David James Fleet

35 Playter Blvd.
Toronto, Ontario
M4K 2W1
(416) 465-6417

Computer Science
University of Toronto
Toronto Ontario
(416) 946-8485

Computer & Mathematical Sciences
University of Toronto Scarborough
Toronto, Ontario
(416) 287-7219, (416) 287-7201

http://www.cs.toronto.edu/~ fleet/

Employment History

Oct. 2003 – present	Professor, Department of Computer and Mathematical Sciences, University of Toronto at Scarborough and Department of Computer Science, University of Toronto
July. 2012 – June 2017	Professor and Chair, Department of Computer and Mathematical Sciences, University of Toronto at Scarborough
Jan. 1999 – Sept. 2003	Member of Research Staff and Area Manager (as of Feb, 2000), Perceptual Document Analysis and Digital Video Analysis Groups Palo Alto Research Center
July 1995 – Dec. 2002	Associate Professor Department of Computing and Information Science, Cross-appointed to Psychology and Electrical Engineering, Queen's University, Kingston
July 1996 – Aug. 1997	Visiting Research Scientist Image Understanding Group, Xerox Palo Alto Research Center (PARC)
July 1996 – Aug. 1997 Jan. 1994 – May 1994	Visiting Scholar Department of Psychology, Stanford University
July 1990 – June 1995	Assistant Professor Department of Computing and Information Science, Cross-appointed to Psychology, Queen's University, Kingston

Academic Degrees

Ph.D.	**Department of Computer Science, University of Toronto, 1991** Supervision: Allan D. Jepson Thesis: *Measurement of Image Velocity*
M.Sc.	**Department of Computer Science, University of Toronto, 1984** Supervision: Allan D. Jepson and John K. Tsotsos Thesis: *Early Processing of Spatiotemporal Visual Information*
B.Sc.	**Queen's University, 1982** Mathematics and Computer Science

Awards / Honours

2017	Faculty Member, Vector Institute
2012	Senior Fellow, Canadian Institute for Advanced Research
2010	Koenderink Prize in Computer Vision
2009	Best Paper Award, British Machine Vision Conference '09
2005	Fellow, Canadian Institute for Advanced Research
2003	Best Paper Award, UIST '03 ACM Symposium on User Interface Software and Technology
2003	Excellence in Science Award, Palo Alto Research Center
2001	Best Paper Runner-Up Award, IEEE Conference on Computer Vision and Pattern Recognition '01
1999	Marr Prize Honorary Mention (runner-up for best paper), IEEE International Conference on Computer Vision '99
1996 – 99	Alfred P. Sloan Research Fellowship
1985 – 87	NSERC Postgraduate Scholarship
1982 – 84	NSERC Postgraduate Scholarship
1982	NSERC Summer Research Grant
1981 – 82	James H. Rattray Memorial Scholarship

Professional Activities

Editorial Boards

- Advisory Board, *IEEE Trans. on Pattern Analysis and Machine Intelligence* (2009–)
- Associate Editor-In-Chief, *IEEE Trans. on Pattern Analysis and Machine Intelligence* (2005–2008)
- Associate Editor, *IEEE Trans. on Pattern Analysis and Machine Intelligence* (2000–2004)
- Guest Editor, *International Journal of Computer Vision*, Special Issue on Human Activity Understanding from 2D and 3D Data (2016)
- Guest Editor, *Computer Vision and Image Understanding*, Special Issue on Human Pose and Gesture (2016)
- Guest Editor, *Computer Vision and Image Understanding*, Special Issue on Motion Analysis (2005)

Conference Chairs (and Senior Program Committess)

Program Co-Chair, European Conference on Computer Vision (2014)
Program Co-Chair, IEEE Conference on Computer Vision and Pattern Recognition (2003)
Area Chair, European Conference on Computer Vision (2012)
Area Chair, Conference on Computer Vision and Pattern Recognition (2004,2007,2008,2010,2012)
Area Chair, International Conference on Computer Vision (2009, 2011, 2013)
Area Chair, International Conference on Machine Learning (2004)
Area Chair, Neural Information Processing Systems, NIPS (2007, 2008)

Program Committees

Asian Conference on Computer Vision, ACCV (2004)
European Conference on Computer Vision, ECCV (2002, 2004)
Workshop on Human Motion: Understanding, Modeling, Capture and Animation (2010)
IAPR International Conference on Pattern Recognition, ICPR (2002, 2004)
IEEE International Conference on Computer Vision, ICCV (1999, 2001, 2003, 2005)
IEEE Conference on Computer Vision and Pattern Recognition, CVPR (1998, 2000, 2001, 2005, 2011, 2013)
IEEE Workshop on Visual Motion (1991)
IEEE Workshop on Motion and Video Computing (2002, 2007)
IEEE Workshop on Statistical Methods for Video Processing (2003)
IEEE Workshop on Event Mining: The Detection and Recognition of Events in Video (2003)
IEEE Workshop on Spatial Coherence for Visual Motion Analysis (2004)
Vision Interface, VI (1993, 1996, 1997, 2000)
Conference on Articularted Motion and Deformable Objects (2006)
Canadian Conference on Computer and Robot Vision, CRV (2004, 2005, 2006)
Robotics Systems and Science, RSS (2006)
NIPS Workshop on Evaluation of Articulated Human Motion and Pose Estimation (2006)
CVPR Workshop on Evaluation of Human Motion and Pose Estimation (2007)
ICCV Inverse Rendering Workshop (2015)

Conference Reviewing

Most major vision conferences (CVPR, ECCV, and ICCV) since 1990, and machine learning conferences (ICML, NIPS), along with miscellaneous vision workshops, and related conferences occasionally (SIGGRAPH, IJCAI and AAAI).

Journal Reviewing

Frequent reviewing:
Computer Vision and Image Understanding;
IEEE Transactions on Image Processing;
IEEE Transactions on Pattern Analysis and Machine Intelligence;
Image and Vision Computing;
International Journal of Computer Vision;
Vision Research

Occasional reviewing:
Artificial Intelligence; Electronics Letters; Electronic Imaging; IEEE Transactions on Robotics and Automation; IEE Proceedings: Vision, Image and Signal Processing; Journal of the Optical Society of America; Journal of Visual Communication and Image Representation; Optical Engineering; Pattern Recognition Letters; Psychological Research; Psychological Science; Visual Neuroscience

Other Technical Reviewing

ACM Distinguished Dissertation Awards;
Kluwer Academic Press: Robotics Monograph Series;
NSERC Discovery and Strategic Grant Proposals;
FCAR Grant Proposals;
NSF Grant Proposals

Membership Activities

ARVO (Assoc. Research on Vision and Ophthalmology),
Senior Member, IEEE (Inst. Electrical and Electronic Engineers),
Chair, IEEE Computers & Communications Kingston Chapter (1994/95)

Research Programme

My research programme addresses fundamental problems in vision, spanning four related areas: machine vision; image processing; visual perception; and visual neuroscience. This includes the development of reliable algorithms for machine vision systems, and theoretical work on the neural basis of visual perception in biological systems. Most of my work to date has focused on *visual motion analysis binocular stereopsis*, *tracking*, and *human pose estimation*. Understanding these visual processes is central to the development of algorithms that would, for example, facilitate the determination of 3D scene layout, the detection and tracking of objects, the inference of the 3d motion of an observer (camera), and motion-based recognition of objects and their activities. Potential applications include novel human-computer interfaces, computer graphics and special effects, assisted automotive navigation, security, safety and surveillance, and markerless human motion capture.

Computer Vision and Image Processing:

The general analysis and interpretation of visual motion requires that one model the relation between world properties and spatiotemporal variations in image intensity. This includes models of the scene dynamics, of object appearance, and of image formation. Early work in visual motion analysis focused on the task of 2D motion estimation, assuming simple models in which intensity is conserved and motion is unique and smooth in local image regions.

My research on visual motion, and early visual processing in general, is aimed at building richer mathematical models that allow us to estimate and interpret the motion of increasingly complex objects in natural image sequences. This research dates from my graduate work on the estimation of optical flow, to my most recent work on probabilistic methods for motion-based segmentation, image-based visual tracking, and the detection and tracking of 3D human pose.

Phase-Based Visual Estimation:
Among my most significant early contributions in machine vision are the development and analysis of phase-based techniques for measuring optical flow (i.e. 2D image velocity), image orientation, and binocular disparity. With Allan Jepson (University of Toronto) and Michael Jenkin (York University), my research on local phase information has involved both experimental and theoretical work. In addition to developing algorithms for phase-based matching, we have examined the stability of phase information, showing how unreliable measurements can be identified and rejected from subsequent analyses. Leif Haglund (Post-doc at Queen's Univ., 1994; now at Saab Research), Keith Langley (University College London) and I have since extended these techniques with the use recursive spatiotemporal filters toward real-time motion estimation, and to the measurement of orientation in small image regions for estimating surface orientation.

Performance Analysis of Optical Flow Techniques:
In collaboration with John Barron and Steven Beauchemin (University of Western Ontario), I conducted an extensive performance analysis of well-known optical flow algorithms. With independent implementations and a common set of synthetic and natural inputs, we showed how well-known techniques produce results which differed significantly in their accuracy and robustness. Phase-based and gradient-based techniques produced the most reliable results. This work was one of the first comprehensive benchmark tests in machine vision, and helped generate an interest in the theoretical and empirical comparison of visual algorithms. Our test suite and results, along with our implementations, have been available to others on the web. They have make a significant impact on the field of visual motion analysis and they remain the benchmark against which new methods are judged.

Modeling Complex Motion and Appearance Changes in Image Sequences:
Optical flow estimation methods work well with smooth, textured surfaces but not with complex deformations, occlusions, and significant changes in image appearance. To deal with more general classes of 2D motion,

Michael Black (Brown University), Yaser Yacoob (University of Maryland), Allan Jepson and I extended the use of parameterized motion models to handle discontinuous motion features (e.g., motion boundaries), as well as domain-specific classes of motion (e.g., the motion of human mouths and the motion of legs of a walking human). We learned linear parameterized models from robust optical flow estimates. The resulting models were then used to constrain the estimation of optical flow, and to detect occurrences of specific classes of motion in image sequences.

To deal with complex forms of appearance changes, Black, Yacoob and I proposed a formulation in which appearance changes were modeled by a linear mixture of causes. These included photometric causes such as lighting variations and specularities, as well as object-specific variations in appearance caused by complex local patterns of motion and occlusion. This research represents an encouraging first step toward the inference of the different physical causes of intensity variations in image sequences, but there exist many open problems in this area.

Horst Haussecker (PARC postdoc, then member of research staff 1999/2000; now at Intel Research) and I derived a linear formulation of a more restricted class of appearance changes based on physical models of image formation and photometry. Such physical models include changes in surface reflection caused by changes in surface orientation with respect to a direction light source, and thermal diffusion in infrared (IR) imaging. The latter formulation was used for tracking unmarked paper in a copier with laser-induced thermal emittance patterns on the paper surface.

Motion Boundaries, Occlusions and Surface Depth:
One of the most significant remaining problems in early motion estimation concerns occlusion boundaries, where the two key assumptions of current optical flow methods, brightness conservation and motion smoothness, are typically violated. Black and I recently formulated a probabilistic solution to this problem with a hybrid state-space model and a particle filter for approximate inference. The state space model included a discrete random variable to represent different motion classes (e.g., smooth motion, or discontinuous motion), and continuous variables to represent the parameters of each motion class. This work won Honorable Mention for the Marr Prize (runner-up for best paper) at the International Conference on Computer Vision in 1999.

Although these initial experiments produced encouraging results the method proved to be unreliable. Oscar Nestares (PARC post-doc 1999-2001; now at Intel Research) and I have since improved the method by introducing a random field of local neighborhoods to encourage spatiotemporal continuity of the inferred surface boundaries, and an empirical edge-based likelihood function to improve boundary localization.

2.5D Visual Tracking and Layered Image Models:
There exists a theoretical gap between model-based tracking of complex objects and early motion estimation, namely, the selection and initialization of models. It seems clear that some form of early analysis should help one select and initialize models, but this remains a largely unexplored topic. Allan Jepson, Michael Black and I have been working on methods for extracting effective representations of visual motion that provide an efficient characterization of the principal moving components of a scene and their relative depths. In future research we hope to use this representation to infer the occurrence of plausible models for subsequent model selection and refinement.

To provide stable image descriptions through time, Allan Jepson, Thomas El-Maraghi (University of Toronto) and I have developed an approach to learning 2D models of image appearance. The approach works with an on-line version of the EM algorithm to identify stable image structure during tracking. In this way the most stable regions can play the most significant role in the tracking, to facilitate tracking over long image sequences with precise image alignment. This work was awarded Best Paper Runner-Up at the IEEE Conference on Computer Vision and Pattern Recognition in 2001.

Perceptual Organization for Rough Document Image Analysis:
Most recently, my group at PARC has focused on the use of perceptual organization and graph-based recognition technologies for parsing the salient visual structure in a general class of document images. This class

includes both formatted documents as well as informal documents such as whiteboard images, annotated documents, sketches, and diagrams. We have produced software for rough document manipulation, especially targetted at whiteboard material, along with a parsing system for hand-drawn flow charts. This work was given the best paper award at UIST in 2003.

3D People Tracking:
The inference of human shape and motion in 3D has become a topic of great interest in the vision community. The problem is difficult because people move in complex ways, having with many degrees of freedom. Their appearance is similarly hard to model due to variations in lighting, to deformations of clothing, and to occlusions. To constrain people tracking, most existing methods assume one or more constraints, such as knowledge of a static background, the existence of multiple views of the person, or that color can be used to find skin regions. Hedvig Sidenbladh (KTH, Sweden), Michael Black and I proposed a Bayesian approach to tracking people in 3D from 2D video. With a motion-based likelihood function based on a robust form of intensity conservation, a particle filter to deal with nonlinear dynamics and observations, a learned parameterized model of human walking motion and manual initialization of the model, we were able to infer the time-varying 3D structure of a single person in unknown cluttered backgrounds in monocular, greyscale video.

More recently, in an attempt to provide more efficient stochastic sampling so that we could handle weaker models of human dynamics, Kiam Choo and I began to consider the use of particle filters with MCMC updates. Applied to the inference of 3D joint configuration from 2D motion capture point data, we found that a particle filter with hybrid Monte Carlo updates produced an estimator more than 2,000 times faster than a conventional particle filter, with similar estimator variance. In combination with richer likelihood functions, combining motion and edge information, I hope this will lead to more effective tracking in high-dimensional spaces with complex dynamics and observation equations.

Learning Latent Variable Models for Human Motion:
With Raquel Urtasun, Jack Wang, Pascal Fua, Aaron Hertzmann, and Neil Lawrence, I have recently worked on a new class of activity-specific latent variable model for representing priors over kinematic pose and motion. There are three variants on the original Gaussian Process Latent Variable Model of Lawrence (2004): the Gaussian Process Dynamical Model (GPDM) extends the GPLVM with a subspace dynamical model (Wang, Fleet, and Hertzmann 2005,2008), the balanced-GPDM which has been applied successfully to monocular people tracking (Urtasun, Fleet and Fua, 2006), a topologically-constrained GPLVM with which one can influence the topology and coherence of the latent space to suit different classes of motion (such as periodic locomotion), to facilitate transitions between activities, and to generalize better; and the multifactor GPLVM (Wang, Fleet and Hertzman, 2007) which allows one to model known factors that combine to determine the pose (such as the subject's identity and the type of activity).

Such latent variable models have been shown to be effective for gait-based recognition (Urtasun, Fleet and Fua, 2006). In 2010 we showed that with such models and a straightforward video-based 3D pose tracker, one can classify physical and emotional attributes of people from monocular video (Livne, Sigal, Troje and Fleet, 2012). We have also shown that a variant of the Restricted Boltzman Machine (called conditional RBMs) can be used to learn rich models of human motion that are very effective for people tracking (Taylor, Sigal, Fleet and Hinton, 2010).

Physics-Based Locomotion Models for Human Pose Tracking:
Despite the extensive use of kinematic models learned from mocap data, existing people trackers still produce unrealistic motions, with feet sliding on or above the ground, or implausible out-of-plane rotations of the body. We have begun to begun to explore an entirely new approach in which physics-based models are used to ensure that the estimated poses and motions are physically realistic. With physics-based models we can enforce physical laws of motion and object interactions (e.g., people maintain balance and make contact with the ground during locomotion), Marcus Brubaker, Hertzmann and I developed a model based on the Anthropomorphic Walker of Kuo (2002) and used it to track people walking on level ground. Brubaker and I developed a more

general model called *Kneed Walker* which can walk and run. It has been shown to track walking people with different speeds, on different slopes and with changing direction. It is also robust to occlusion.

These basic 2D models have since been generalized in a number of interesting ways to 3D full-body locomotion. Motivated by well-known biomechanics properties of human locomotion, Jack Wang formulated objective criteria, which when optimized, lead to an open-loop controller that produces human-like walking (Wang, Fleet and Hertzmann, 2009). Later work incorporated uncertainty, showing how uncertainty and task constraints lead to different, yet natural, styles of human walking (Wang, Fleet and Hertzmann, 2010).

Physics-Based Models for Modeling Human Interactions:
Brubaker, Sigal and I (2009) formulated a natural measure of physical plausibility in order to identify net forces acting on a body that are inconsistent with a given model. This is useful as it provides a way to determine, given a 3D motion, whether or not unexplained contact exists, and it provides a natural measure of error in a model fit to data that one can optimize in order to estimation contact properties (e.g., the ground plane geometry and dynamics).

Large-Scale Indexing and Image Search:
As datasets grow larger the issue of storing, indexing and searching massive collections of images and video becomes ever more important. One potential approach to the problem in general is to map images (or image features, etc) onto compact binary codes (Norouzi and Fleet, 2011; Norouzi, Fleet and Salakhutdinov, 2012). With compact binary codes one can store billions of items in memory at a given time. And with our new multi-index hashing algorithm we can search such datasets with exact nearest neighbours in sub-linear time (Norouzi, Punjani and Fleet, 2013,2014). Norouzi and I (2013) also formulated a family of compositional models for discrete image embedding, generalizing Iterative Quantization and Product Quantization to achieve state-of-the-art nearest neighbour retrieval on large high-dimensional datasets.

3D Structure Determination in CryoEM:
Cryo-EM is a relatively new method for 3D structure determination of single macromolecules (proteins, viruses, or protein-drug complexes) from electron microscopy. Ali Punjani, Marcus Brubaker and I We have a developed a suite of new algorithms comprising effective importance sampling and quadrature for pose estimation, stochastic gradient descent for fast multi-view reconstruction with robustness to the initial guess, and a branch and bound algorithm for fast maximum likelihood pose estimation and tomographic reconstruction. Together, these led to the development of a major software framework called cryoSPARC that is now widely used in the field.

Perception and Neuroscience:

The study of biological visual systems faces many of the same problems as does research on machine vision. Most notably, the theoretical underpinings of how one might infer properties of the world from images is common to both domains of enquiry. In biological vision my work has focused on mathematical formulations of visual processing, to elucidate the functional role of neural activity in the visual cortex.

Second-Order Motion and Stereo Perception:
With Keith Langley, I developed a formal description of so-called non-Fourier (or second-order) perceptual phenomena. Non-Fourier motion stimuli are signals whose perceived motion is not predicted by typical energy or gradient-based models. These stimuli include the motion of contrast envelopes (e.g., shadows), occlusion boundaries, and some illusory motions due to aliasing. We showed that many non-Fourier stimuli, when viewed as multiplicative combinations of elementary signals, have simple descriptions in the Fourier domain. Our formal description is derived, in part, from the definition of group velocity in wave mechanics.

This modeling effort has since led to several psychophysical studies. Langley, Paul Hibbard (University of St. Andrews), and I reported evidence supporting the hypothesis that contrast envelopes are processed by the visual system after orientation- and scale-specific filtering (in visual cortex). This showed that their perception is not an artifactual consequence of an early nonlinearity (e.g., in photo-transduction), which had been proposed to explain these percepts.

In binocular vision, Langley, Hibbard and I showed that non-Fourier stimuli sometimes produce a percept of transparency, distinct in its properties from transparent percepts that arise from a superposition of two signals. This supported the hypothesis of a non-Fourier channel in stereo depth perception. One consequence of this result, in conjunction with our earlier modeling efforts, was a hypothesis that non-Fourier stimuli are related to distinct physical properties of natural scenes (e.g., multiplicative transparency and occlusions), and therefore non-Fourier processing channels may not subserve all the same visual tasks (such as egomotion) as conventional first-order models. However, Rick Gurnsey, Cindy Potechin (Concordia University) and I found that non-Fourier motion can be used to induce a percept of self-motion (vection). In this study we found a dissociation between motion-aftereffects (non-existent with non-Fourier stimuli) and vection as the relative amounts of Fourier and nonFourier motion energy were varied in the stimuli.

Neural Model of Binocular Disparity:
In visual neuroscience, with Hermann Wagner (Aachen University) and David Heeger (Stanford University), I developed a neural model for the processing of binocular disparity and 3d depth in visual cortex. The goal was to explain the binocular interaction of cells in the primary visual cortex from a functional perspective, in terms of the computation and representation of binocular disparity. The model involves linear neurons and energy neurons, interocular position-shifts and/or phase-shifts, monocular and binocular normalization, pooling in local spatial neighbourhoods, and pooling across orientation- and scale-specific channels. The basic computational framework was derived as a modified form of *phase-correlation.* This work produced a number of theoretical findings, including the fact that conventional energy models are not disparity detectors (as they regularly respond strongly to false matches). This work also makes specific predictions concerning how one might measure the source of disparity selectivity in V1 neurons, and how one might construct disparity detectors from their outputs. I have not continued this work since moving to Xerox PARC, except for my collaboration with Herman Wagner and his students.

fMRI Studies of Human Binocular Depth Perception:
More recently, Ben Backus (University of Pennsylvania), Geoff Boynton (Salk Institute), Andrew Parker (Oxford University), David Heeger and I began to use functional magnetic resonance imaging (fMRI) to study the neural basis of binocular vision and depth perception in humans. It is well known that disparity tuned neurons are widespread in several visual cortical areas, but it is not clear which of these neurons are involved in stereo depth perception, as opposed to binocular fusion or eye movement control for example. Using different 3d surface configurations, and varying amounts of binocularly uncorrelated noise, we identified several brain areas where the brain activity measured with fMRI correlates well with psychophysical tests of binocular acuity limits and disparity upper depth limits. Among these different visual areas, area V3A showed striking sensitivity to stereoscopic stimuli, suggesting that V3A may play a special role in the stereo pathway.

Spinal fMRI:
With David Cadotte, Julien Cohen-Adad, and Micha Livne, I have been working on ways to analysize Spinal fMRI data. The general goal is to provide an atlas which can be used widely within the spinal FMRI community. This entails segmentation of Spinal fMRI images, as well as the identification of key anatomical and functional landmarks, such as vertabrea and spinal nerve rootlets.

Research Contributions

Books and Book Chapters

Fleet, D., Pajdla, T., Schiele, B., and Tuytelaars, T. (Editors), **Proceedings of the 13th European Conference on Computer Vision**, Parts I–VII, Zurich, Switzerland, September 6-12, 2014. Lecture Notes in Computer Science, Volumes 8689–8695, Springer

Fleet, D.J. (2011) Motion models for people tracking. **Guide to Visual Analysis of Humans: Looking at People**, T. Moeslund, A. Hilton, V. Krueger, L. Sigal (editors). Springer, pp. 171–198.

Brubaker, M.A., Sigal, L. and Fleet, D.J. (2009) Video-based people tracking. **Handbook of Ambient Intelligence and Smart Environments**, H. Nakashima, H. Aghajan, J.C. Augusto (editors). Springer Verlag, pp. 57–88.

Fleet, D.J. and Weiss, Y. (2005) Optical flow estimation. **Mathematical Models of Computer Vision: The Handbook**, N. Paragios, Y. Chen, and O. Faugeras (editors), Springer Verlag, Chapter 15, pp. 239-258

Fleet, D.J., Black, M.J. and Nestares, O. (2002) Bayesian inference of visual motion boundaries. **Exploring Artificial Intelligence in the New Millennium**, G. Lakemeyer and B. Nebel (editors), Morgan Kaufmann Press (edited volume containing the invited Distinguished Papers from IJCAI 2001), pp. 139–173

Weiss, Y. and Fleet, D.J. (2001) Velocity likelihoods in biological and machine vision. In **Probabilistic Models of the Brain: Perception and Neural Function**, R.P.N. Rao, B.A. Olshausen and M.S. Lewicki (editors), MIT Press, pp. 81–100.

Fleet, D.J., Heeger, D.J. and Wagner, H. (1997) Neural encoding of binocular disparity. In **Computational and Biological Mechanisms of Visual Coding**, M. Jenkin and L. Harris (editors), Cambridge University Press, pp. 103-130

Fleet, D.J. (1992) **Measurement of Image Velocity**. Kluwer Academic Publishers, Norwell MA

Tsotsos, J.K., Fleet, D.J., Jepson, A.D. (1988) Towards a theory of motion understanding in man and machine. In **Motion Understanding: Robot and Human Vision**, W. Martin and J. Aggarwal (editors), Kluwer Academic Publishers, pp. 353-417

Refereed Journal Publications

Punjani, A., Zhang, H., Rubinstein, J., Brubaker, M.A., and Fleet, D.J. (2018) Algorithmic advances in single particle caryo-EM data processing. *Microscopy and Microanalysis*, 24(S1), 868-869.

Punjani, A., Rubinstein, J., Fleet, D.J. and Brubaker, M.A. (2017) cryoSPARC: Algorithms for rapid unsupervised cryo-EM structure determination. *Nature Methods*, 14 (3): 290-296.

Punjani, A., Brubaker, M.A., and Fleet, D.J. (2016) Building proteins in a day: Efficient 3D molecular reconstruction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(4):706-718, 2017.

Cao, Y., Brubaker, M., Fleet, D.J., and Hertzmann, A. (2015) Efficient optimization for sparse Gaussian process regression. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(12):2415-2427.

Cadotte, A., Cadotte, D.W., Livne, M., Cohen-Adad, J., Fleet, D.J., Mikulis, D. and Fehlings, M.G. (2015) Spinal cord segmentation by one dimensional normalized template matching: A novel, quantitative technique to analyze advanced magnetic resonance imaging data. *PLoS ONE*, 10(10): e0139323 (doi: 10.1371/journal.pone.0139323)

Cadotte, D., Cadotte, A., Cohen-Adad, J., Fleet, D.J., Livne, M., Wilson, J.R., Mikulis, D., Nugaeva, N., and Fehlings, M.G. (2014) Characterizing the location of spinal and vertebral levels in the human cervical spinal cord. *American Journal of Neuroradiology*, December, 2014, A4192.

Norouzi, M., Punjani, A., and Fleet, D.J. (2014) Fast exact search in Hamming space with multi-index hashing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(6):1107–1119

Memisevic, R., Sigal, L., and Fleet, D.J. (2012) Shared kernel information embedding for discriminative inference. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(4):778–790

Livne, M., Sigal, L., Troje, N., and Fleet, D.J. (2012) Human attributes from video-based pose tracking. *Computer Vision and Image Understanding*, 116:648–660

Macrini, D., Dickinson, S., Fleet, D.J., and Siddiqi, K. (2011) Shape categorization using bone graphs. *Computer Vision and Image Understanding*, 115(8):1187–1206

Macrini, D., Dickinson, S., Fleet, D.J., and Siddiqi, K. (2011) Bone Graphs: Medial shape parsing and abstraction. *Computer Vision and Image Understanding*, 115(7):1044–1061

de La Gorce, M., Fleet, D.J., Paragios, N. (2011) Hand tracking with occlusion, lighting and texture. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(9):1793–1805

Wang, J., Fleet, D.J. and Hertzmann, A. (2010) Optimizing walking controllers with uncertain user inputs and environments. *ACM Transactions on Graphics (SIGGRAPH)*, 29(4), Article 73

Brubaker, M.A., Fleet, D.J. and Hertzmann, A. (2010) Physics-based pose tracking with the Anthropomorphic Walker. *International Journal of Computer Vision*, 87(1):140–155

Wang, J., Fleet, D.J. and Hertzmann, A. (2009) Optimizing walking controllers. *ACM Transactions on Graphics (SIGGRAPH Asia)*, 28(5), Article 168

Levinshtein, A., Stere, A., Kutulakos, K., Fleet, D.J., Dickinson, S. and Sidiqqi, K. (2009) TurboPixels: Fast superpixels using geometric flows. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(12):2290–2297

Wang, J., Fleet, D.J. and Hertzmann, A. (2008) Gaussian process dynamical models. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 30(2):283–298

Urtasun, R., Fleet, D.J. and Fua, P. (2006) Motion models for 3D people tracking. *Computer Vision and Image Understanding* 104(2):157–177

Nunes, P., Haines, N., Kuppuswamy, V., Fleet, D.J. and Stewart, B. (2006) Synaptic Vesicle Mobility and Presynaptic F-Actin Are Disrupted in an NSF Allele of Drosophila. *Molecular Biology of the Cell* 17:4709–4719

Jepson, A.D., Fleet, D.J. and El-Maraghi, T. (2003) Robust on-line appearance models for vision tracking. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 25(10):1296–1311

Backus, B., Fleet, D.J., Parker, A.J. and Heeger, D.J. (2001) Human cortical activity correlates with stereoscopic depth perception. *Journal of Neurophysiology* 86:2054-2068

Haussecker, H.W. and Fleet, D.J. (2001) Estimating optical flow with physical models of brightness variation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23(6):661–673

Gurnsey, R., and Fleet, D.J. (2001) Texture space. *Vision Research* 41(3):745–757

Black, M.J. and Fleet, D.J. (2000) Probabilistic detection and tracking of motion boundaries. *International Journal of Computer Vision* 38(3):229–243

Fleet, D.J., Black, M.J., Yacoob, Y., and Jepson, A.D. (2000) Design and use of linear models for image motion analysis. *International Journal of Computer Vision* 36(3):171–193

Lippert, J., Fleet, D.J., and Wagner, H. (2000) Disparity tuning as simulated by a neural net. *Biological Cybernetics* 83(1):61–72

Black, M.J., Fleet, D.J., and Yacoob, Y. (2000) Robustly estimating changes in image appearance. *Computer Vision and Image Understanding* 78(1):8–31

Langley, K., Fleet, D.J., and Hibbard, P.B. (1999) Stereopsis from contrast envelopes. *Vision Research* 39(14):2313–2324.

Langley, K., Fleet, D.J., and Hibbard, P.B. (1998) Linear and nonlinear transparencies in stereopsis. *Proceedings of the Royal Society London B.* 265:1837–1845.

Gurnsey, R., Fleet, D.J. and Potechin, C. (1998) Second-order motions contribute to vection. *Vision Research* 38(18):2801–2816

Fleet, D.J., Wagner, H., and Heeger, D.J. (1996) Neural encoding of binocular disparity: Energy models, position-shifts and phase-shifts. *Vision Research* 36(12):1839–1857

Langley, K., Fleet, D.J., and Hibbard, P.B. (1996) Linear filtering precedes nonlinear processing in early vision. *Current Biology* 6(7):891–896

Fleet, D.J. and Langley, K. (1995) Recursive filters for optical flow. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(1):61–67

Fleet, D.J. and Langley, K. (1994) Computational analysis of non-Fourier motion. *Vision Research*, 34(22):3057–3079

Barron, J.L., Fleet, D.J., and Beauchemin, S.S. (1994) Performance of optical flow techniques. *International Journal of Computer Vision*, 12(1):43–77

Fleet, D.J. and Jepson, A.D. (1993) Stability of phase information. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(12):1253–1268

Jepson, A.D. and Fleet, D.J. (1991) Phase singularities in scale-space. *Image and Vision Computing Journal*, 9(5):338–343

Fleet, D.J., Jepson, A.D., and Jenkin, M. (1991) Phase-based disparity measurement. *Computer Vision, Graphics and Image Processing: Image Understanding*, 53(2):198–210

Fleet, D.J. and Jepson, A.D. (1990) Computation of component image velocity from local phase information. *International Journal of Computer Vision*, 5(1):77–104

Fleet, D.J. and Jepson, A.D. (1989) Hierarchical construction of orientation and velocity selective filters. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(3):315–325

Fleet, D.J., and Jepson, A.D. (1985) Spatiotemporal inseparability in early vision: centre-surround models and velocity selectivity. *Computational Intelligence*, 1(3):89–102

Fleet, D.J., Hallett, P.E., and Jepson, A.D. (1985) Spatiotemporal inseparability in early visual processing. *Biological Cybernetics*, 52(2):153–164

Refereed Conferences/Workshops: Papers in Proceedings

ICCV: *IEEE International Conference on Computer Vision*, IEEE Press
CVPR: *IEEE Conference on Computer Vision and Pattern Recognition*, IEEE Press
ICIP: *IEEE International Conference on Image Processing*, IEEE Press
ECCV: *European Conference on Computer Vision*, Springer Verlag
BMVC: *British Machine Vision Confeence*
NIPS: *Advances in Neural Information Processing Systems*, MIT Press
ICML: *International Conference on Machine Learning*, ACM Press
ICLR: *International Conference on Learning Representations* (online)

Faghri, F., Fleet, D.J., Kiros, J.R., and Fidler, S.: VSE++: Improved visual-semantic embeddings with hard negatives. *BMVC*, Newcastle, September 2018 (14 pages)

Livne, M., Sigal, L., Brubaker, M. and Fleet, D.J.: Walking on thin air: Environment-free physics-based markerless motion capture. *Candaian Conference on Robot and Computer Vision (CRV)*, May 2018

Asgarian, A., Ashraf, A.B., Fleet, D.J. and Taati, B.: Subspace selection to suppress confounding source domain information in AAM transfer learning. *Int. J. Conf. Biometrics (IJCB)*, Oct 2017

Sabour, S., Cao, Y., Faghri, F. and Fleet, D.J.: Adversarial manipulation of deep representations. *ICLR*, Puerto Rico, May 2016 (9 pages)

Norouzi, M., Collins, M.D., Johnson, M., Fleet, D.J. and Kohli, P.: Efficient non-greedy optimization of decision trees. *NIPS*, Montreal, December 2015 (9 pages)

Brubaker, M.A., Punjani, A. and Fleet, D.J.: Building proteins in a day: Efficient 3D molecular reconstruction *CVPR*, Boston, June 2015 (10 pages)

Pons-Moll, G., Fleet, D.J., and Rosenhahn, B.: Posebits for human pose estimation. *CVPR*, Columbus, June 2014 (9 pages)

Cao, Y., Brubaker, M., Fleet, D.J., and Hertzmann, A.: Efficient optimization for sparse Gaussian process regression. *NIPS*, Lake Tahoe, December 2013 (9 pages)

Norouzi, M. and Fleet, D.J.: Cartesian k-means. *CVPR*, Portland, June 2013 (8 pages)

Norouzi, M., Fleet, D.J., and Salakhutdinov, R.: Hamming space metric learning. *NIPS*, Lake Tahoe, December 2012 (9 pages)

Norouzi, M., Punjani, A. and Fleet, D.J.: Fast search in Hamming space with multi-index hashing. *CVPR*, Providence, June 2012 (8 pages)

Norouzi, M., and Fleet, D.J.: Minimal loss hashing for compact binary codes. *ICML*, Bellevue, June 2011 (8 pages)

Manfredotti, C., Fleet, D.J., Hamilton, H.J. and Zilles, S.: Simultaneous tracking and activity recognition. *Int. Conf. Tools with A.I.* Boca Raton, November 2011

Sigal, L., Fleet, D.J., Troje, N. and Livne, M.: Human attributes from 3D pose tracking. *ECCV*, Crete, September 2010 (14 pages)

Taylor, G., Sigal, L., Fleet, D.J. and Hinton, G.: Dynamical binary latent variable models for 3D pose tracking. *CVPR*, San Francisco, June 2010 (8 pages)

Manfredotti, C., Fleet, D.J., and Messina, E.: Relations to improve multi-target tracking in an activity recognition system. *Int. Conf. on Imaging for Crime Detection and Prevention*, London, December 2009

Brubaker, M., Sigal, L. and Fleet, D.J.: Estimating contact dynamics. *ICCV*, Kyoto, October 2009 (8 pages)

Estrada, F., Fleet, D.J. and Jepson, A.D.: Stochastic image denoising. *BMVC*, London, September 2009 (11 pages) **[Best Paper Award]**

Darby, J., Li, B., Costens, N., Fleet, D.J. and Lawrence, N.: Backing Off: Hierarchical decomposition of activity for 3D novel pose recovery. *BMVC*, London, September 2009 (11 pages)

Sigal, L., Memsevic, R., and Fleet, D.J.: Shared kernel information embedding for discriminative inference. *CVPR*, Miami, June 2009 (8 pages)

Urtasun, R., Fleet, D.J., Geiger, A., Popovic, J., Darrell, T. and Lawrence, N.: Topologically constrained latent variable models. *ICML*, Helsinki, June 2008, pp. 96-103

de La Gorce, M., Paragios, N., and Fleet, D.J.: Model-based hand tracking with texture, shading and self-occlusions. *CVPR*, Anchorage, June 2008 (8 pages)

Brubaker, M. and Fleet, D.J.: The Kneed Walker for human pose tracking. *CVPR*, Anchorage, June 2008 (8 pages)

Urtasun, R., Fleet, D.J. and Lawrence, N.: Modeling human locomotion with topologically constrained latent variable models. *Workshop on Human Motion*, Rio de Janeiro, October 2007, pp. 104-118 (Lecture Notes in Computer Science, v. 4814)

Hyndman, M., Jepson, A.D., Fleet, D.J.: Higher-order autoregressive models for dynamic textures. *BMVC*, Coventry, September 2007 (11 pages)

Wang, J.M., Fleet, D.J., and Hertzmann, A.: Multifactor Gaussian Process models for style-content separation. *ICML*, Corvallis, Oregon, June 2007, pp. 975-982

Brubaker, M., Fleet, D.J., and Hertzmann, A.: Physics-based people tracking with simplified lower-body dynamics. *CVPR*, Minneapolis, June 2007 (8 pages)

Urtasun, R., Fleet, D.J., and Fua, P.: 3D people tracking with Gaussian process dynamical models. *CVPR*, New York, June 2006, Vol. 1, pp. 238-245

Wang, J.M., Fleet, D.J., and Hertzmann, A.: Gaussian Process dynamical models. *NIPS*, Vancouver, December 2006, pp. 1441-1448

Urtasun, R., Fleet, D.J., Hertzmann, A. and Fua, P.: Priors for people tracking from small training sets. *ICCV*, Bejing, October 2005, Vol. I, pp. 403-410

Urtasun, R., Fleet, D.J. and Fua, P.: Monocular 3D tracking of the golf swing. *CVPR*, San Diego, June 2005, Vol. II, pp. 932-938

Marinak, D., Dudek, G. and Fleet, D.J.: Learning sensor network topology through Monte Carlo expectation maximization. *Int. Conf. Robotics and Automation*, Barcelona, April 2005, pp. 4581-4587.

Saund, E., Fleet, D.J., Larner, D. and Mahoney, J.V.: Perceptually-supported image editing of text and graphics. *ACM Symposium on User Interface Software and Technology (UIST '03)*, Vancouver, November 2003, *ACM CHI Letters* 5(2):183-192 **[Best Paper Award]**

Nestares, O. and Fleet, D.J.: Error-in-variables likelihood functions for motion estimation. *IEEE International Conference on Image Processing*, Barcelona, September 2003, vol. III, pp. 77-80

Poon, E. and Fleet, D.J.: Hybrid Monte Carlo filtering: Edge-based people tracking. *IEEE Workshop on Motion and Video Computing*, Orlando, December 2002, pp. 151–158

Jepson, A.D., Fleet, D.J. and Black, M.J.: A layered motion representation with occlusion and compact spatial support. *ECCV*, Copenhagen, May 2002, Vol. I, pp. 692–706

Bhotika, R., Fleet, D.J., and Kutulakos, K.: A probabilistic theory of occupancy and emptiness. *ECCV*, Copenhagen, May 2002, Vol. III, pp. 112–130

Jepson, A.D., Fleet, D.J. and El-Maraghi, T.: Robust on-line appearance models for vision tracking. *CVPR*, Kauai, December 2001, Vol. I, pp. 415–422 **[Best Paper Runner-Up Award]**

Nestares, O. and Fleet, D.J.: Probabilistic tracking of motion boundaries with spatiotemporal predictions. *CVPR*, Kauai, December 2001, Vol. II. pp. 358–365

Saund, E., Mahoney, J., Larner, D., and Fleet, D.J.: Perceptual organization as a foundation for graphics recognition. *IAPR Int. Workshop on Graphics Recognition*, Kingston, Canada, Sept. 2001, pp. 175–179

Ormoneit, D., Lemieux, C., and Fleet, D.J.: Lattice particle filters. *Conference on Uncertainty in Artificial Intelligence (UAI)*, Seattle, August 2001, Morgan Kaufmann Press, pp. 395–402

Choo, K. and Fleet, D.J.: Tracking people using hybrid Monte Carlo. *ICCV*, Vancouver, July 2001, Vol II, pp. 321-328

Sidenbladh, H., Black, M.J., and Fleet, D.J.: Stochastic tracking of 3D human figures using 2d image motion. *ECCV*, Dublin, June 2000, Vol. II, pp. 702–718 **[Koenderink Prize]**

Haussecker, H. and Fleet, D.J.: Computing optical flow with physical models of brightness variation. *CVPR*, Hilton Head, South Carolina, June 2000, Vol. II, pp. 760-767

Nestares, O., Fleet, D.J., and Heeger, D.J.: Likelihood functions and confidence bounds for total-least-squares problems. *CVPR*, Hilton Head, South Carolina, June 2000, Vol. I, pp. 523–530

Ma, B., Ellis, R.E. and Fleet, D.J.: Spotlights: A robust method for surface-based registration in orthopedic surgery. *Int. Conf. Medical Image Comp. and Computer-Assisted Intervention*, Cambridge, UK, Springer, September 1999, pp. 936-944

Black, M.J. and Fleet, D.J.: Probabilistic detection and tracking of motion discontinuities. *ICCV*, Corfu, Greece, September 1999, pp. 551–558 **[Marr Prize Honorable Mention]**

Fleet, D.J., Black, M.J. and Jepson, A.D.: Motion feature detection using steerable flow fields. *CVPR*, Santa Barbara, June 1998, pp. 274-281

Black, M.J., Fleet, D.J., and Yacoob, Y.: A framework for modeling appearance change in image sequences. *ICCV*, Mumbai, India, January 1998, pp. 660-667

Fleet, D.J. and Heeger, D.J.: Embedding invisible information in color images. *ICIP*, Santa Barbara, October 1997, Vol. I, pp. 532-535

Black, M.J., Yacoob, Y., Jepson, A.D., and Fleet, D.J.: Learning parameterized models of image motion. *CVPR*, Puerto Rico, June 1997, pp. 561-567

Black, M.J., Yacoob, Y., and Fleet, D.J.: Modeling appearance change in image sequences. *Third Int. Workshop on Visual Form*, Capri, Italy, May 1997, C. Arcelli, L.P. Cordella, and G.S. di Baja, Eds., World Scientific Pub., pp. 11–20.

Ellis, R., Fleet, D.J., Bryant, T., Rudan, J., and Fenton, P.: A method for evaluating CT-based surgical registration. *Proc. Int. Conf. Medical Robotics and Computer Assisted Surgery*, Grenoble, March 1997 pp. 141-150

Clifford, C., Langley, K. and Fleet, D.J.: Centre-frequency adaptive IIR temporal filters for phase-based image velocity estimation. *IEE Int. Conf. Image Proc. and Applic.*, Edinburgh, July 1995, pp. 173-178

Haglund, L. and Fleet, D.J.: Stable estimation of image orientation. *ICIP*, Austin, November 1994, Vol. III, pp. 68-72

Fleet, D.J.: Disparity from local weighted phase-correlation. *IEEE International Conf. Systems, Man, & Cybernetics*, San Antonio, October 1994, pp. 48-56

Barron, J.L., Beauchemin, S.S., and Fleet, D.J.: On optical flow. *6th Int. Conf. AI and Inform.-Control Systems of Robots*, Bratislava, Slovakia, Sept. 1994, World Scientific Publ., pp. 3-14

Fleet, D.J. and Langley, K.: Toward real-time optical flow. *Vision Interface*, Toronto, May 1993, pp. 116-124

Langley, K. and Fleet, D.J.: Recursive filters for phase-based optical flow. *Isreali Conf. on Vision and AI*, Ramat Gan, Isreal, December 1992, pp. 255-264

Langley, K. and Fleet, D.J.: Multiple binaural time delay estimation. *ESCA Workshop on Speech Processing in Adverse Conditions*, Cannes, November 1992, pp. 159-162 (ISSN 1018-4554)

Langley, K., Fleet, D.J., and Atherton, T.J.: On transparent motion computation. *BMVC*, Leeds, September 1992, pp. 245-255

Barron, J.L., Fleet, D.J., Beauchemin, S.S., and Burkitt, T.: Performance of optical flow techniques. *CVPR*, Champaign, June 1992, pp. 236-242

Langley, K., Fleet, D.J., and Atherton, T.J.: Multiple motions from instantaneous frequency. *CVPR*, Champaign, June 1992, pp. 846-849

Fleet, D.J. and Jepson, A.D.: Stability of phase information. *IEEE Workshop on Visual Motion*, Princeton, October 1991, pp. 52-60

Jepson, A.D., and Fleet, D.J.: Scale-space singularities. *ECCV*, Antibes, April 1990, Springer-Verlag, pp. 50-55

Fleet, D.J. and Jepson, A.D.: Computation of normal velocity from local phase information, *CVPR*, San Diego, June 1989, pp. 379-386

Fleet, D.J. and Jepson, A.D.: Computation of normal velocity from local phase information, *OSA Image Understanding and Machine Vision Meeting*, Cape Cod, June 1989, pp. 58-61

Mohnhaupt, M. and Fleet, D.J.: Raum-zeitliche Filter fur eine top-down Steuerung der Bewegungs-analyse, *German AI Workshop*, August 1988, Springer-Verlag, pp. 296-305

Fleet, D.J. and Jepson, A.D.: Velocity extraction without form interpretation, *IEEE Workshop on Computer Vision*, Bellaire, October 1985, pp. 179-185

Tsotsos, J.K., Jepson, A.D., and Fleet, D.J.: Motion understanding meets early vision: an introduction, *IEEE Applications of AI*, Denver, December 1984, pp. 239-244

Conference Video Proceedings

Urtasun, R., Fleet, D.J., and Fua, P.: 3D people tracking with Gaussian Process dynamical models. *CVPR*, New York, June 2006

Urtasun, R., Fleet, D.J., Hertzmann, A. and Fua, P.: Priors for people tracking from small training sets. *ICCV*, Bejing, October 2005, Vol. I, pp. 403-410

Urtasun, R., Fleet, D.J. and Fua, P.: Monocular 3D tracking of the golf swing. *CVPR*, San Diego, June 2005, Vol. II, pp. 932-938

Lightly Refereed Workshops/Symposia: Papers in Proceedings

Mahdisoltani, F., Memisevic, R., and Fleet, D.J.: Hierarchical video understanding. *BMVC Video Workshop*, Munich, 2018

Cao, Y. and Fleet, D.J.: Transductive log opinion pool of Gaussian Process experts. *NIPS Workshop on Nonparametric Methods for Large Scale Representation Learning*, Montreal, 2015.

Brubaker, M.A., Punjani, A. and Fleet, D.J.: Efficient 3D macromolecular reconstruction with electron cryo-microscopy. *CVPR BioImage Computing Workshop (BIC)*, Boston, 2015.

Cao, Y. and Fleet, D.J.: Generalized product of experts for automatic and principled fusion of Gaussian process predictions. *NIPS Workshop on Automating the Learning Pipeline*. Montreal, 2014.

Manfredotti, C., Messina, E. and Fleet, DJ.: Relations as context to improve multi-target tracking and activity recognition. Proc. of 1st Workshop Log-IC09, Potsdam, September 2009. (CEUR Workshop Proceedings, ISSN 1613-0073)

Urtasun, R., Fleet, D.J., Darrell, T. and Lawrence, N.: Modeling human locomotion with topologically constrained latent variable models. *NIPS Workshop on Topology Learning*, Vancouver, December 2007

Brubaker, M.A., Fleet, D.J. and Hertzmann A.: Physics-based human pose tracking. *NIPS Workhop: Evaluation of Articulated Human Motion and Pose Estimation*, Vancouver, December 2006

Saund, E., Fleet, D., Mahoney, J., and Larner, D.: Rough and degraded document interpretation by perceptual organization. *Symposium on Document Image Understanding Technology*, Maryland, April 2003

Saund, E., Mahoney, J., Fleet, D.J., Larner, D. and Lank, E.: Perceptual organization as a foundation for intelligent sketch editing. *AAAI Spring Symp. on Sketch Understanding*, March 2002, Stanford University, pp. 118–125

Ormoneit, D., Sidenbladh, H., Black, M., Hastie, T., and Fleet, D.: Learning and tracking human motion using functional analysis. *Proc. IEEE Workshop on Human Modeling, Analysis and Synthesis*, Hilton Head, SC, June 2000.

van der Willigen, R.F., Lippert, J., Fleet, D.J., Wagner, H.: Binocular information processing in the owl. In: *GISI'97*, M. Jarke and K. Pasedach (editors), Springer Verlag, Berlin, Sept, 1997

Fleet, D.J. and Jepson, A.D.: The extraction of orientation and 2-d velocity through hierarchical processing, *SPIE Conf. on Image Coding (vol. 594)*, Cannes, December 1985, pp. 10-20

Conferences with Published Abstracts

ARVO: Association for Research in Vision and Ophthalmology Annual Meeting
ECVP: European Conference on Visual Perception
OSA: Optical Society of America Annual Meeting

Cadotte, D., Cadotte, A., Cohen-Adad, J., Fleet, D.J., Livne, M., Mikulis, D., and Fehlings, M.G. Resolving the anatomic variability of the human cervical spinal cord: A solution to facilitate advanced neural imaging. Annual Meeting for the *International Society for Magnetic Resonance in Medicine* (ISMRM), Milan. May, 2014.

Cohen-Adad, J., Cadotte, D., Cadotte, A., Fleet, D.J., Livne, M., and Fehlings, M.G. Spinal Cord Toolbox: an open-source framework for processing spinal cord MRI data. Annual Meeting for the *Organization of Human Brain Mapping* (OHBM), Hamburg, June 2014.

Saund, E., Fleet, D., Larner, D, and Mahoney, J.: Perceptually Supported Image Editing of Text and Graphics. *ACM Transactions on Graphics: Proceedings of ACM SIGGRAPH 2004*, V. 23, No. 3, August, 2004.

Backus, B., Fleet, D.J., Parker, A.J., Heeger, D.J.: Cortical activity correlates with stereoscopic depth perception. *Conf. on Functional Brain Imaging in Vision*, Fort Lauderdale, Florida, May 2000

Weiss, Y. and Fleet, D.J.: Velocity likelihoods from generative models. *ARVO*. Fort Lauderdale, May 2000 (see *Invest. Ophth. and Vis. Res.*, vol. 41)

Backus, B.T., Fleet, D.J., Tyler, C.W., and Heeger, D.J.: FMRI correlates of stereo depth discrimination. *Society for Neuroscience Annual Meeting*, October 1999

Backus, B.T., Fleet, D.J. and Heeger, D.J.: Differential fMRI response to absolute and relative disparity in area V3/V4. *ARVO*, Fort Lauderdale, May 1999 (see *Invest. Ophth. and Vis. Res.*, vol. 40)

Langley, K., Fleet, D.J., and Hibbard, P.: A comparison of first-order and second-order transparency thresholds in stereopsis. *ECVP*, Oxford, August 1998 (see *Perception*, vol. 27 suppl., p. 102)

Fleet, D.J.: Binocular energy models and the encoding of binocular disparity. *OSA*, Long Beach, October 1997 (see *Optics and Photonics News*, vol. 8 supplement, p. 92)

Potechin, C., Gurnsey, R. and Fleet, D.J.: Vection and motion after-effects with nonFourier stimuli. *ARVO*, Fort Lauderdale, May 1997 (see *Invest. Ophth. and Vis. Res.*, vol. 38)

Khan, R., Boynton, G., Fleet, D.J., Heeger, D.J.: Neural basis of stereo depth perception measured with fMRI. *ARVO*, Fort Lauderdale, May 1997 (see *Invest. Ophth. and Vis. Res.*, vol. 38)

Gurnsey, R. and Fleet, D.J.: A multidimensional scaling study of texture perception. *ARVO*, Fort Lauderdale, May 1996 (see *Invest. Ophth. and Vis. Res.*, vol. 37)

Langley, K. and Fleet, D.J.: Combined multiplicative/additive model of plaid transparency. *ARVO*, Fort Lauderdale, May 1996 (see *Invest. Ophth. and Vis. Res.*, vol. 37)

Potechin, C., Gurnsey, R. and Fleet, D.J.: Vection, motion aftereffects and first and second-order motion signals. *ARVO*, Fort Lauderdale, May 1996 (see *Invest. Ophth. and Vis. Res.*, vol. 37)

Langley, K., Fleet, D.J. and Hibbard, P.: Linearity of early visual analysis. *Proc. Ann. Meeting of Applied Vision Assoc.* Reading, April 1996

Hibbard, P., Langley, K., and Fleet, D.J.: Transparent asymmetry in stereopsis. *ECVP*, Tuebingen, August 1995

Langley, K., Fleet, D.J., and Hibbard, P.: Scale dependence of transparency in RDS. *ECVP*, Tuebingen, August 1995

Fleet, D.J., Heeger, D.J. and Wagner, H.: Computational model of binocular vision. *ARVO*, Fort Lauderdale, May 1995 (see *Invest. Ophth. and Vis. Res.*, vol. 36)

Langley, K. and Fleet, D.J.: A model for coherent and multiplicatively transparent plaids. *ARVO*, Fort Lauderdale, May 1995 (see *Invest. Ophth. and Vis. Res.*, vol. 36)

Fleet, D.J. and Langley, K.: Non-Fourier channels in stereopsis and motion. *ECVP*, Eindhoven, September 1994 (see *Perception*, vol. 23 supplement, p. 83)

Hibbard, P., Langley, K, and Fleet, D.J.: Computational model for stereoscopic slant using orientational differences from Fourier and non-Fourier mechanisms. *ECVP*, Eindhoven, September 1994 (see *Perception*, vol. 23 supplement, p. 35)

Fleet, D.J. and Langley, K.: Computational analysis of non-Fourier motion. *ARVO*, Sarasota, May 1994 (see *Invest. Ophth. and Vis. Res.*, vol. 35, no. 4, p. 1267)

Langley, K. and Fleet, D.J.: Analysis of subsampled image motion. *ARVO*, Sarasota, May 1994 (see *Invest. Ophth. and Vis. Res.*, vol. 35, no. 4, p. 1406)

Langley, K. and Fleet, D.J.: A post-filtering logarithmic transformation applied to multiplicative transparency and motion discontinuities. *Proc. Applied Vision Assoc.* Bristol, April, 1994 (see *Opthal. Physiol. Optics* 14, p. 441, 1994)

Langley, K. and Fleet, D.J.: Phase and energy velocity: An alternative to Fourier/non-Fourier motion mechanisms. *ECVP*, Edinburgh, August 1993 (see *Perception*, vol. 22 supplement, p. 83-84)

Langley, K. and Fleet, D.J.: On D_{max} in Optic Flow. *ECVP*, Pisa, September 1992 (see *Perception*, vol. 21 supplement, p. 41)

Langley, K. and Fleet, D.J.: Using group and phase velocity to explain coherent and transparent motion. *Proc. Ann. Meeting of Applied Vision Assoc.* Manchester, 1992, pp. 1-2 (see *Opthal. Physiol. Optics*)

Fleet, D.J. and Langley, K.: Computing orientational disparities from phase gradients. *U.K. SERC III Workshop*, Glasgow, September 1991

Langley, K. and Fleet, D.J.: Causal velocity mechanisms applied to transparent surfaces. *ECVP*, Moscow, August 1991 (see *Perception*, vol. 20 supplement, p. 79)

Fleet, D.J. and Jepson, A.D.: Measurement of orientation and image velocity through hierarchical processing. *OSA*, Washington DC, October 1985 (see *J. Optical Society of America A*, vol. 2, No. 13, p. 19)

Technical Reports

RPL: Robotics and Perception Laboratory, Queen's University
RBCV: Research in Biological and Computational Vision, University of Toronto
FBI-HH: Department of Computer Science, University of Hamburg

Bhotika, R., Fleet, D.J. and Kutulakos, K.: Probabilistic theory of occupancy and emptiness. University of Rochester, TR 753, December 2001.

Fleet, D.J., Wagner, H. and Heeger, D.J.: Neural encoding of binocular disparity: Energy models, position shifts and phase shifts. RPL-TR-9510

Darrell, T. and Fleet, D.J.: Second-order method for occlusion relationships in motion layers. MIT Media Lab TR: 314, 1995

Fleet, D.J. and Langley, K.: Computational analysis of non-Fourier motion. RPL-TR-9309

Barron, J.L., Fleet, D.J., and Beauchemin, S.S.: Performance of optical flow techniques. RPL-TR-9207 – Revised, July 1993

Fleet, D.J. and Langley, K.: Recursive filters for optical flow. RPL-TR-9308

Barron, J.L., Fleet, D.J., and Beauchemin, S.S.: Performance of optical flow techniques. RPL-TR-9207

Fleet, D.J. and Jepson, A.D.: Stability of Phase Information. RPL-TR-9105

Fleet, D.J., Jepson, A.D. and Jenkin, M.: Phase-based disparity measurement. RBCV-TR-89-29

Fleet, D.J. and Jepson, A.D.: Computation of normal velocity from local phase information. RBCV-TR-89-27

Fleet, D.J.: Implementation of velocity-tuned filters and image encoding. FBI-HH-M-159/88

Fleet, D.J. and Jepson, A.D.: On the hierarchical construction of orientation and velocity selective filters. RBCV-TR-85-8

Fleet, D.J.: The early processing of spatio-temporal visual information. RBCV-TR-84-7

Fleet, D.J. and Jepson, A.D.: A cascaded filter approach to the construction of velocity selective mechanisms. RBCV-TR-84-6

Fleet, D.J., Jepson, A.D., and Hallett, P.E.: A spatio-temporal model for early visual processing. RBCV-TR-84-1

Patents

Patent Application US 15/596,173: ”Methods and systems for image alignment of at least one image to a model”, Ali Punjani, Marcus Brubaker and David Fleet (Filed: May 16, 2017).

Patent US 20,170,103,161: ”Methods and system for 3D structure estimation”, Marcus Brubaker, Ali Punjani and David Fleet (Filed: October 13, 2016; Awarded April 13, 2017).

Patent EP 1,361,544: ”System and method for editing electronic images” Eric Saund, Tom Moran, Dan Larner, James Mahoney, David Fleet and Chris Popat (Filed: Mar. 19, 2003, Awarded: Oct. 6, 2010).

Patent Application EP 1,349,054: ”Method and system for interpreting imprecise object selection paths” Eric Saund, Ed Lank, David Fleet, James Mahoney, Dan Larner, and Tom Moran (Filed: Mar. 19, 2003).

Patent US 7,093,202: ”Method and system for interpreting imprecise object selection paths”, Eric Saund, Ed Lank, David Fleet, James Mahoney, Dan Larner, and Tom Moran (Filed: Mar. 22, 2002, Awarded: Aug. 15, 2006).

Patent US 7,058,205: ”Robust on-line appearance models for visual tracking”, Allan D. Jepson, David J. Fleet, and Thomas F. El-Maraghi (Filed: December 7, 2001, Awarded: June 6, 2006)

Patent EP 1,318,477: ”Robust on-line appearance models for visual tracking”, Allan D. Jepson, David J. Fleet, and Thomas F. El-Maraghi (Filed: Dec. 6, 2002, Awarded: Mar. 22, 2006)

Patent US 6,954,544: ”Visual motion analysis method for detecting arbitrary numbers of moving objects in image sequences”, Allan D. Jepson, David J. Fleet, and Michael J. Black (Filed: May 22, 2002, Awarded: Oct. 11, 2005).

Patent US 6,903,751: ”System and method for editing electronic images” Eric Saund, Tom Moran, Dan Larner, James Mahoney, David Fleet, Chris Popat (Filed: Mar. 22, 2002, Awarded: June 7, 2005).

Patent EP 0,912,042: ”Method for Embedding Signals in a Color Image”, David J. Fleet, David J. Heeger, Todd A. Cass, David L. Hecht (Filed: Oct. 16, 1998; Awarded: Aug. 4, 2004)

Patent US 5,949,055: ”Automatic Geometric Image Transformations Using Embedded Signals”, David J. Fleet, David J. Heeger, Todd A. Cass, David L. Hecht (Filed: Oct. 23, 1997; Awarded: Sept. 7, 1999)

Invited Conference/Workshop/Keynote Presentations

Punjani, A., Fleet, D.J. & Brubaker, M.A.: Advances in cryoSPARC for Cryo-EM. Gordon Conference on 3D Electron Microscopy, Switzerland, June 2017

Fleet, D.J.: Multi-view reconstruction using Cryo-EM. *German Conference on Patter Recognition*, [Keynote Presentation], September 2016

Punjani, A., Rubinstein, J., Fleet, D.J. & Brubaker, M.A.: Algorithms for reducing the computational burden of Cryo-EM. *Gordon Conference on 3D Electron Microscopy*, Hong Kong, June 2016.

Punjani, A., Rubinstein, J., Fleet, D.J. & Brubaker, M.A.: Outlier rejection methods and branch-and-bound image alignment in cryoSPARC. *Gordon Conference on 3D Electron Microscopy*, Hong Kong, June 2016 [Poster].

Fleet, D.J.: Estimation of human pose and interaction using physics-based models. *At the Intersection of Vision, Graphics, Learning and Sensing*, Cambridge, May 2012

Fleet, D.J.: Hashing for large-scale image retrieval. *International Workshop on Computer Vision*, Sicily, May 2012

Fleet, D.J.: Tracking and understanding human motion. *British Machine Vision Conference*, [Keynote Presentation], September 2011

Fleet, D.J.: Physics-Based models for human motion analysis. *ECCV Workshop on Human Motion: Understanding, Modeling, Capture and Animation*, [Keynote Presentation], September 2010

Fleet, D.J.: Tracking human pose and motion. *IEEE CVPR Workshop on Evaluation of Human Motion and Pose Estimation*, [Keynote Presentation], June 2007

Fleet, D.J.: Motion models for 3D human tracking. *BIRS Workshop on Mathematical Methods in Computer Vision*, Banff, October 2006

Fleet, D.J.: Vision as Bayesian Inference: Looking at people. *APICS Conference on Mathematics and Computer Science* [Keynote Presentation], University of Cape Breton, October 2006

Fleet, D.J.: Gaussian process dynamical models for modeling and tracking human motion. *Workhsop on Learning, Representation and Context for Human Sensing in Video*, New York, June 2006

Fleet, D.J.: Inference of visual motion boundaries. *Early Cognitive Vision Workshop* [Keynote presentation], Isle of Skye, Scotland, May 2004

Fleet, D.J.: Bayesian inference of visual motion boundaries on random fields. *Vision Interface 2003* [Keynote presentation], Halifax, June 2003

Fleet, D.J.: Appearance models for visual tracking. *Workshop on Recent Advances and Future Trends in Computer Vision*, Stanford University, March 2002

Fleet, D.J. and Black, M.J.: Bayesian inference of visual motion boundaries. *IJCAI: International Joint Conference on Artificial Intelligence* [Distinguished Paper Track], Seattle, August 2001

Fleet, D.J.: Bayesian image sequence analysis. *Workshop on Image Sequence Processing for Studying Dynamic Systems*, University of Heidelberg, Germany, Sept. 2000

Fleet, D.J.: Bayesian detection and tracking of motion boundaries. *BASICS: Banff Annual Seminar in Cognitive Science*, May 2000

Fleet, D.J.: Binocular energy models and the encoding of binocular disparity. *Annual Meeting of the Optical Society of America*, Long Beach, October 1997

Fleet, D.J., Black, M.J., and Jepson, A.D.: Learning parameterized models for optical flow. *Workshop on Image Sequence Processing for Studying Dynamic Systems*, Heidelberg, Germany, June 1997

Fleet, D.J.: Extraction and representation of binocular disparity. *International Conference on Visual Coding*, Toronto, June 1995

Fleet, D.J.: Models of binocular interaction and disparity estimation. *Computational Neuroscience of Stereoscopic Depth Perception*, Max-Planck Inst. for Biological Cybernetics, Germany, July 1994

Fleet, D.J. and Jepson, A.D.: Hierarchical construction of velocity-tuned filters, *University of Toronto Symposium on Vision*, Toronto, May 1986

Other Invited Workshop/Symposium Presentations and Panels

Fleet, D.J., Punjani, A., and Zhang, H.: Non-local refinement for Cryo-EM structure determination. *Machine Learning for Cryo-EM*, New York Structural Biology Center, Apr 2018

Fleet, D.J., Punjani, A., and Brubaker, M.: Barnch and Bound Optimization for Cryo-EM. *CIFAR Annual Meeting for Program on Learning in Machines and Brains.* Dec., 2017

Fleet, D.J., Punjani, A., and Brubaker, M.: Fast and accurate Cryo-EM structure determination. *CIFAR Workshop on Deep Learning and Medical Image Analysis*, Amsterdam, October, 2016

Fleet, D.J.: Cartesian k-means. *CIFAR Workshop on Neural Computation and Adaptive Perception*, Montreal, December 2014

Fleet, D.J.: Physics-based models for people tracking. *Bellairs Workshop on Computer Animation*, Barbadoes, February 2011

Fleet, D.J.: Gaussian Process models for human motion. *International Workshop on Recent Trends in Computer Vision*, Kyoto, June 2009

Fleet, D.J.: Gaussian Process models for human motion. *CIAR Workshop on Neural Computation and Adaptive Perception*, Vancouver, December 2007

Fleet, D.J.: People Tracking with Simplified Lower-Body Dynamics. *Workshop on Computer Vision*, CMU, Pittsburgh, April 2007

Fleet, D.J.: Motion models for 3D people tracking. *CIAR Workshop on Neural Computation and Adaptive Perception*, Vancouver, December 2005

Fleet, D.J.: Priors for people tracking from small datasets. *CIAR Workshop on People Tracking*, Toronto, July 2005

Fleet, D.J.: Bayesian inference of visual motion boundaries. *CIAR Workshop: Learning to See*, Vancouver, December 2003

Fleet, D.J.: Appearance model for visual tracking. *Computational Neuroscientists of Upper Canada (CNUC) Meeting*, Fields Institute, Toronto, October 2003

Fleet, D.J.: Structure in motion. *IEEE Workshop on Motion and Video Computing, Panel on Visual Motion Analysis*, Orlando, December 2002

Fleet, D.J., Black, M.J., and Nestares, O.: Probabilistic detection and tracking of motion boundaries. *Annual Interdisciplinary Conference*, Jackson Hole, WY, January 2001

Haussecker, H. and Fleet, D.J.: Computing optical flow with physical models of brightness variation. *Bay Area Vision Meeting*, IBM Almaden, May, 2000

Sidenbladh, H., Black, M.J., and Fleet, D.J.: Stochastic tracking of 3D human figures using 2d image motion. *Bay Area Vision Meeting*, IBM Almaden, May, 2000

Barron, J.L., Fleet, D.J., Beauchemin, S.S., and Burkitt, S.: Performance of optical flow techniques. *IEEE Visual Motion Workshop: Experimental Session*, Princeton, October 1991

Fleet, D.J.: Phase-based optical flow. *IEEE Visual Motion Workshop: Experimental Session*, Princeton, October 1991

Fleet, D.J. and Jepson, A.D.: Computation of normal velocity from local phase information, *Univ. Buffalo Graduate Conf. in Computer Science*, March 1989 (Proceedings in SUNY Buffalo TR)

Fleet, D.J. and Jepson, A.D.: Velocity extraction without form interpretation, *Canadian Institute for Advanced Research Workshop on Vision*, Halifax, March 1986

Invited Talks/Seminars/Colloquia

Teaching Computers to See
TedxToronto, October 2014

Estimating Human Pose and Contact Dynamics
Department of Computer Science Distinguished Talk, University of Manitoba, April 2013
Department of Computer Science, University of Copenhagen, January 2012
Institute for Information Processing, Leibniz University of Hannover, January 2012
Max Plank Institute for Intelligent Systems, Tubingen, January 2012
Department of Applied Mathematics, University of Bonn, December 2009
Department of Applied Mathematics, Ecole Centrale de Paris, December 2009

Physics-Based 3D People Tracking
Department of Science, UOIT, March 2011
Department of Computer Science Colloquim, McGill University, Feburary 2010
Department of Computer Science, University of Aachen, December 2009
Department of Computer Science, University of Darmstadt, December 2009
Department of Applied Mathematics, Ecole Centrale de Paris, May 2008
Honda Research Institute, Mountain View, CA, March 2008
Computer Vision Colloqium, CSAIL, MIT, October 2007
Department of Computer Science, Boston University, October 2007

Motion Models for 3D People Tracking
Department of Computer Science Colloquim, Dartmouth College, April 2009
Computer Vision Distinguished Seminar Series, EECS, University of Central Florida, Febuary 2008
Department of Computer Science, Swiss Federal Inst. of Technology Zurich (ETH), May 2006
Department of Computing and Information Science, University of Pennsylvania, November 2005.
Department of Computer Science, York University, November 2005

Appearance Models for Visual Tracking

Computer and Communication Science, Swiss Federal Inst. of Technology Lausanne (EPFL), June 2004
Department of Psychology, Queen's University, Febuary 2004
IBM Almaden Research Center, San Jose, October 2002
Department of EECS, University of California at Berkeley, February 2002
Fuji-Xerox Research Labs, Palo Alto, February 2002
Department of Computer Science, Vision Seminar, Stanford University, February 2002

Bayesian People Tracking from Image Sequences

Department of Computer Science, University of British Columbia, June 2002
Department of Computer Science, University of Toronto, May 2002
School of Computing, Queen's University, April 2002
School of Computer Science, University of Waterloo, March 2002
School of Engineering, University of California at Santa Cruz, November 2001

Bayesian Detection and Tracking of Motion Boundaries

Los Alamos National Laboratory, Sante Fe, New Mexico, December 2000
Department of Biology, Universitat Aachen, September 2000
Department of Computer Science, York University, August 2000
Xerox Research Center Europe, Cambridge, U.K., June 2000
Department of Computer Science, University of Waterloo, March 2000
Xerox Wilson Research Center, Rochester, NY, March 2000
Department of Computer Science, University of Rochester, March 2000
Department of Computing and Information Science, Queen's University, March 2000
Department of EECS, University of California at Berkeley, February 2000
Broad-Area Colloquium, Department of Computer Science, Stanford University, December 1999

Parameterized Motion Models for Image Sequence Analysis

Department of Computer Science, McGill University, Montreal, October 1998
Department of Computing and Information Science, Queen's University, September 1998
Xerox Palo Alto Research Center, August 1998

Neural Basis of Stereo Depth Perception With fMRI

Institute for Zoology, Aachen University, June 1997

Embedding Invisible Information in Color Images

Department of Computing and Information Science, Queen's University, March 1998
Image Understanding Seminar, Xerox PARC, February 1997

Neural Encoding of Binocular Disparity

Department of Psychology, York University, September 1998
Department of Psychology, Stanford University, November 1996
Neuroscience Seminar Series, Queen's University, November 1995
Smith-Kettlewell Research Institute, San Francisco, July 1995

Computational Analysis of Non-Fourier Motion

Interval Research Corp., Palo Alto, CA, January 1997
Department of EECS, University of California at Berkeley, April 1994
Department of Psychology, Stanford University, Stanford, March 1994
Machine Perception Seminar, Xerox PARC, Palo Alto, February 1994
Centre for Intelligent Machines, McGill University, Montreal, December 1993
Department of Pyschology, University College London, England, October 1993

Max-Planck Institute for Biological Cybernetics, Tubingen, Germany, October 1993
Department of Computer Science, University of Hamburg, Germany, October 1993

Stability of Phase for Signal Matching

Department of Electrical Engineering, University of Linkoping, Sweden, June 1992
Department of Pyschology, University College London, England, May 1992
Department of Computer Science, York University, Downsview, January 1992
Department of Electrical Engineering, Yale University, New Haven, December 1991
Department of Electrical Engineering, Brown University, Providence, December 1991
Siemans Research Centre, Princeton, November 1991
Department of Computer Science, University of British Columbia, Vancouver, June 1991
Centre for Intelligent Machines, McGill University, Montreal, April 1991

Phase-Based Measurement of Binocular Disparity

Computer Science Department, University of Rochester, Rochester, November 1989

Phase-Based Measurement of Image Velocity

Center for Scientific Computation, University of Heidelberg, December 1990
Department of Computer Science, Brown University, Providence, April 1990
Department of Computing and Information Science, Queen's University, February 1990
Centre for Intelligent Machines, McGill University, Montreal, December 1989
Department of Computer Science, University of Western Ontario, London, November 1989
Fraunhofer Research Institute, Karlsruhe, West Germany, May 1989
Department of Computer Science, University of Hamburg, West Germany, April 1989

Measurement of Image Properties

Department of Psychology, Cornell University, Ithica, April 1992
Department of Psychology, Queen's University, Kingston, March 1990

Spatiotemporal Inseparability in Early Visual Processing

Department of Psychology, New York University, New York, June 1985
Sarnoff Research Labs, Princeton, June 1985

Velocity Extraction Using Velocity-Tuned Filters

Computer Science Department, Carnegie-Mellon University, Pittsburg, June 1985