. *accept rate*: $\frac{209}{905} = 23.1\%$.
- [C2] G. Mori and J. Malik. Estimating Human Body Configurations using Shape Context Matching. In *European Conference on Computer Vision*, ECCV 2002. Copenhagen, Denmark, May 2002. *accept rate*: $\frac{226}{600} = 37.7\%$.
- [C1] G. Mori, S. Belongie, and J. Malik. Shape Contexts Enable Efficient Retrieval of Similar Shapes. In *IEEE Conf. on Computer Vision and Pattern Recognition*, CVPR 2001. Kauai, HI, December 2001. *accept rate*: $\frac{273}{920} = 29.7\%$.

Refereed Workshop Papers

- [W21] **R. Deng**, M. Brubaker, G. Mori, and A. Lehrmann. Continuous Latent Process Flows. *Time Series Workshop (at ICML)*, Virtual, July 2021.
- [W20] **N. Mehrasa, R. Deng**, M. Ahmed, B. Chang, J. He, T. Durand, M. Brubaker, and G. Mori. Point Process Flows. *Workshop on Learning with Temporal Point Processes (at NeurIPS)*, Vancouver, December 2019.
- [W19] **Y. Gong**, H. Hajimirsadeghi, J. He, **M. Nawhal**, T. Durand, and G. Mori. Variational Selective Autoencoder. *Symposium on Advances in Approximate Bayesian Inference*, Vancouver, December 2019.
- [W18] **M. Carvalho**, T. Durand, J. He, **N. Mehrasa**, and G. Mori. Arbitrarily-conditioned Data Imputation. *Symposium on Advances in Approximate Bayesian Inference*, Vancouver, December 2019.
- [W17] **L. Chen, M. Zhai, J. He**, and G. Mori. Object Grounding via Iterative Context Reasoning. *Workshop on Multi-Discipline Approach for Learning Concepts - Zero-Shot, One-Shot, Few-Shot and Beyond (at ICCV)*, Seoul, October 2019.
- [W16] **J. He, Y. Gong**, J. Marino, G. Mori, and A. Lehrmann. Variational Latent Dependency Learning. *Bayesian Deep Learning Workshop (at NeurIPS)*, Montreal, December 2018.

- [W15] J. Butepage, **J. He**, C. Zhang, L. Sigal, G. Mori, and S. Mandt. Informed Priors for Deep Representation Learning. *Symposium on Advances in Approximate Bayesian Inference (at NeurIPS)*, Montreal, December 2018.
- [W14] **L. Chen**, **M. Zhai**, and G. Mori. Attending to Distinctive Moments: Weakly-supervised Attention Models for Action Localization in Video. *5th Workshop on Web-scale Vision and Social Media (at ICCV)*, Venice, Italy, October 2017.
- [W13] **N. Nauata**, **J. Smith**, and G. Mori. Hierarchical Label Inference for Video Classification. *CVPR Workshop on Youtube-8M*, Honolulu, Hawaii, July 2017.
- [W12] **M. Khodabandeh**, **A. Vahdat**, **G.-T. Zhou**, **H. Hajimirsadeghi**, M. Roshtkhari, G. Mori, and S. Se. Discovering Human Interactions in Videos with Limited Data Labeling. *Workshop on Group and Crowd Behavior Analysis and Understanding (at CVPR)*, Boston, MA, June 2015.
- [W11] **T. Lan**, **L. Chen**, **Z. Deng**, **G.T. Zhou**, and G. Mori. Learning Action Primitives for Multi-Level Video Event Understanding. *Workshop on Visual Surveillance and Re-Identification (at ECCV)*, Zurich, Switzerland, September 2014.
- [W10] **A. Vahdat**, **B. Gao**, **M. Ranjbar**, and G. Mori. A Discriminative Key Pose Sequence Model for Recognizing Human Interactions. *Eleventh IEEE International Workshop on Visual Surveillance (at ICCV)*, Barcelona, Spain, November 2011.
- [W9] **T. Lan**, **Y. Wang**, G. Mori, and S. Robinovitch. Retrieving Actions in Group Contexts. *International Workshop on Sign Gesture Activity (at ECCV)*, Hersonissos, Greece, September 2010.
- [W8] **W. Yang**, **Y. Wang** and G. Mori. Human Action Recognition from a Single Clip per Action. *2nd International Workshop on Machine Learning for Vision-based Motion Analysis (at ICCV)*. Kyoto, Japan, September 2009.
- [W7] G. Mori, **M. Moslemi Naeini**, **A. Rova**, **P. Sabzmeydani**, and **J. Wawerla**. Monitoring Creatures Great and Small: Computer Vision Systems for Looking at Grizzly Bears, Fish, and Grasshoppers. *Workshop on Visual Observation and Analysis of Animal and Insect Behavior (at ICPR)*. Tampa, FL, December 2008.
- [W6] **B. Chen**, **N. Nguyen**, and G. Mori. Geometric Blur in Human Pose Estimation. *IEEE Workshop on Applications of Computer Vision, WACV 2008*. Copper Mountain, CO, January 2008.
- [W5] **B. Chen**, **W. Ma**, **Y. Tan**, A. Fedorova, and G. Mori. GreenRT: A Framework for the Design of Power-Aware Soft Real-Time Applications. *Workshop on the Interaction between Operating Systems and Computer Architecture, WIOSCA 2008*. Beijing, China, June 2008.
- [W4] **Y. Wang** and G. Mori. Boosted Multiple Deformable Trees for Parsing Human Poses. *2nd Workshop on HUMAN MOTION Understanding, Modeling, Capture and Animation (at ICCV)*. Rio de Janeiro, Brazil, October 2007. *accept rate*: $\frac{11}{38} = 28.9\%$.
- [W3] **Y. Wang**, **P. Sabzmeydani**, and G. Mori. Semi-Latent Dirichlet Allocation: A Hierarchical Model for Human Action Recognition. *2nd Workshop on HUMAN MOTION Understanding, Modeling, Capture and Animation (at ICCV)*. Rio de Janeiro, Brazil, October 2007. *accept rate*: $\frac{11}{38} = 28.9\%$.
- [W2] **C. McIntosh**, G. Hamarneh, and G. Mori. Human Limb Delineation and Joint Position Recovery Using Localized Boundary Models. *IEEE Workshop on Motion and Video Computing, WMVC 2007*. Austin, TX, February 2007.
- [W1] G. Mori and J. Malik. Estimating Human Body Configurations using Shape Context Matching. *Workshop on Models versus Exemplars in Computer Vision (at CVPR)*. Kauai, HI, December 2001.

Non-refereed Publications

- [N11] G. Mori, C. Pantofaru, N. Kothari, T. Leung, G. Toderici, A. Toshev, W. Yang. Pose Embeddings: A Deep Architecture for Learning to Match Human Poses. *arXiv:1507.00302*, July, 2015.
- [N10] S. Oh, A. Perera, I. Kim, M. Pandey, **K. Cannons**, **H. Hajimirsadeghi**, **A. Vahdat**, G. Mori, B. Miller, S. McCloskey, Y. Cheng, Z. Huang, C. Lee, C. Xu, R. Kumar, W. Chen, J. Corso, L. Fei-Fei, D. Koller, V. Ramanathan, K. Tang, A. Joulin, A. Alahi. TRECVID2013 GENIE: Multimedia Event Detection and Recounting *TREC Video Retrieval Evaluation Workshop (TRECVID)*, November 2013.
- [N9] **O. Aziz**, S. Robinovitch, E. Park, and G. Mori. Distinguishing Near-Falls From Activities Of Daily Living Using Triaxial Accelerometers. *Canadian Society for Biomechanics*, June, 2012.
- [N8] A. Perera, S. Oh, M. Leotta, I. Kim, B. Byun, C.-H. Lee, S. McCloskey, J. Liu, B. Miller, **Z. F. Huang**, **A. Vahdat**, **W. Yang**, G. Mori, K. Tang, D. Koller, L. Fei-Fei, K. Li, G. Chen, J. Corso, Y. Fu, R. Srihari. GENIE TRECVID2011 Multimedia Event Detection: Late-Fusion Approaches to Combine Multiple Audio-Visual features. *TREC Video Retrieval Evaluation Workshop (TRECVID)*, November 2011.
- [N7] **Z. F. Huang** and G. Mori. SFU at TRECVID 2010: Surveillance Event Detection. *TREC Video Retrieval Evaluation Workshop (TRECVID)*, November 2010.
- [N6] **W. Yang**, **T. Lan**, and G. Mori. SFU at TRECVID 2009: Event Detection. *TREC Video Retrieval Evaluation Workshop (TRECVID)*, November 2009.
- [N5] **W. Ma**, G. Hamarneh, G. Mori, K. Dinelle, and V. Sossi. Motion Estimation for Functional Medical Imaging Studies Using a Stereo Head Pose Tracking System. *IEEE Medical Imaging Conference*, Dresden, Germany, October 2008.
- [N4] **C. Johnson** and G. Mori. Responsive Video-Based Motion Synthesis. ACM SIGGRAPH / Eurographics Symposium on Computer Animation (poster), San Diego, CA, August 2007.
- [N3] G. Mori. Detecting and Localizing Human Figures. Ph.D. thesis, Computer Science Division, University of California at Berkeley, 2004.
- [N2] G. Mori, A. Berg, A. Efros, A. Eden, and J. Malik. Video Based Motion Synthesis by Splicing and Morphing. *University of California, Berkeley Tech Report: UCB//CSD-04-1337*, June 2004.
- [N1] G. Mori, L. Walker, S.R. Bharadwaj, C. Schor, J. Malik. Do object viewing strategies change when parts are ambiguous?. *European Conference on Visual Perception*, Paris, France, September 2003.

2.5 Selected Invited Talks

- [T41] Deep Structured Models for Human Activity Recognition. *International Conference on Image Processing Theory, Tools, and Applications (IPTA)*, Montreal, Quebec, November 2017.
- [T40] Deep Structured Models for Group Activities and Label Hierarchies. *ACCVAC Workshop*, Keelung, Taiwan, August 2016.
- [T39] Deep Structured Models for Group Activities and Label Hierarchies. *Ecole Normale Supérieure*, Paris, France, June 2016.
- [T38] Deep Structured Models for Group Activities and Label Hierarchies. *National Institute of Informatics (NII)*, Tokyo, Japan, May 2016.
- [T37] Deep Structured Models for Group Activities and Label Hierarchies. *ICCV2015 Workshop: Multi-Sensor Fusion for Dynamic Scene Understanding*, Santiago, Chile, December 2015.
- [T36] Structured Models for Group Activity Analysis. *CVPR GROW Workshop*, Boston, MA, June 2015.
- [T35] Structured Models for Recognition: Towards Sub-Category and Interaction Discovery. *Stanford University*, Stanford, CA, March 2015.
- [T34] Structured Models for Recognition: Towards Sub-Category and Interaction Discovery. *Holistic Scene Understanding Seminar*, Dagstuhl, Germany, February 2015.
- [T33] Social Roles in Hierarchical Models for Human Activity Recognition. *CVPR Workshop on Perceptual Organization in Computer Vision*, Columbus, OH, June 2014.
- [T32] Discriminative Latent Variable Models for Human Action Recognition. *CRV Symposium on Activity Recognition*, Montreal, Quebec, May 2014.
- [T31] Discriminative Latent Variable Models for Human Action Recognition. *ICCV Workshop on Understanding Human Activities: Context and Interactions*, Sydney, Australia, December 2013.
- [T30] Discriminative Latent Variable Models for Human Action Recognition. *CVPR Workshop on Action Similarity in Unconstrained Videos*, Portland, OR, June 2013.
- [T29] Discriminative Latent Variable Models for Human Action Recognition. *Nara Institute of Science and Technology*, Nara, Japan, May 2013.
- [T28] Discriminative Latent Variable Models for Human Action Recognition. *Second Short Spring School in Surveillance S5*, University of Modena and Reggio Emilia, Italy, May 2013.
- [T27] Looking at People in Surveillance Video: Detecting Actions and Vehicle Interactions. *MacDonald, Dettwiler and Associates Ltd. (MDA)*, Richmond, BC, November 2012.
- [T26] Computer Vision Algorithms for Fall Detection. *Canadian Association of Gerontology Symposium*, Vancouver, BC, October 2012.
- [T25] Max-margin Learning of Models of Human Action. *University of British Columbia*, Vancouver, BC, January 2012.
- [T24] Learning Structured Models for Recognizing Human Actions. *University of Waterloo*, Waterloo, ON, July 2011.
- [T23] Learning Structured Models for Recognizing Human Actions. *CVPR Workshop on Gesture Recognition*, Colorado Springs, CO, June 2011.

- [T22] Video Technology for Monitoring and Preventing Falls in Long-term Care. *RESNA Conference Workshop*, Toronto, ON, June 2011.
- [T21] Learning Structured Models for Recognizing Human Actions. *Zhejiang University*, Hangzhou, China, April 2011.
- [T20] Learning Structured Models for Recognizing Human Actions. *MacDonald, Dettwiler and Associates Ltd. (MDA)*, Richmond, BC, January 2011.
- [T19] Recognizing Human Actions and Face Engagement for Human-Robot Interaction. *Keynote Speaker at First International Workshop on Computer Vision for Human-Robot Interaction (CVforHRI 2010) at CVPR*, San Francisco, CA, June 2010.
- [T18] Learning Structured Models for Recognizing Human Actions. *Keynote Speaker at Seventh Canadian Conference on Computer and Robot Vision (CRV 2010)*, Ottawa, ON, June 2010.
- [T17] Recognizing Human Actions from Video Data. *SFU Webcasts in Communication / IEEE Circuits and Systems Society Joint Chapter of the Vancouver/Victoria Sections Colloquium Series*, SFU, April 2009.
- [T16] Recognizing Human Actions from Video Data. *CSMG/MoCSSy Colloquium Series*, SFU, January 2009.
- [T15] Monitoring Creatures Great and Small: Computer Vision Systems for Looking at Grizzly Bears, Fish, and Grasshoppers. *Workshop on Visual Observation and Analysis of Animal and Insect Behavior (at ICPR)*, Tampa, FL, December 2008.
- [T14] Boosted Multiple Deformable Trees for Parsing Human Poses. *EHuM2: 2-nd Workshop on Evaluation of Articulated Human Motion and Pose Estimation*, Minneapolis, MN, June 2007.
- [T13] Detecting Pedestrians by Learning Shapelet Features and Boosted Multiple Deformable Trees for Parsing Human Poses. *Carnegie Mellon University VASC Seminar*, Pittsburgh, PA, May 2007.
- [T12] Detecting Pedestrians by Learning Shapelet Features and Boosted Multiple Deformable Trees for Parsing Human Poses. *Toyota Technical Institute - Chicago*, Chicago, IL, May 2007.
- [T11] Detecting Pedestrians by Learning Shapelet Features. *Tokyo Institute of Technology*, Tokyo, Japan, May 2007.
- [T10] Estimating Human Body Pose in Still Images. *BIRS 2006 Workshop on Mathematical Methods in Computer Vision*, Banff, AB, October 2006.
- [T9] Looking at People... and Animals. *York University CVR & Computer Science Colloquium*, Toronto, ON, June 2006.
- [T8] Human Body Pose Estimation in Static Images. *Canadian Institute for Advanced Research (CIAR) Neural Computation & Adaptive Perception*, Toronto, ON, July 2005.
- [T7] Recognizing Human Figures and Actions. *UBC Vision and Robotics Group*, Vancouver, BC, October 2004.
- [T6] Recognizing Human Figures and Actions. *Center for Scientific Computing*, Burnaby, BC, September 2004.
- [T5] Recovering Human Body Configurations: Combining Segmentation and Recognition. *IEEE Conf. on Computer Vision and Pattern Recognition*, Washington, D.C., June 2004.
- [T4] Recognizing Objects in Adversarial Clutter: Breaking a Visual CAPTCHA. *IEEE Conf. on Computer Vision and Pattern Recognition*, Madison, WI, June 2003.
- [T3] Recognizing Objects in Adversarial Clutter: Breaking a Visual CAPTCHA. *Bay Area Vision Meeting*, U.C. Santa Cruz, June 2003.

- [T2] Estimating Human Body Configurations using Shape Context Matching. *Workshop on Models versus Exemplars in Computer Vision (at CVPR01)*, Kauai, HI, December 2001.
- [T1] Estimating Human Body Configurations using Shape Context Matching. *Bay Area Vision Meeting*, Compaq, Palo Alto, CA, October 2001.

2.6 Research Funding

The following summary table shows external research funding for which I was a **principal investigator**.

	Type	Source	Awarded	End	Annual	Total
G49	Operating	MITACS	2017	2017	\$15,000	\$15,000
G48	Operating	MITACS	2017	2017	\$30,000	\$30,000
G47	Operating	NSERC-EPG	2016	2016	\$25,000	\$25,000
G46	Operating	NSERC-DAS	2016	2019	\$40,000	\$120,000
G45	Operating	NSERC-DG	2016	2021	\$60,000	\$300,000
G44	Operating	Research Contract	2016	2017	\$102,250	\$102,250
G43	Operating	NSERC-EG	2016	2016	\$25,000	\$25,000
G42	Operating	NSERC-SPG	2015	2018	\$176,500	\$529,500
G41	Operating	NSERC-SPG	2015	2018	\$175,899	\$527,697
G40	Operating	Disney Research	2015	2016	\$260,000	\$260,000
G39	Operating	NSERC-EG	2015	2015	\$25,000	\$25,000
G38	Equipment	NSERC-RTI	2015	2015	\$112,922	\$112,922
G37	Operating	MITACS	2015	2015	\$120,000	\$120,000
G35	Operating	NSERC-EG	2014	2014	\$25,000	\$25,000
G34	Operating	NSERC-EPG	2014	2014	\$25,000	\$25,000
G33	Operating	Disney Research	2014	2014	\$12,500	\$12,500
G32	Operating	NSERC-EG	2014	2014	\$25,000	\$25,000
G31	Operating	MITACS	2014	2014	\$15,000	\$15,000
G30	Operating	NSERC-EG	2014	2014	\$25,000	\$25,000
G29	Operating	MITACS	2013	2013	\$15,000	\$15,000
G28	Operating	MDA-ISTPC/NSERC	2013	2014	\$72,800	\$72,800
G27	Operating	Google	2013	2014	\$32,791	\$32,791
G26	Equipment	NSERC-RTI	2013	2013	\$60,124	\$60,124
G25	Operating	MITACS	2012	2012	\$15,000	\$15,000
G23	Operating	Nokia	2012	2012	\$12,500	\$12,500
G22	Operating	NSERC-EG	2012	2012	\$25,000	\$25,000
G21	Operating	MITACS	2011	2012	\$35,000	\$35,000
G20	Operating	IARPA	2011	2014	\$151,695	\$455,085
G19	Operating	Google	2011	2011	\$49,582	\$49,582
G18	Operating	NSERC-EG	2011	2011	\$25,000	\$25,000
G17	Operating	NSERC-DG	2011	2016	\$42,000	\$210,000
G16	Operating	NSERC-EG	2011	2011	\$24,778	\$24,778
G15	Operating	NSERC-EG	2010	2010	\$25,000	\$25,000
G14	Operating	BCFRST-NRAS	2010	2013	\$99,466	\$298,398
G13	Operating	MITACS	2010	2010	\$15,000	\$15,000
G11	Operating	MITACS	2008	2009	\$15,000	\$15,000
G10	Operating	NVIDIA	2008	2009	\$25,000	\$25,000
G9	Operating	NSERC-DAS	2008	2011	\$40,000	\$120,000
G8	Operating	NSERC-DG	2008	2011	\$24,000	\$72,000
G7	Operating	NSERC-CRD	2007	2008	\$30,000	\$30,000
G5,G6	Operating	MITACS	2006	2007	\$15,000	\$15,000
G4	Operating	NSERC-CRD	2006	2007	\$40,500	\$40,500
G3	Equipment	CFI/BCKDF	2005	2006	\$300,000	\$300,000
G2	Operating	NSERC-DG	2005	2008	\$19,000	\$57,000
G1	Operating	SFU-PRG	2004	2005	\$10,000	\$10,000
Total:					\$3,902,381 Operating + \$473,046 Equipment	

The descriptions below detail grants in which I have played a major role, either as principal investigator or primary co-investigator. It does not include grants in which I was listed as collaborator. SFU start-up funds are also omitted.

Descriptions

- [G49] Accelerate BC (MITACS) Internship. Applying Deep Learning to Optimize 3D Pose Estimation from Monocular Video. J. Smith and G. Mori: \$15,000 (2017)
- [G48] Accelerate BC (MITACS) Internship. Human Detection from Videos based on Deep Learning. X. Liu and G. Mori: \$30,000 (2017)
- [G47] NSERC Engage Plus Grants (EPG). Deep Learning for Dermoscopy Images (with MetaOptima). G. Mori: \$25,000 (2016)
- [G46] NSERC Discovery Accelerator Supplement. Structured Models for Human Activity Recognition. \$120,000 (\$40,000 per year, 2016-2019)
- [G45] NSERC Discovery Grant. Structured Models for Human Activity Recognition. \$300,000 (\$60,000 per year, 2016-2021)
- [G44] Pantoscope Media Research Contract. Utilizing User Behavior Data in Tag-Based Photo Retrieval. G. Mori: \$102,250 (2016-2017)
- [G43] NSERC Engage Grants (EG). Adaptive Deep Learning for Visualisation of Diagnostic Features of Dermoscopy Images (with MetaOptima). G. Mori: \$25,000 (2016)
- [G42] NSERC Strategic Project Grants (SPG). DEEPCISION: Seeing and Understanding Humans with Deep Structured Models. G. Taylor and G. Mori: \$529,500 (50%) (2015-2018)
- [G41] NSERC Strategic Project Grants (SPG). Changing the Game: Novel Computer Vision Algorithms for Automated Sports Analytics. G. Mori, J. Little, R. Urtasun: \$527,697 (33%) (2015-2018)
- [G40] Disney Research. Deep Structured Models for Video Analytics: Understanding Individual and Group Activities in Video Data. G. Mori: \$260,000 (2015)
- [G39] NSERC Engage Grants (EG). Learning to Analyze Sentiment in Unconstrained Internet Images (with Netra Inc.). G. Mori: \$25,000 (2015)
- [G38] NSERC Research Tools and Instruments (RTI). Vision based autonomous navigation and HRI for a team of flying robots P. Tan, G. Mori, and R. Vaughan: \$112,922 (25%) (2015)
- [G37] MITACS Accelerate Internship Cluster. Human Activity Analysis in Sports Videos. G. Mori and M. Levine. \$120,000 (50%) (2015)
- [G36] Brain Canada Foundation. Novel Retinal Biomarkers for Alzheimer's Disease. M. F. Beg, M. Sarunic, J. Graham, G. Mori, J. Matsubara, G.-Y. Hsiung, P. Mackenzie, A. Merkur, I. Mackenzie. \$1,496,634 (5%) (2014-2017)
- [G35] NSERC Engage Grants (EG). Utilizing User Behavior Data in Tag-Based Photo Retrieval (with Pantoscope Media). G. Mori: \$25,000 (2014)
- [G34] NSERC Engage Plus Grants (EPG). Improved face analysis for school photo collections (with Vidigami). G. Mori: \$25,000 (2014)
- [G33] Disney Research. Hierarchical Max-margin Clustering and Applications. G. Mori: \$12,500 (2014)

- [G32] NSERC Engage Grants (EG). Matching of social media profile photographs (with ThinkCX). G. Mori: \$25,000 (2014)
- [G31] Accelerate BC (MITACS) Internship. Using Machine Learning Techniques to Improve Automatic Keyword Extraction from Textual Web Content. H. Hajimirsadeghi and G. Mori: \$15,000 (2014)
- [G30] NSERC Engage Grants (EG). Photo Quality and Content Detection for Image Collections (with Vidigami). G. Mori: \$25,000 (2014)
- [G29] Accelerate BC (MITACS) Internship. Using Machine Learning Methods to Improve Image Suggestion and Image Retrieval Results. A. Bakhtiari and G. Mori: \$15,000 (2013)
- [G28] MDA Corporation - ISTPC / NSERC Joint Program. Complex Behavior Recognition from Aerial Video (CBRAV). G. Mori: \$72,800 (2013)
- [G27] Google Research Awards. Weakly Supervised Object Discovery for Learning Image and Video Classifiers. G. Mori: \$32,791 (2013)
- [G26] NSERC Research Tools and Instruments (RTI). HRI in the Sky: Command and Control of Teams of Flying Robots by Uninstrumented Operators. R. Vaughan and G. Mori: \$60,124 (33%) (2013)
- [G25] Accelerate BC (MITACS) Internship. Machine Learning for User Behaviour Prediction in Mobile Games. A. Bakhtiari and G. Mori: \$15,000 (2012)
- [G24] NSERC Strategic Project Grant (SPG). GreenPhones: Energy-smart Software for Ubiquitous Mobility. A. Fedorova, A. Shriraman, E. Stroulia, G. Mori, A. Sarkar: \$441,928 (15%) (2012)
- [G23] Nokia University Cooperation Funds. Event Recognition in Cell Phone Videos. K. Cannons and G. Mori: \$12,500 (2012)
- [G22] NSERC Engage Grants (EG). Human activity recognition from aerial surveillance video (with MDA). G. Mori: \$25,000 (2012)
- [G21] MITACS Elevate Strategic Fellowship. Learning Structured Labels and their Spatiotemporal Extents in Unstructured Internet Videos. K. Cannons and G. Mori: \$35,000 (2011)
- [G20] IARPA ALADDIN. General Engine for Indexing Events (GENIE), SFU sub-contract from Kitware Inc. (prime). G. Mori: \$455,085 (2011-2014)
- [G19] Google Research Awards. Structured and Localized Tagging of Internet Videos. G. Mori: \$49,582 (2011)
- [G18] NSERC Engage Grants (EG). Optical character recognition in the wild (with Trusterra Technologies Inc.). G. Mori: \$25,000 (2011)
- [G17] NSERC Discovery Grant. Recognizing Human Figures and Actions. \$210,000 (\$42,000 per year, 2011-2016)
- [G16] NSERC Engage Grants (EG). Developing a high-performance logo recognition algorithm for videos (with BroadbandTV Corp.). G. Mori: \$24,778 (2011)
- [G15] NSERC Engage Grants (EG). Content Retrieval and Visualization for Visual Effects (with Spin Pro Vfx). G. Mori: \$25,000 (2010)
- [G14] BCFRST Natural Resources and Applied Sciences Research Team Program (NRAS). Automatic Detection and Tracking of Pedestrians for Transportation and Traffic Engineering. G. Mori and T. Sayed: \$298,398 (\$99,466 per year 2010-2013) (40%)

- [G13] Accelerate BC (MITACS) Internship. Video Analytics Technology for Fishery Monitoring Services. F. Stefanus and G. Mori: \$15,000 (2010)
- [G12] CIHR Emerging Team Grant. CIHR Team in the Cause and Prevention of Falls in Residential Care. S. Robinovitch (PI) and 14 others. G. Mori (co-applicant) \$1,500,000 (\$300,000 per year 2009-2014) (5%)
- [G11] Accelerate BC (MITACS) Internship. Face Recognition in Personal Photo Collections using Binary Quadratic Programming. M. Ranjbar and G. Mori: \$15,000 (2008)
- [G10] NVIDIA Professor Partnership. Efficient Pedestrian Detection on the GPU. \$25,000 (2008)
- [G9] NSERC Discovery Accelerator Supplement. Recognizing Human Figures and Actions. \$120,000 (\$40,000 per year, 2008-2010)
- [G8] NSERC Discovery Grant. Recognizing Human Figures and Actions. \$72,000 (\$24,000 per year, 2008-2011)
- [G7] NSERC Collaborative Research and Development (CRD). Semi-Automatic Video Motion Capture. G. Mori and T. Möller: \$30,000 (2007) (100%)
- [G6] MITACS Internship. Semi-Automatic Video Motion Capture. C. Johnson and G. Mori: \$7,500 (2006)
- [G5] MITACS Internship. Semi-Automatic Video Motion Capture. P. Sabzmejdani and G. Mori: \$7,500 (2006)
- [G4] NSERC Collaborative Research and Development (CRD). Semi-Automatic Video Motion Capture. G. Mori and T. Möller: \$40,500 (2006) (100%)
- [G3] Canadian Foundation for Innovation (CFI) and BC Knowledge Development Fund (BCKDF). Research Facility for Scientific Data Acquisition, Transmission and Storage (SDATS). R. Vaughan, G. Mori, A. Bulatov, F. Ergun, J. Liu, and J. Pei: \$300,000 (2005) (25%)
- [G2] NSERC Discovery Grant. Recognizing Human Figures and Actions. \$57,000 (\$19,000 per year, 2005-2007)
- [G1] Simon Fraser University President's Research Grant: \$10,000 (2004)

3 Teaching

3.1 Award

Excellence in Undergraduate Teaching Award 2006

This award is handed out annually by the SFU Undergraduate Computing Science Student Society (CSSS) to Computing Science faculty to recognize their efforts in teaching undergraduate courses in our department. I received this award in 2006.

3.2 Graduate Student Supervision

Summary of graduate student senior supervisory duties:

Ph.D.		M.Sc.		Total
Active	Complete	Active	Complete	
6	15	0	36	57

SFU operates on a trimester system. In the tables below -1, -2, and -3 denote Spring, Summer, and Fall terms respectively. My graduate students who have completed their degrees did so in a timely fashion (avg. 5.9 trimesters for MSc students, 13 trimesters for PhD students).

Senior Supervision – Graduated Students

Legend: Name, Degree, Tenure, Thesis/Project Title, Publications, Future positions

15. Lei Chen. Ph.D. 2016-1 to 2022-1. Deep Networks for Weakly-supervised Localization and Visual Grounding. [C101, C97, C93, C91, C80, C64, C60, W17, W14, W11]. Next Amazon.
14. Mengyao Zhai. Ph.D. 2016-1 to 2021-3. Towards Lifelong Learning for Generative Adversarial Networks [C101, C96, C97, C91, C80, C64, C60, W17, W14]. Next Borealis AI.
13. Nazanin Mehrasa. Ph.D. 2017-2 to 2021-1. Towards Event Analysis in Time-series Data: Asynchronous Probabilistic Models and Learning from Partial Labels. [C89, C88, C76, C71, W18, W20]. Next Borealis AI.
12. Srikanth Muralidharan. Ph.D. 2016-2 to 2020-2. Multi-person Video Understanding with Deep Neural Networks. [C75, C71, C70, C64]. Next Huawei.
11. Jiawei He. Ph.D. 2015-3 to 2019-3. Theoretical and Applicational Advances in Variational Autoencoders. [C90, C91, C87, C88, C82, C78, W16, W15]. Next Borealis AI.
10. Zhiwei Deng. Ph.D. 2016-1 to 2019-3. Learning Deep Structured Models for Visual Understanding. [J21, C86, C84, C80, C78, C74, C69, C68, C70, C64, W11]. Next Postdoc at Princeton University.
9. Moustafa Ibrahim. Ph.D. 2015-1 to 2018-3. Deep Models for Multi-Person Activity Understanding. [C85, C78, C70]. Next Flex AI.
8. Guang-Tong Zhou. Ph.D. 2012-2 to 2015-3. Toward Scene Recognition by Discovering Semantic Structures and Parts. [J21, C58, C55, C45, W12, W11]. Next Oracle Labs Vancouver.
7. Hossein Hajimirsadeghi. Ph.D. 2011-2 to 2015-3. Multiple Instance Learning for Visual Recognition: Learning Latent Probabilistic Models. [J18, J16, J13, C65, C61, C47, C41, W12, N10]. Next Oracle Labs Vancouver.

6. Arash Vahdat. Ph.D. 2011-3 to 2014-3. Weakly Supervised Models For Recognizing And Clustering High-Level Complex Events In Video. [J17, J13, C61, C58, C55, C53 C52, C48, C42, C39, C36, W10, N10, N8]. Next University Research Associate (soft money faculty) at SFU.
 5. Nataliya Shapovalova. Ph.D. 2011-3 to 2014-3. Towards Action Recognition and Localization in Videos with Weakly Supervised Learning. [C62, C54, C39]. Next Amazon.
 4. Tian Lan. Ph.D. 2010-3 to 2013-2. From Flat to Hierarchical: Modeling Structures in Visual Recognition. [J12, J9, C55, C51, C45, C46, C44, C40, C39, C37, C34, C30, W9, N6]. Next postdoc at Stanford.
 3. Weilong Yang. Ph.D. 2010-2 to 2012-3. Discriminative Latent Variable Models For Visual Recognition. [J9, C45, C42, C40, C33, C30, C26, C21, B2, W8, N6, N8]. Next Google Research.
 2. Mani Ranjbar². Ph.D. 2008-2 to 2012-1. Optimizing Non-Decomposable Loss Functions In Structured Prediction. [J12, J8, C36, C28, C25, C19, W10]. Next D-Wave Systems.
 1. Yang Wang. Ph.D. 2005-3 to 2009-2. Learning Structured Models for Human Actions and Poses. [J6, J5, C30, C29, C27, C28, C26, C21, C20, C17, C16, C15, C9, B2, W8, W4, W3]. Next NSERC Postdoctoral Fellowship at University of Illinois at Urbana-Champaign, now Assistant Professor at University of Manitoba.
36. Mohammad Hadi Salari. M.Sc. 2019-3 to 2021-2. Kronecker-factored Hessian Approximation for Continual Learning. Next MetaOptima.
 35. Hamed Shirzad. M.Sc. 2019-3 to 2021-2. Graph Generation Using Tree Decomposition. Next UBC PhD student.
 34. Yu Gong. M.Sc. 2017-3 to 2019-2. Latent Structure Learning in Variational Autoencoders. [C87, W16]. Next SFU PhD student.
 33. Ruizhi Deng. M.Sc. 2017-3 to 2019-2. Adaptive Appearance Rendering. [C84, C80]. Next SFU PhD student.
 32. Yifang Fu. M.Sc. 2016-3 to 2018-3. Deep Video Visual Relation Detection. [C86]. Next Microsoft.
 31. Akash Abdu Jyothi. M.Sc. 2016-3 to 2018-2. Generating Natural Language Summaries for Image Sets. Next SFU PhD student.
 30. Nelson Nauata³. M.Sc. 2016-3 to 2018-1. Structured Label Inference for Visual Understanding. [J21, W13]. Next SFU PhD student.
 29. Jon Smith. M.Sc. 2016-3 to 2017-3. REP3D: 3D Human Motion Capture Dataset for Athletic Movement. [W13]. Next Deep Learning Engineer at Canadian Broadcasting Corporation.
 28. Xiaoyu Liu. M.Sc. 2015-3 to 2017-2. Joint Constrained Clustering and Feature Learning based on Deep Neural Networks. Next Deep Learning Researcher at Altumview Systems.
 27. Nazanin Mehrasa⁴. M.Sc. 2015-3 to 2017-1. Learning Person Trajectory Representations for Team Activity Analysis. [C76, C71]. Next SFU PhD student.
 26. Karoon Rashedi Nia. M.Sc. 2015-3 to 2017-1. Automatic Building Damage Assessment Using Deep Learning and Ground-Level Image Data. [C72]. Next Oracle Labs Vancouver.

²co-supervised with Prof. Ze-Nian Li

³Co-supervised with Prof. Richard Zhang

⁴co-supervised with Prof. Luke Bornn

25. Yatao Zhong. M.Sc. 2015-3 to 2017-1. Learning Person Trajectory Features for Sports Video Analysis. [C76] Next Oracle Labs Vancouver.
24. Srikanth Muralidharan. M.Sc. 2014-3 to 2016-1. A Hierarchical Deep Temporal Model for Group Activity Recognition. [C70, C64]. Next SFU PhD student.
23. Lei Chen. M.Sc. 2013-3 to 2015-3. Learning Action Primitives for Multi-Level Video Event Understanding. [C64, C60, W11]. Next SFU PhD student.
22. Zhiwei Deng. M.Sc. 2014-1 to 2015-3. Deep Structured Models for Group Activity Recognition. [C64, W11]. Next SFU PhD student.
21. Mengyao Zhai. M.Sc. 2013-3 to 2015-3. Object Detection in Surveillance Video from Dense Trajectories. [C64, C60]. Next SFU PhD student.
20. Mehran Khodabandeh⁵. M.Sc. 2013-1 to 2015-1. Discovering Human Interactions in Videos with Limited Data Labeling. [C60, W12]. Next SFU PhD student.
19. Jinling Li. M.Sc. 2012-3 to 2014-2. Road User Detection and Analysis in Traffic Surveillance Videos. [J16, C60, C47, C59]. Next Netra.
18. Yasaman Sefidgar. M.Sc. 2012-1 to 2014-1. Discriminative Key-Segment Model for Interaction Detection. [J17]. Next The Jonah Group.
17. Amir Hossein Bakhtiari. M.Sc. 2011-3 to 2013-3. Detecting Pedestrians Using Motion Patterns: A Latent Tracking Approach. Next Vidigami.
16. Zhi Feng Huang. M.Sc. 2010-1 to 2012-2⁶. Latent Boosting For Action Recognition. [C33, N7, N8]. Next Apple.
15. Brian Milligan⁷. M.Sc. 2009-3 to 2012-2⁸. Selecting And Commanding Groups Of Robots Using A Vision-Based Natural User Interface. [C31]. Next Big Park / Microsoft.
14. Pengfei Yu. M.Sc. 2009-3 to 2011-2. Image Classification Using Latent Spatial Pyramid Matching. Next Microsoft, now Facebook.
13. Bo Gao. M.Sc. 2009-3 to 2011-2. Exemplar-Based Human Interaction Recognition: Features And Key Pose Sequence Model. [W10]. Next Trusterra Inc., now Microsoft.
12. Arash Vahdat⁹. M.Sc. 2009-3 to 2011-1. A Key Pose Model For Human Interaction Recognition And Color From Gray By Optimized Color Ordering. [W10]. Next SFU Ph. D. student.
11. Ferdinand Stefanus¹⁰. M.Sc. 2009-1 to 2010-3. Automatic Pedestrian Detection and Tracking with a Multiple-Cue Max-Margin Framework. [C25, J8, J10]. Next MacDonald, Dettwiler and Associates Ltd. (MDA).
10. Tian Lan. M.Sc. 2009-1 to 2010-2. Beyond Actions: Discriminative Models For Contextual Group Activities. [C30, N6, W9]. Next SFU Ph.D. student.
9. Bahman Yari Saeed Khanloo. M.Sc. 2008-3 to 2010-2. Combining Simple Trackers Using Structural SVMs For Offline Single Object Tracking. [C25, J8, J10]. Next Aalto University Ph.D. student.

⁵co-supervised with Prof. Ze-Nian Li

⁶On leave 2011-3 to 2012-1.

⁷co-supervised with Prof. Richard Vaughan

⁸On leave 2011-1 to 2012-2.

⁹co-supervised with Prof. Mark Drew

¹⁰co-supervised with Prof. Ze-Nian Li

8. Weilong Yang. M.Sc. 2008-3 to 2010-1. Transfer Learning for Human Action Recognition. [C26, C21, B2, W8, N6]. Next SFU Ph.D. student.
7. Mohammad Norouzi. M.Sc. 2008-1 to 2009-3. Convolutional Restricted Boltzmann Machines for Feature Learning. [C19]. Next University of Toronto Ph.D. student.
6. William Ma¹¹. M.Sc. 2007-3 to 2009-2. Motion Estimation For Functional Medical Imaging Studies Using A Stereo Video Head Pose Tracking System. [W5, N5]. Next computer vision software engineer at ASM Pacific.
5. Alireza Fathi. M.Sc. 2006-3 to 2008-2. Efficient Human Figure Tracking Using Motion Exemplars. [C14, C13]. Next Georgia Tech Ph.D. student.
4. Maryam Moslemi. M.Sc. 2006-1 to 2007-2. Clustering and Visualizing Actions of Humans and Animals Using Motion Features. [C10]. Next M.Sc. student at University of Illinois at Urbanan-Champaign, now Salesforce.com.
3. Christopher Johnson. M.Sc. 2005-3 to 2007-2. Responsive Video-Based Motion Synthesis Using Motion Graphs. [N4]. Next Blast Radius.
2. Andy Rova¹². M.Sc. 2005-1 to 2007-1. Eigen-CSS Shape Matching And Recognizing Fish In Underwater Video. [C11].
1. Payam Sabzmejdani. M.Sc. 2004-3 to 2006-3. Detecting Pedestrians in Still Images Using Learned Shape Features. [J4, C12, W3]. Next Koolhaus Games, now AirG.

Senior Supervision – Currently Active Students

	Name	Degree	Pubs.	Began
1.	Ruizhi Deng	Ph.D.	[C98, C84, C80, W20, W21]	2019-3
2.	Yu Gong	Ph.D.	[C99, C87, W19, W16]	2019-3
3.	Sha Hu	Ph.D.	[C100, C94]	2016-3
4.	Akash Abu Jyothi	Ph.D.	[C90, C88]	2018-3
5.	Mehran Khodabandeh	Ph.D.	[C71, W12]	2015-3
6.	Megha Nawhal	Ph.D.	[C96, C97, C91, W19]	2017-3

Visiting Students

	Name	Institution	Project Title	Tenure
1.	Fabien Baradel	INSA Lyon - LIRIS	Human Action Recognition	2018-1
2.	Antoine Boiteau	Université de Tours	Motion-based Human Action Recognition	2009-2
3.	Louis-Kenzo Cahier	University of Nice Sophia-Antipolis	Experiments in Pedestrian Tracking	2008-2

¹¹co-supervised with Prof. Ghassan Hamarneh

¹²co-supervised with Prof. Mark Drew

SFU Supervisory Committee Membership

	Name	Degree	Role	Status	Completed
1.	Ye Lu	Ph.D.	SFU Examiner	Completed	2005-1
2.	Mauricio Zuluaga	M.Sc.	SFU Examiner	Completed	2005-1
3.	Lilong Shi	M.Sc.	Supervisor	Completed	2005-2
4.	Roosbeh Mottaghi	M.A.Sc.	SFU Examiner	Completed	2006-1
5.	Hao Jiang	Ph.D.	SFU Examiner	Completed	2006-2
6.	Cheng Lu	Ph.D.	SFU Examiner	Completed	2006-2
7.	Varun Jain	M.Sc.	Supervisor	Completed	2006-2
8.	Xiaoxing Li	M.Sc.	Supervisor	Completed	2006-3
9.	Pooya Karimian	M.Sc.	Supervisor	Completed	2007-2
10.	Pawel Zebrowski	M.Sc.	Supervisor	Completed	2007-2
11.	Omer Ishaq	M.Sc.	Supervisor	Completed	2008-1
12.	Yan Lu	M.A.Sc.	SFU Examiner	Completed	2008-2
13.	Brian Fraser	Ph.D.	SFU Examiner	Completed	2009-1
14.	Muntaseer Salahuddin	M.Sc.	SFU Examiner	Completed	2009-1
15.	Gustavo Frigo	M.Sc.	SFU Examiner	Completed	2009-2
16.	Jiawei Huang	M.Sc.	SFU Examiner	Completed	2009-2
17.	Farhad Hormozdiari	M.Sc.	SFU Examiner	Completed	2009-2
18.	Gholamreza Haffari	Ph.D.	Supervisor	Completed	2009-2
19.	Nasim Hajari	M.Sc.	Supervisor	Completed	2009-3
20.	Ajeet Grewal	M.Sc.	SFU Examiner	Completed	2009-3
21.	Alex Couture-Beil	M.Sc.	Supervisor	Completed	2010-1
22.	Ziming Zhang	M.Sc.	SFU Examiner	Completed	2010-1
23.	Ali Kamali	M.Sc.	Supervisor	Completed	2010-2
24.	Andrei Missine	Ph.D.	SFU Examiner	Completed	2010-2
25.	Ann Clifton	M.Sc.	SFU Examiner	Completed	2010-2
26.	Abbas Sadat	M.Sc.	Supervisor	Completed	2010-3
27.	Matt Olson	Ph.D.	SFU Examiner	Completed	2011-1
28.	Porus Patell	M.Sc.	Supervisor	Completed	2011-1
29.	Paul Wighton	Ph.D.	Supervisor	Completed	2011-2
30.	Sebastian Schmidt	M.Sc.	SFU Examiner	Completed	2011-2
31.	Chris McIntosh	Ph.D.	SFU Examiner	Completed	2011-3
32.	Oliver van Kaick	Ph.D.	SFU Examiner	Completed	2011-3
33.	Yue Meng Chen	Ph.D.	SFU Examiner	Completed	2012-1
34.	Naghmeh Khodabakhshi	M.Sc.	SFU Examiner	Completed	2012-2
35.	Max Whitney	M.Sc.	SFU Examiner	Completed	2012-2
36.	Tyler Dwyer	M.Sc.	SFU Examiner	Completed	2012-3
37.	Tom Torsney-Weir	M.Sc.	SFU Examiner	Completed	2012-3
38.	Peng Peng	M.Sc.	Supervisor	Completed	2012-3
39.	Xiaochen Dai	M.A.Sc.	SFU Examiner	Completed	2012-3
40.	Jianqiao Li	M.Sc.	Supervisor	Completed	2013-2
41.	Majid Razmara	Ph.D.	SFU Examiner	Completed	2013-2
42.	Vahid Zakeri	Ph.D.	SFU Examiner	Completed	2013-2
43.	Baskaran Sankaran	Ph.D.	Supervisor	Completed	2013-2
44.	Milan Mosny	Ph.D.	SFU Examiner	Completed	2013-2
45.	Farnaz Agahian	M.Sc.	Supervisor	Completed	2013-3
46.	Rohit Dholakia	M.Sc.	SFU Examiner	Completed	2014-1
47.	Siddharth Oli	M.A.Sc.	Supervisor	Completed	2014-1

48.	Shawn Andrews	Ph.D.	Supervisor	Completed	2014-2
49.	Ahmed Abdelsadek	M.Sc.	Supervisor	Completed	2014-3
50.	Pradeep Raamana	Ph.D.	Supervisor	Completed	2014-3
51.	Mirzaalain Dastjerdi Hengameh	Ph.D.	SFU Examiner	Completed	2014-3
52.	Seyed Masoud Nosrati	Ph.D.	Supervisor	Completed	2015-1
53.	Omar Aziz	Ph.D.	Supervisor	Completed	2015-1
54.	Tien Jung (Matthew) Lee	M.A.Sc.	Supervisor	Completed	2015-2
55.	Ann Clifton	Ph.D.	SFU Examiner	Completed	2015-3
56.	Te Bu	M.Sc.	SFU Examiner	Completed	2015-3
57.	Jake Bruce	M.Sc.	Supervisor	Completed	2015-3
58.	Haoyu Ren	Ph.D.	SFU Examiner	Completed	2016-1
59.	Abbas Sadat	Ph.D.	SFU Examiner	Completed	2016-2
60.	Mani Monajjemi	Ph.D.	Supervisor	Completed	2016-2
61.	Shokoofeh Pourmehr	Ph.D.	Supervisor	Completed	2016-3
62.	Lingkang Zhang	M.Sc.	Supervisor	Completed	2017-1
63.	Colin Brown	Ph.D.	Supervisor	Completed	2017-1
64.	Zhaopeng Cui	Ph.D.	Supervisor	Completed	2017-2
65.	Rui Huang	Ph.D.	SFU Examiner	Completed	2017-2
66.	SeyedMehdi (Sepehr) MohaimenianPour	M.Sc.	Supervisor	Completed	2018-1
67.	Jeremy Kawahara	Ph.D.	Supervisor	Completed	2019-2
68.	Jiaxi Tang	Ph.D.	SFU Examiner	Completed	2019-3
69.	Mengliu Zhao	Ph.D.	Supervisor	Completed	2020-1
70.	Ashkan Alinejad	Ph.D.	SFU Examiner	Completed	2021-2

External Examiner

	Name	Degree	Institution	Thesis Title	Completed
1.	G. Gill	Ph.D.	McGill University	Building a model for a 3D object class in a low dimensional space for object detection	Spring 2009
2.	A. Shabani	Ph.D.	University of Waterloo	Human Action Recognition in Video	Summer 2011
3.	K. Okuma	Ph.D.	University of British Columbia	Active Exploration of Training Data for Improved Object Detection	Spring 2012
4.	F. Shi	Ph.D.	University of Ottawa	Local Part Model for Action Recognition in Realistic Videos	Spring 2014
5.	M. Demirkus	Ph.D.	McGill University	A hierarchical temporal probabilistic graphical model for labeling and classifying faces in real-world videos	Summer 2014
6.	I. Radwan	Ph.D.	University of Canberra	Occlusion-Aware, Robust Human Pose Estimation with Applications for Action Recognition	Fall 2014
7.	P. Bojanowski	Ph.D.	Ecole Normale Supérieure	Learning to annotate dynamic video scenes	Summer 2016
8.	M. Ramathan	Ph.D.	Nanyang Technological University	Pose-Invariant Action Recognition for Automated Behaviour Analysis	Fall 2016
9.	M. Shafiee	Ph.D.	University of Waterloo	Randomly-connected Non-Local Conditional Random Fields	Spring 2017
10.	T. Robert	Ph.D.	Sorbonne Université	Improving Latent Representations of ConvNets for Visual Understanding	Fall 2019
11.	E. Crawford	Ph.D.	McGill University	Learning Object-Oriented Models of the Visual World	Summer 2021

3.3 Postdoctoral Fellow Supervision

6. Xiaobin Chang. 2019-3 to 2021-1. Multi-view learning. [C102]. Next Assistant Professor at Sun Yat-sen University
5. Micael Carvalho. 2019-2 to 2019-3. Multi-modal machine learning. [W18]. Next Tesla.
4. Thibaut Durand. 2018-1 to 2019-1. Weakly-supervised learning. [C90, C89, C88]. Next Borealis AI.
3. Fred Tung. 2017-1 to 2019-1. Human action recognition, deep network compression. [J20, C91, C92, C83, C79, C75]. Next Borealis AI.
2. Wang Yan. 2013-3 to 2017-3. Video retrieval. [C63, C61]. Next ScopeMedia.
1. Kevin Cannons. 2011-3 to 2013-2. Human action recognition. [J13, C52, C48, C39, N10]. Next Ikomed.

3.4 Undergraduate Student Supervision

I have supervised research projects by undergraduate students working under the NSERC Undergraduate Student Research Award (USRA) program, directed studies courses (CMPT 415/416), and directly via research assistantships (RA).

List of Undergraduates Supervised at Research

	Name	Program	Notes	Tenure
1.	Jiacheng Chen	RA	[C86, C80], next USC PhD	2017-2 to 2018-2
2.	Changan Chen	RA, CMPT 415	[C94, C83], next UT Austin PhD	2017-2 to 2018-1
3.	Bicheng Xu	VPR-USRA	next UBC MSc	2016-2 to 2017-2
4.	Hexiang Hu	RA	[J21, C68, C69], next UCLA PhD	2015-2 to 2016-1
5.	Jeff Hsu	NSERC USRA		2014-2
6.	Yuhao Liu	RA	[C64, C59], next UC Berkeley MSc	2014-2 to 2014-3
7.	Jordan Yap	NSERC USRA; RA; CMPT 415, 416	[C63]	2013-2 to 2015-3
8.	Yuke Zhu	RA	[C44], next Stanford MSc	2012-2 to 2012-3
9.	Pouria Saghalati	NSERC USRA		2012-2
10.	Youyou Yang	RA	next Cornell MSc, LinkedIn	2011-3 to 2012-1
11.	Wesley May	NSERC USRA	next Toronto MSc	2011-2
12.	Ben Reilly	RA	next Toronto MSc	2011-1
13.	Jia Sun	RA		2011-1
14.	Shinsuke Murata ¹³	RA		2010-3
15.	Zhi Feng Huang	volunteer	next SFU MSc	2009-2
16.	Aditya Ramesh	NSERC USRA	next Stanford MSc	2009-2
17.	Mark Bayazit	NSERC USRA; RA	[C18], next ShipSmartly.com	2008-2 to 2009-1
18.	Angelica Lim	NSERC USRA	Monbukagakusho at Kyoto Univ.	2008-1
19.	Bo Chen	CMPT 415; RA	[W6,W5], next UBC MSc, Caltech PhD	2007-1 to 2008-2
20.	Nhan Nguyen	CMPT 415; 416	[W6], next UBC MSc	2007-1 to 2007-2
21.	Jennifer Fernquist	NSERC USRA	next UBC MSc, now Google	2006-2
22.	Chris Lundgren	NSERC USRA	next Safe Software	2006-2
23.	Russell Warneboldt	CMPT 415	next SFU MSc, Electronic Arts	2005-1
24.	Arseniy Akuney	NSERC USRA	next UBC MSc	2005-2
25.	Ben Hull	CMPT 415		2005-2

3.5 Course Teaching at Simon Fraser University

For each of my courses I maintain a webpage with lecture slides, homework assignments, and other course material. These webpages are linked from my webpage: <http://www.cs.sfu.ca/~mori>.

Summary of SFU teaching schedule with final enrollments

Legend:

1 Spring trimester: Jan. to Apr.

2 Summer trimester: May to Aug.

3 Fall trimester: Sep. to Dec.

¹³Co-supervised with Prof. Torsten Möller

Course	2004	2005			2006			2007			2008			2009			2010		Totals	
	3	1	2	3	1	2	3	1	2	3	1	2	3	1	2	3	1	2	UG	Grad
(1) CMPT 225																	41		41	
(2) CMPT 310	34		34	39			51			57			66						281	
(3) CMPT 419												28				8			36	
(3) CMPT 726												32				55				87
(4) CMPT 882	6			12			10			12					17					57
(5) CMPT 888																	14			14
(6) CMPT 415/6		1	1					2	1										5	
Sub totals:																	363	158		

Course	2011	2012			2013			2014			2015			2016			2017		Totals	
	3	1	2	3	1	2	3	1	2	3	1	2	3	1	2	3	1	2	UG	Grad
(1) CMPT 225		122			131														253	
(3) CMPT 419	19						27			47			46			65			204	
(3) CMPT 726	50						56			89			115			116				426
(6) CMPT 415/6								1				1	1						3	
Sub totals:																	460	426		

Course	2018	2019														Totals				
	3	1	2	3	1	2	3	1	2	3	1	2	3	1	2	UG	Grad			
(3) CMPT 419	148															148				
(3) CMPT 726	187			201														388		
Sub totals:																	148	388		

Grand totals:																	971	972		
----------------------	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	-----	-----	--	--

Course descriptions

- CMPT 225 - Data Structures and Programming:** This is the core second year computer science course in our undergraduate program. It explores ideas of data and program organization that allow complex tasks to be solved in simple and elegant ways. In order to manage the complexity of programs, it presents program design and organization ideas such as abstract data types, data structures and recursion. Standard data structures such as linked lists, binary search trees, red-black trees, hash tables, and heaps are presented. Practical experience of these ideas is illustrated by considering their applications and implementations in Java and C++.
- CMPT 310 - Introduction to Artificial Intelligence:** This course is an upper-division AI survey course. In my offering of this course I provide students with an introduction to important ideas in four aspects of AI: search (e.g. DFS, BFS, IDS, A^*), logic (first-order logic), uncertainty (probability, Bayesian networks, hidden Markov models), and machine learning (e.g. decision trees, neural nets). My programming assignments include competitions that have proven very popular with the students – a tournament of automatic backgammon-playing agents written by students, a hand-written digit recognition contest.
- CMPT 419/726 - Machine Learning:** This is a graduate course on machine learning that I created with Prof. Sarkar and Prof. Schulte to fill a hole in SFU's graduate curriculum. Machine learning techniques are part of the standard toolbox of artificial intelligence, and are essential for areas including computer vision, natural language processing, data mining, bioinformatics, and robotics. Its first offering, that I taught in Summer 2008, attracted 60 students (32 grads and 28 undergrads). The enrollment grew to 161 students (115 grads, 46 undergrads) in the Fall 2015 offering. The total number of students enrolled in the CS grad program is approximately 200.

I developed this course based on Chris Bishop's textbook *Pattern Recognition and Machine Learning*. We covered techniques in the standard machine learning toolkit: generalized linear models (regression, classification), kernel density estimation, nearest-neighbour, neural networks, support vector machines, graphical

models, sampling methods, mixture models, expectation-maximization, combining models (boosting, mixture of experts), dimensionality reduction. Knowledge of this set of tools enables students to conduct research in the aforementioned fields for which machine learning is the lingua franca.

4. **CMPT 882 - Recognition Problems in Computer Vision:** This graduate course is an 800-level seminar-style course. The goal is to introduce students to important problems and approaches in recognition, especially object recognition and looking at people challenges. We discuss classical and recent research papers on these topics. The term’s work culminates in a course project in which students gain in-depth knowledge about solving a recognition problem. Many of these course projects have become conference publications.
5. **CMPT 888 - Human Activity Recognition:** This graduate course is an 800-level seminar-style course. The goal is to teach students the state of the art in vision-based human activity recognition. We cover the major approaches in the literature, for the most part by reading and discussing research papers. Students present recent research papers and carry out a course project exploring an approach to activity recognition in detail.
6. **CMPT 415/416 - Special Research Projects:** I have supervised undergraduate directed studies courses, denoted by CMPT 415/416. These projects ranged from implementing existing research papers (e.g. Warneboldt implemented Kenji Okuma’s ECCV 2004 paper on coupled detection and tracking) under my supervision, to conducting novel research (e.g. Chen and Nguyen’s work on pose estimation that we published at WACV, a reputable regular vision workshop).

3.6 Summary of Student Course Evaluations

Instructor Evaluation

(average over all offerings)	Undergraduate Courses [0 = poor, 4 = excellent]			Graduate Courses [0 = poor, 4 = excellent]		
	CMPT 225	CMPT 310	CMPT 419	CMPT 726	CMPT 882	CMPT 888
Organization and Preparation	3.78	3.74	3.90	3.83	3.81	3.90
Teaching Ability	3.87	3.79	3.94	3.89	3.98	4.00

Course Evaluation

Course Number	(average over all offerings of each course)	
	Difficulty [0 = too difficult, 4 = too easy]	Rating [0 = poor, 4 = excellent]
CMPT 225	1.74	3.45
CMPT 310	1.53	3.44
CMPT 419	1.37	3.62
CMPT 726	1.49	3.63
CMPT 882	1.75	3.89
CMPT 888	1.60	4.00

3.7 Other Teaching

- **Statistical and Structural Recognition of Human Actions: Tutorial at European Conference on Computer Vision (ECCV) 2010, Ivan Laptev and Greg Mori:** This course gave an introduction into novel trends in statistical and structural action recognition and illustrated ideas with examples of successful methods from recent literature. In particular, we covered bag-of-features action recognition and discussed alternative local feature representations and their extensions. We considered current issues in human actions datasets and addressed weakly supervised and unsupervised approaches for human actions. We presented advances in structural modeling of human poses and covered recent structured learning methods for action recognition. While this course mostly covered action recognition in video, we also discussed action recognition from still images such as in the Action Classification Taster Competition of PASCAL VOC 2010.

- **Emerging Topics in Human Activity Recognition: Tutorial at Computer Vision and Pattern Recognition (CVPR) 2014, Michael Ryoo, Ivan Laptev, Sangmin Oh, and Greg Mori:** This tutorial offered a sequence of lectures on active and emerging topics in activity recognition. Starting with the general motivation, history overview and basic bag-of-words techniques, we presented advances in several subproblems of action recognition. In particular, we covered (i) modeling spatio-temporal structure of actions, (ii) group activity recognition, (iii) activity recognition from the first-person view, and (iv) real-world applications of activity recognition.

4 Service

4.1 Service To Simon Fraser University

I served as the Director of the School of Computing Science at SFU. I have served on a variety of regular departmental committees and participated in a number of ad-hoc committees. In addition, I have played an active role in undergraduate recruitment work.

Departmental Administration

- Director, School of Computing Science, SFU 2015-2018
- Associate Director, Research and Industrial Relations, School of Computing Science, SFU AY 2011-2014

University Committees

- Dean of Graduate Studies Search Committee, AY 2012-13
- Senate Graduate Awards Adjudication Committee, SGAAC AY 2009-10

Departmental Committees

- Tenure and Promotion Committee AY 2012-13, 2013-14
- Faculty Search Committee AY 2005-06, 2011-12, 2012-13, 2013-14
- CS Strategic Research Committee AY 2009-10
- CS Director Search Consultation Committee AY 2009-10
- CS Distinguished Lecture Series Committee AY 2009-10
- Graduate Program Committee AY 2004-05, 2008-09
- Undergraduate Program Committee AY 2006-07, 2007-08
- Teaching Quality Committee AY 2006-07
- Social Committee AY 2006-07
- Led ad-hoc committee for FAS file server purchase 2004-06

Recruitment work

- SFU Undergraduate Orientation Faculty Mentor 2009
- SFU Graduate Orientation Faculty Panelist 2009
- SFU Open House presenter 2006, 2008
- High school visit presenter 2007, 2008
- Scientific I presenter 2006
- FAS Open House presenter 2006

4.2 Service to the Academic Community

I served on the editorial boards of IEEE Transactions on Pattern Analysis and Machine Intelligence (T-PAMI) and the International Journal of Computer Vision (IJCV), the top journals in computer vision. I regularly serve on the program committees of the major computer vision conferences (CVPR, ICCV, ECCV). I was the program co-chair (for computer vision) of the Canadian Conference on Computer and Robot Vision in 2006 and 2007. I also review research grant proposals for Canadian granting councils (NSERC, CHRP), and was invited as a grant review panelist to the US National Science Foundation in 2008 and 2009. A listing of service duties appears below.

I am also involved in a variety of technology transfer / industry consulting roles. I serve as a Lab Scientist for Creative Destruction Labs (CDL), a national network of seed stage tech incubators.

Journal Editor:

- Associate Editor, IEEE Transactions on Pattern Analysis and Machine Intelligence, T-PAMI 2010-2018
- Editorial Board, International Journal of Computer Vision, IJCV 2013-current
- Associate Editor, IPSJ Transactions on Computer Vision and Applications, CVA 2013-2016
- Guest Editor, Computer Vision and Image Understanding, Special Issue on Group Activities 2015
- Guest Editor, Journal of Image and Vision Computing, Special Issue on Computer and Robot Vision 2007

Conference Program Chair

- IEEE/CVF Computer Vision and Pattern Recognition, CVPR 2020
- Asian Conference on Computer Vision, ACCV 2018
- IAPR Conference on Machine Vision Applications, MVA 2017
- IEEE Winter Conference on Applications of Computer Vision, WACV 2016
- Canadian Conference on Computer and Robot Vision, CRV 2006, 2007

Conference Area Chair

- IEEE International Conference on Computer Vision, ICCV 2015
- IEEE Computer Vision and Pattern Recognition, CVPR 2012, 2014, 2017, 2018
- European Conference on Computer Vision, ECCV 2014, 2018
- Neural Information Processing Systems, NIPS 2015, 2017, 2018, 2019
- Asian Conference on Computer Vision, ACCV 2016
- British Machine Vision Conference, BMVC 2018
- IEEE International Conference on Advanced Video and Signal-based Surveillance, AVSS 2012

Conference Organization

- Program Coordination Co-Chair, IEEE Computer Vision and Pattern Recognition, CVPR 2016
- Finance Co-Chair, IEEE Computer Vision and Pattern Recognition, CVPR 2015
- Publicity Co-Chair, IAPR Conference on Machine Vision Applications, MVA 2015

Program Committee Member:

- IEEE Computer Vision and Pattern Recognition¹⁴, CVPR 2005-2011
- IEEE International Conference on Computer Vision, ICCV 2007, 2009, 2011
- European Conference on Computer Vision, ECCV 2006, 2008, 2010, 2012, 2016
- Asian Conference on Computer Vision, ACCV 2010, 2012, 2014
- IEEE Winter Conference on Applications of Computer Vision, WACV 2017
- IAPR International Conference on Pattern Recognition, ICPR 2008, 2010, 2012, 2014
- IEEE Conference on Automatic Face and Gesture Recognition, FG 2011, 2013
- IAPR Conference on Machine Vision Applications, MVA 2009, 2011, 2013, 2015
- International Joint Conference on Artificial Intelligence, IJCAI 2009
- Canadian Conference on Computer and Robot Vision, CRV 2005, 2008-2018
- British Machine Vision Conference, BMVC 2008-2015
- IEEE Workshop on Applications of Computer Vision, WACV 2011-2015
- Asian Conference on Pattern Recognition, ACPR 2013

¹⁴Use of the term Program Committee at the major vision conferences became deprecated circa 2012 or so.

- IEEE Workshop on Motion and Video Computing, WMVC 2009, 2011
- IEEE International Conference on Advanced Video and Signal-based Surveillance, AVSS 2013-2014
- International Conference on Image Analysis and Processing, ICIAP 2015
- DeepVision Workshop at CVPR 2017
- Workshop on Web-scale Vision and Social Media 2015, 2016, 2017
- 5th International Workshop on Human Behavior Understanding 2014
- International Workshop on Web-scale Vision and Social Media 2014
- ECCV Workshop on Web-scale Vision and Social Media 2012
- Second Workshop on Egocentric Vision 2012
- Second International Workshop on Stochastic Image Grammars 2011
- CVPR Fine-Grained Visual Categorization Workshop 2011
- CVPR Gesture Recognition Workshop 2011
- CVPR International Workshop on Human Activity Understanding from 3D Data 2011, 2012
- Workshop on Action Recognition and Pose Estimation in Still Images 2012
- ICPR Workshop on Computer Vision for Analysis of Underwater Imagery 2014
- International Workshop on Situation, Activity and Goal Awareness 2012
- Workshop on Person Oriented Vision 2011
- International Conference on Computer Vision Theory and Applications 2008
- IEEE International Conference On Advanced Video and Signal Based Surveillance 2008
- International Workshop on Machine Learning for Vision-based Motion Analysis, MLVMA 2008, 2009, 2011
- Workshop on Computer Vision Based Analysis in Sport Environments 2006
- First International Workshop on Multimedia Analysis of User Behaviour and Interactions 2008
- First International Workshop on Video Mining 2008
- First Workshop on Egocentric Vision 2009
- First International Workshop on Stochastic Image Grammars 2009
- First International Workshop on Visual Scene Understanding 2009
- Doctoral Spotlight Committee CVPR 2009
- First International Workshop on Computer Vision for Human-Robot Interaction 2010
- Visual Observation and Analysis of Animal and Insect Behavior, VAIB 2010, 2012
- Third Workshop on Human Motion: Understanding, Modeling, Capture and Animation 2010

Journal Reviewer:

- IEEE Transactions on Pattern Analysis and Machine Intelligence, T-PAMI 2004-2010
- International Journal of Computer Vision, IJCV 2005, 2007-2012
- Computer Vision and Image Understanding, CVIU 2006-08, 2010-15
- Journal of Machine Learning Research 2013
- Journal of Artificial Intelligence Research 2014
- IEEE Transactions on Visualization and Computer Graphics 2008, 2009
- IEEE Transactions on Robotics 2008, 2009
- IEEE Transactions on Multimedia 2004, 2010, 2012, 2013
- IEEE Transactions on Image Processing 2006, 2010, 2012, 2013
- IEEE Transactions on Systems, Man and Cybernetics Part C 2008, 2009, 2011
- IEEE Transactions on Information Forensics & Security 2012
- Image and Vision Computing 2006, 2008, 2009, 2011, 2013

- Science 2011
- Communications of the ACM 2013
- Machine Vision and Applications 2014
- IEEE Security & Privacy 2008, 2009
- International Journal of Image and Graphics 2006
- IEEE Transactions on Circuits and Systems for Video Technology 2007, 2014
- IEEE Transactions on Biomedical Engineering 2006
- IEEE Computer Graphics and Applications 2007
- Journal of Mathematical Imaging and Vision 2007
- Pattern Recognition Letters 2007
- Neurocomputing 2007
- Journal of Electronic Imaging 2003
- Computer-Aided Design 2004
- Computer Graphics Forum 2004, 2005

Conference Reviewer:

- IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2004, 2013, 2015
- IEEE International Conference on Computer Vision, ICCV 2005, 2013
- European Conference on Computer Vision, ECCV 2002, 2004
- Neural Information Processing Systems, NIPS 2009, 2010, 2012-14
- International Conference on Learning Representations, ICLR 2017
- IEEE Winter Conference on Applications of Computer Vision, WACV 2018
- IEEE Conference on Automatic Face and Gesture Recognition, FG 2017
- AI & Statistics, AISTATS 2011
- Eurographics 2009
- IEEE International Conference on Robotics and Automation, ICRA 2008, 2009, 2012, 2014
- ACM SIGGRAPH 2004-06, 2009-10
- ACM SIGGRAPH Asia 2015
- Advanced Concepts for Intelligent Vision Systems 2006, 2007
- Workshop on Analysis of Functional Medical Images 2008
- Transportation Research Board (TRB) Annual Meeting 2009-14, 2016
- IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2014
- Eurographics Symposium on Rendering 2011
- Workshop on Automatic Traffic Surveillance 2015

Grant Proposal Reviewer:

- NSERC Collaborative Research and Development 2010
- NSERC Collaborative Health Research Projects 2009, 2010, 2011
- NSERC Discovery Grant 2008-10, 2012, 2014
- Mathematics of Information Technology and Complex Systems (MITACS) 2007-09, 2011-13
- US National Science Foundation 2008, 2009, 2015
- European Research Council (ERC) 2015
- Spain Ministry of Science and Innovation 2009, 2014
- Portuguese Foundation for Science and Technology (FCT) 2012
- Czech Science Foundation 2012

February 2, 2022