

Curriculum Vitae

Jeffrey Mark Siskind

<http://engineering.purdue.edu/~qobi>
<http://upplysingaoflun.ecn.purdue.edu/~qobi/cccp/>

Tuesday 13 December 2016

Education

Technion, Israel Institute of Technology

B.A. *cum laude* in computer science, November 1979

Massachusetts Institute of Technology

S. M. in computer science, January 1989

Thesis title: The Culprit Pointer Method for Selective Backtracking

Advisor: David A. McAllester

Massachusetts Institute of Technology

Ph.D. in computer science, February 1992

Thesis title: Naive Physics, Event Perception, Lexical Semantics, and Language Acquisition

Advisor: Robert C. Berwick

Academic Employment

Weizmann Institute of Science

July 1974–February 1975

Researcher, Chemical Physics Department

Rehovot, Israel

MIT Lincoln Laboratory

March 1980–August 1982

Staff Member, Digital Integrated Circuits Group and Speech Systems Technology Group

Lexington, Massachusetts

Xerox PARC

June–September 1987

February 1988

June–August 1988

July–September 1989

December 1989–January 1990

Research Intern, System Sciences Laboratory

Palo Alto, California

University of Pennsylvania

February 1992–July 1993

Postdoctoral Fellow, Institute for Research in Cognitive Science

Philadelphia, Pennsylvania

University of Toronto

August 1993–December 1995

Assistant Professor, Department of Computer Science

Toronto, Ontario

Technion, Israel Institute of Technology

February–August 1996

Senior Lecturer, Department of Electrical Engineering

Haifa, Israel

University of Vermont

September 1996–May 1997

Visiting Assistant Professor, Department of Computer Science and Electrical Engineering

Burlington, Vermont

NEC Research Institute
June 1997–December 2001
Research Scientist, Computer Science Division

Princeton, New Jersey

Purdue University
January 2002–present
Associate Professor, School of Electrical and Computer Engineering
September 2003–present
Affiliate, Linguistics Program
September 2009–present
Courtesy Appointment, Department of Computer Science
September 2016–present
Courtesy Appointment, Department of Speech, Language, and Hearing Sciences

West Lafayette, Indiana

Industrial Employment

MetaLogic, Inc.

Cambridge, Massachusetts

April 1982–August 1985
Founder, Chairman, Treasurer, Vice President, and CTO
Founded MetaLogic to commercialize the MACPITTS technology developed at MIT. Acting as CEO for the first year, wrote business plan and negotiated contracts with venture capital firms that successfully raised \$1.5 million. Set up corporate infrastructure, which included: planning, supervising construction, and furnishing of office space; financing, purchase, and installation of computer facilities; establishing relationships with legal and accounting firms; and growing the staff to a total of ten people. As CTO, managed a technical staff of five through the development and delivery of METASYN, a complex digital integrated circuit design automation system comprising over 50,000 lines of LISP code as well as a large nMOS cell library. Was an active participant in, and major contributor to, that design and implementation effort. Supervised successful completion and installation of the product at customers' facilities, resulting in the design of over a dozen fabricated nMOS ICs, all of which proved fully functional at first silicon. Actively participated in the marketing effort which included preparation of sales literature, giving presentations to potential customers, performing in-the-field market research, formulating product strategy, and organizing and giving trade show presentations.

Ontologic, Inc.

Billerica, Massachusetts

January–May 1987
Consultant
Performed consulting services pertaining to the design of a persistent object-oriented database for representing VLSI designs.

International Systems Services Company

New York, New York

June 1987
Consultant
Performed consulting services pertaining to automatic test vector generation from high-level behavioral descriptions of digital systems.

Chestnut Software

Boston, Massachusetts

June–August 1990
Consultant
Performed consulting services pertaining to the design and implementation of a natural-language front end for database applications.

Grants Received

1994

J.M. Siskind, PI, 'Computational Models of Lexical Acquisition and Visual Event Perception,' NSERC operating grant, April 1994–April 1997, \$48,000 (Canadian).

1996

J.M. Siskind, PI, C. Colbourn, co-PI, and R. Snapp, co-PI, 'A Computer Facility to Support Compute-Intensive Research,' University of Vermont Research Advisory Council, November 1996, \$7,497, my share \$7,497.

2001

J.M. Siskind, co-PI, and R.L. Givan, co-PI, 'Learning Program Behaviors using Temporal Event Logic,' Center for Education and Research in Information Assurance and Security, September 2001–August 2002, \$50,000, my share \$30,698.59.

2002

J.M. Siskind, PI, 'A World-Wide Distributed Repository to Support Fine-Grained Sharing of Source Code,' Cisco University Research Program, December 2002, \$83,400.

2003

J.M. Siskind, PI, C.A. Bouman, co-PI, and I. Pollak, co-PI, 'Hierarchical Perceptual Organization with the Center-Surround Algorithm,' NSF award 0329156IIS, September 2003–August 2006, \$500,000, my share \$264,000. Robotics and Computer Vision, the program that funded this proposal, funded only 27 out of the 172 proposals submitted (16%) in 2003.

2004

J.M. Siskind, PI, 'Stochastic Spatio-Temporal Grammars for Event Recognition in Video,' Honda Initiation Grant, November 2004, \$49,878. Only 5 out of the 125 proposals submitted (4%) were funded in this competition.

2005

J.M. Siskind, PI, 'Algorithmic Differentiation of Functional Programs,' NSF award CCF-0438806, May 2005–April 2010, \$420,000. Science of Design, the program that funded this proposal, funded only 10% of the proposals submitted in this competition.

2010

J.M. Siskind, PI, AWS research grant, Amazon.com, March 2010–February 2011, 3000 EC2 credits.

J.M. Siskind, PI, 'Emergent Intelligent Behavior through Integrated Investigation of Embodied Natural Language, Reasoning, Learning, Computer Vision, and Robotic Manipulation,' NRL contract N00173-10-1-G023 (DARPA seedling), July–November 2010, \$50,000.

J.M. Siskind, lead PI (subcontracts to S. Khan and H. Cheng, Sarnoff Corporation, S. Wang, University of South Carolina, and S. Dickinson, University of Toronto), 'Robust and Principled Visual Intelligence,' ARL cooperative agreement W911NF-10-2-0060 (DARPA Mind's Eye program), July 2010–March 2016, \$3,600,000.

2015

J.M. Siskind, PI, 'NRI: Collaborative Research: RobotSLANG: Simultaneous Localization, Mapping, and Language Acquisition,' NSF award 1522954-IIS, September 2015–August 2018, \$650,000. This is in collaboration with Jason J. Corso of the University of Michigan, who received a separate award 1522904-IIS of \$650,000.

Honors

1986

One of 25 annual nationwide recipients of a four year AT&T Ph.D. fellowship, 1986–1990.

1992

Nominated for the ACM Distinguished Dissertation Award.

1993

The George M. Sprowls Award for ‘Outstanding contribution in the field of Electronic Computer and Investigation Research by an EECS Student,’ Department of Electrical Engineering and Computer Science, MIT.

1996

The Louis and Miriam Benjamin Academic Lectureship, Department of Electrical Engineering, Technion.

2013

Best paper award, Fifty First Annual Meeting of the Association for Computational Linguistics.

2015

Paper accepted to senior member track, blue-sky ideas subtrack, Conference on Artificial Intelligence.

Paper published in award winning papers track, Journal of Artificial Intelligence Research.

Refereed Journal Publications

Citation counts in this and other sections are from Google Scholar on Tuesday 1 November 2016. Impact factors are from JCR (2014).

1995

J.M. Siskind, ‘Grounding Language in Perception,’ *Artificial Intelligence Review*, 8(5–6):371–391, 1995, impact factor 2.111, 101 citations.

<http://engineering.purdue.edu/~qobi/papers/aireview1995.pdf>

1996

J.M. Siskind, ‘A Computational Study of Cross-Situational Techniques for Learning Word-to-Meaning Mappings,’ *Cognition*, 61(1–2):39–91, October–November 1996, impact factor 3.479. Also appeared in *Computational Approaches to Language Acquisition*, M.R. Brent, ed., Elsevier, pp. 39–91, 1996, 477 citations.

<http://engineering.purdue.edu/~qobi/papers/cognition1996.pdf>

1997

R. Mann, A.D. Jepson, and J.M. Siskind, ‘The Computational Perception of Scene Dynamics,’ *Computer Vision and Image Understanding (CVIU)*, 65(2):113–128, February 1997, impact factor 1.540, 79 citations.

<http://engineering.purdue.edu/~qobi/papers/cviu1997.pdf>

2000

J.M. Siskind, ‘Grounding the Lexical Semantics of Verbs in Visual Perception Using Force Dynamics and Event Logic,’ *Journal of Artificial Intelligence Research (JAIR)*, 15:31–90, August 2001, impact factor 1.257. Also available as Technical Report 2000-105, NEC Research Institute, Inc., July 2000, 253 citations.

<http://engineering.purdue.edu/~qobi/papers/jair2001.pdf>

M.R. Brent and J.M. Siskind, ‘The Role of Exposure to Isolated Words in Early Vocabulary Development,’ *Cognition*, 81(2):B33–B44, September 2001, impact factor 3.479. Also available as Technical Report 99-107, NEC Research Institute, Inc., July 1999. Revised as Technical Report 2000-067R, NEC Research Institute, Inc., May 2000, 369 citations.

<http://engineering.purdue.edu/~qobi/papers/cognition2001.pdf>

2002

A.P. Fern, R.L. Givan, and J.M. Siskind, ‘Specific-to-General Learning for Temporal Events with Application to Learning Event Definitions from Video,’ *Journal of Artificial Intelligence Research (JAIR)*, 17:379–449, December 2002, impact factor 1.257, 56 citations.

<http://engineering.purdue.edu/~qobi/papers/jair2002.pdf>

2003

S. Wang and J.M. Siskind, ‘Image Segmentation with Ratio Cut,’ *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 25(6):675–690, June 2003, impact factor 5.781, 323 citations.

<http://engineering.purdue.edu/~qobi/papers/pami2003.pdf>

J.M. Siskind, ‘Reconstructing Force-Dynamic Models from Video Sequences,’ *Artificial Intelligence (AIJ)*, 151(1–2):91–154, December 2003, impact factor 3.371, 39 citations.

<http://engineering.purdue.edu/~qobi/papers/aij2003.pdf>

2005

S. Wang, T. Kubota, J.M. Siskind, and J. Wang, ‘Salient Closed Boundary Extraction with Ratio Contour,’ *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 27(4):546–561, April 2005, impact factor 5.781, 155 citations.

<http://engineering.purdue.edu/~qobi/papers/pami2005.pdf>

2006

W. Wang, I. Pollak, T.-S. Wong, C.A. Bouman, M.P. Harper, and J.M. Siskind, ‘Hierarchical Stochastic Image Grammars for Classification and Segmentation,’ *IEEE Transactions on Image Processing (TIP)*, 15(10):3033–3052, October 2006, impact factor 3.625, 40 citations.

<http://engineering.purdue.edu/~qobi/papers/tip2006.pdf>

2007

J.M. Siskind, J. Sherman, Jr., I. Pollak, M.P. Harper, and C.A. Bouman, ‘Spatial Random Tree Grammars for Modeling Hierarchical Structure in Images with Regions of Arbitrary Shape,’ *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 29(9):1504–1519, September 2007, impact factor 5.781, 25 citations.

<http://engineering.purdue.edu/~qobi/papers/pami2007.pdf>

2008

B.A. Pearlmutter and J.M. Siskind, ‘Reverse-Mode AD in a Functional Framework: Lambda the Ultimate Backpropagator,’ *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 30(2):1–36, March 2008, impact factor 0.897, 21 citations.

<http://engineering.purdue.edu/~qobi/papers/toplas2008.pdf>

J.M. Siskind and B.A. Pearlmutter, ‘Nesting Forward-Mode AD in a Functional Framework,’ *Higher-Order and Symbolic Computation (HOSC)*, 21(4):361–376 December 2008, 29 citations.

<http://engineering.purdue.edu/~qobi/papers/hosc2008.pdf>

2012

N. Siddharth, A. Barbu, and J.M. Siskind, ‘Seeing Unseeability to See the Unseeable,’ *Advances in Cognitive Systems (ACS)*, 2:77–94, December 2012.

<http://engineering.purdue.edu/~qobi/papers/acs2012a.pdf>

A. Barbu, N. Siddharth, A. Michaux, and J.M. Siskind, ‘Simultaneous Object Detection, Tracking, and Event Recognition,’ *Advances in Cognitive Systems (ACS)*, 2:203–220, December 2012, 17 citations.

<http://engineering.purdue.edu/~qobi/papers/acs2012b.pdf>

2015

H. Yu, N. Siddharth, A. Barbu, and J.M. Siskind, ‘A Compositional Framework for Grounding Language Inference, Generation, and Acquisition in Video,’ *Journal of Artificial Intelligence Research (JAIR)*, 52:601–713, April 2015, impact factor 1.257, **award winning papers track**, 10 citations.

<http://engineering.purdue.edu/~qobi/papers/jair2015.pdf>

2016

D.P. Barrett, A. Barbu, N. Siddharth, and J.M. Siskind, ‘Saying What You’re Looking For: Linguistics Meets Video Search,’ *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 38(10):2069–2081, October 2016, impact factor 5.781.

<http://engineering.purdue.edu/~qobi/papers/pami2016.pdf>

D.P. Barrett, R. Xu, H. Yu, and J.M. Siskind, ‘Collecting and Annotating the Large Continuous Action Dataset,’ *Machine Vision and Applications (MVAP)*, 27(7):983–995, October 2016, impact factor 1.351.

<http://engineering.purdue.edu/~qobi/papers/mvap2016.pdf>

D.P. Barrett and J.M. Siskind, ‘Action Recognition by Time-Series of Retinotopic Appearance and Motion Features,’ *IEEE Transactions on Circuits and Systems for Video Technology (TCSVT)*, 26(12):2250–2263, December 2016, impact factor 2.615, 1 citation.

<http://engineering.purdue.edu/~qobi/papers/tcsvt2016.pdf>

in press

O. Manzyuk, B.A. Pearlmutter, A.A. Radul, D.R. Rush, and J.M. Siskind, ‘Confusion of Tagged Perturbations in Forward Automatic Differentiation of Higher-Order Functions,’ *Higher-Order and Symbolic Computation (HOSC)*, accepted.

<http://engineering.purdue.edu/~qobi/papers/hosc-accepted.pdf>

Refereed Conference Publications

1982

J.M. Siskind, J.R. Southard, and K.W. Crouch, ‘Generating Custom High Performance VLSI Designs from Succinct Algorithmic Descriptions,’ *Proceedings of the Conference on Advanced Research in VLSI*, pp. 28–40, January 1982, oral 21/82 (25%), 80 citations.

<http://engineering.purdue.edu/~qobi/papers/arv1982.pdf>

1990

J.M. Siskind, ‘Acquiring Core Meanings of Words, Represented as Jackendoff-Style Conceptual Structures, From Correlated Streams of Linguistic and Non-Linguistic Input,’ *Proceedings of the Twenty Eighth Annual Meeting of the Association for Computational Linguistics (ACL)*, pp. 143–156, June 1990, oral, 36 citations.

<http://engineering.purdue.edu/~qobi/papers/acl90.ps.Z>

1993

J.M. Siskind and D.A. McAllester, ‘Nondeterministic Lisp as a Substrate for Constraint Logic Programming,’ *Proceedings of the Eleventh National Conference on Artificial Intelligence (AAAI)*, pp. 133–138, July 1993, oral 135/>500 (<27%), 91 citations.

<http://engineering.purdue.edu/~qobi/papers/aaai93.ps.Z>

1994

J.M. Siskind, ‘Lexical Acquisition in the Presence of Noise and Homonymy,’ *Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI)*, pp. 760–766, July 1994, oral 222/780 (28%), 28 citations.

<http://engineering.purdue.edu/~qobi/papers/aaai94.ps.Z>

1996

J.M. Siskind and Q. Morris, ‘A Maximum-Likelihood Approach to Visual Event Classification,’ *Proceedings of the Fourth European Conference on Computer Vision (ECCV)*, pp. 347–360, April 1996, oral 43/328 (13%), 84 citations.

<http://engineering.purdue.edu/~qobi/papers/eccv96a.ps.Z>

R. Mann, A.D. Jepson, and J.M. Siskind, ‘The Computational Perception of Scene Dynamics,’ *Proceedings of the Fourth European Conference on Computer Vision (ECCV)*, pp. 528–539, April 1996, poster 123/328 (37%), 44 citations.

<http://engineering.purdue.edu/~qobi/papers/eccv96b.ps.Z>

2000

J.M. Siskind, ‘Visual Event Classification via Force Dynamics,’ *Proceedings of the Seventeenth National Conference on Artificial Intelligence (AAAI)*, pp. 149–155, August 2000, oral 143/432 (33%). Also available as Technical Report 2000-007, NEC Research Institute, Inc., January 2000. Revised as Technical Report 2000-047R, NEC Research Institute, Inc., April 2000, 54 citations.

<http://engineering.purdue.edu/~qobi/papers/aaai2000.pdf>

2001

S. Wang and J.M. Siskind, ‘Image Segmentation with Minimum Mean Cut,’ *Proceedings of the Eighth International Conference on Computer Vision (ICCV)*, pp. 517–524, July 2001, poster 205/596 (34%). Also available as Technical Report 2000-169, NEC Research Institute, Inc., December 2000, 71 citations.

<http://engineering.purdue.edu/~qobi/papers/iccv2001.pdf>

2002

A.P. Fern, R.L. Givan, and J.M. Siskind, ‘Specific-to-General Learning for Temporal Events,’ *Proceedings of the Eighteenth National Conference on Artificial Intelligence (AAAI)*, pp. 152–158, July 2002, poster 120/470 (25%), 8 citations.

<http://engineering.purdue.edu/~qobi/papers/aaai2002a.pdf>

A.P. Fern, J.M. Siskind, and R.L. Givan, ‘Learning Temporal, Relational, Force-Dynamic Event Definitions from Video,’ *Proceedings of the Eighteenth National Conference on Artificial Intelligence (AAAI)*, pp. 159–166, July 2002, oral 47/470 (10%), 19 citations.

<http://engineering.purdue.edu/~qobi/papers/aaai2002b.pdf>

2003

I. Pollak, J.M. Siskind, M.P. Harper, and C.A. Bouman, ‘Modeling and Estimation of Spatial Random Trees with Application to Image Classification,’ *Proceedings of the 2003 IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, volume III, pp. 305–308, Hong Kong, 6–12 April 2003, 1290/2377 (54%), 12 citations.

<http://engineering.purdue.edu/~qobi/papers/icassp2003.pdf>

I. Pollak, J.M. Siskind, M.P. Harper, and C.A. Bouman, ‘Parameter Estimation for Spatial Random Trees Using the EM Algorithm,’ *Proceedings of the International Conference on Image Processing (ICIP)*, pp. 257–260, Barcelona, Spain, September 2003, poster 827/1901 (43%), 14 citations.

<http://engineering.purdue.edu/~qobi/papers/icip2003.pdf>

S. Wang, T. Kubota, and J.M. Siskind, ‘Salient Boundary Detection Using Ratio Contour,’ *Proceedings of the Neural Information Processing Systems Conference (NIPS)*, pp. 1571–1578, Vancouver, Canada, 9–11 December 2003, poster 198/717 (27%), 24 citations.

<http://engineering.purdue.edu/~qobi/papers/nips2003.pdf>

2007

B.A. Pearlmutter and J.M. Siskind, ‘Lazy Multivariate Higher-Order Forward-Mode AD,’ *Proceedings of the 34th Annual Symposium on Principles of Programming Languages (POPL)*, pp. 155–160, Nice, France, 17–19 January 2007, oral 36/200 (18%), 20 citations.

<http://engineering.purdue.edu/~qobi/papers/popl2007a.pdf>

J.M. Siskind and B.A. Pearlmutter, ‘First-Class Nonstandard Interpretations by Opening Closures,’ *Proceedings of the 34th Annual Symposium on Principles of Programming Languages (POPL)*, pp. 71–76, Nice, France, 17–19 January 2007, oral 36/200 (18%), 14 citations.

<http://engineering.purdue.edu/~qobi/papers/popl2007b.pdf>

2008

B.A. Pearlmutter and J.M. Siskind, ‘Using Programming Language Theory to Make AD Sound and Efficient,’ *Proceedings of the 5th International Conference on Automatic Differentiation (AD2008)*, pp. 79–90, Bonn, Germany, 11–15 August 2008, oral, 5 citations.

<http://engineering.purdue.edu/~qobi/papers/ad2008.pdf>

2010

A. Barbu, N. Siddharth, and J.M. Siskind, ‘Learning Physically-Instantiated Game Play Through Visual Observation,’ *Proceedings of the 2010 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1879–1886, Anchorage, AK, 3–8 May 2010, oral 856/2062 (42%), 15 citations.

<http://engineering.purdue.edu/~qobi/papers/icra2010.pdf>

2011

N. Siddharth, A. Barbu, and J.M. Siskind, ‘A Visual Language Model for Estimating Object Pose and Structure in a Generative Visual Domain,’ *Proceedings of the 2011 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4854–4860, Shanghai, China, 9–13 May 2011, oral 982/2004 (49%), 5 citations.

<http://engineering.purdue.edu/~qobi/papers/icra2011.pdf>

D. Wingate, N. Goodman, A. Stuhlmüller, and J.M. Siskind, ‘Nonstandard Interpretations of Probabilistic Programs for Efficient Inference,’ *Proceedings of the Neural Information Processing Systems Conference (NIPS)*, Granada, Spain, 12–15 December 2011, poster 306/1400 (22%), 24 citations.

<http://engineering.purdue.edu/~qobi/papers/nips2011.pdf>

2012

Z. Zhang, S. Fidler, J.W. Waggoner, Y. Cao, S. Dickinson, J.M. Siskind, and S. Wang, ‘Superedge grouping for object localization by combining appearance and shape information,’ *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3266–3273, Providence, RI, 16–21 June 2012, poster 465/1933 (24%), 5 citations.

<http://engineering.purdue.edu/~qobi/papers/cvpr2012.pdf>

A. Radul, B.A. Pearlmutter, and J.M. Siskind, ‘AD in Fortran: Implementation via Preprocessor,’ *Proceedings of the 6th International Conference on Automatic Differentiation (AD2012)*, pp. 273–284, Fort Collins, CO, 23–27 July 2012, oral, 3 citations.

<http://engineering.purdue.edu/~qobi/papers/ad2012.pdf>

A. Barbu, A. Bridge, Z. Burchill, D. Coroian, S. Dickinson, S. Fidler, A. Michaux, S. Mussman, N. Siddharth, D. Salvi, L. Schmidt, J. Shangquan, J.M. Siskind, J. Waggoner, S. Wang, J. Wei, Y. Yin, and Z. Zhang, ‘Video in sentences out,’ *Proceedings of the 28th Conference on Uncertainty in Artificial Intelligence (UAI)*, pp. 102–112, Catalina, CA, 15–17 August 2012, oral 24/304 (8%), 68 citations.

<http://engineering.purdue.edu/~qobi/papers/uai2012.pdf>

N. Siddharth, A. Barbu, and J.M. Siskind, ‘Seeing Unseeability to See the Unseeable,’ *First Annual Conference for Advances in Cognitive Systems (ACS)*, pp. 77–94, Palo Alto, CA, 6–8 December 2012, oral 14/38 (37%).

<http://engineering.purdue.edu/~qobi/papers/acs2012a.pdf>

A. Barbu, N. Siddharth, A. Michaux, and J.M. Siskind, ‘Simultaneous Object Detection, Tracking, and Event Recognition,’ *First Annual Conference for Advances in Cognitive Systems (ACS)*, pp. 203–220, Palo Alto, CA, 6–8 December 2012, oral 14/38 (37%).

<http://engineering.purdue.edu/~qobi/papers/acs2012b.pdf>

2013

Y. Cao, D. Barrett, A. Barbu, N. Siddharth, H. Yu, A. Michaux, Y. Lin, S. Dickinson, J.M. Siskind, and S. Wang, ‘Recognize Human Activities from Partially Observed Videos,’ *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2658–2665, Portland, OR, 25–27 June 2013, poster 472/1870 (25.2%), 51 citations.

<http://engineering.purdue.edu/~qobi/papers/cvpr2013.pdf>

H. Yu and J.M. Siskind, ‘Grounded Language Learning from Video Described with Sentences,’ *Proceedings of the Fifty First Annual Meeting of the Association for Computational Linguistics (ACL)*, pp. 56–63, Sofia, Bulgaria, 4–9 August 2013, oral 175/662 (26.4%), 81 citations, **best paper award**.

<http://engineering.purdue.edu/~qobi/papers/acl2013.pdf>

2014

N. Siddharth, A. Barbu, and J.M. Siskind, ‘Seeing What You’re Told: Sentence-Guided Activity Recognition In Video,’ *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 732–739, Columbus, OH, 24–27 June 2014, poster, 12 citations.

<http://engineering.purdue.edu/~qobi/papers/cvpr2014.pdf>

A. Barbu, D.P. Barrett, W. Chen, N. Siddharth, C. Xiong, J.J. Corso, C.D. Fellbaum, C. Hanson, S.J. Hanson, S. Hélie, E. Malaia, B.A. Pearlmuter, J.M. Siskind, T.M. Talavage, and R.B. Wilbur, ‘Seeing is *Worse* than Believing: Reading People’s Minds Better than Computer-Vision Methods Recognize Actions,’ *Proceedings of the Thirteenth European Conference on Computer Vision (ECCV)*, volume V, pp. 612–627, Zürich, Switzerland, 8–11 September 2014, poster, 6 citations.

<http://engineering.purdue.edu/~qobi/papers/eccv2014.pdf>

2015

J.M. Siskind, ‘Conducting Neuroscience to Guide the Development of AI,’ *Proceedings of the Twenty-Ninth Conference on Artificial Intelligence (AAAI)*, pp. 4067–4072, Austin, TX, 25–30 January 2015, oral, blue sky ideas/senior member track, 1 citation.

<http://engineering.purdue.edu/~qobi/papers/aaai2015a.pdf>

H. Yu and J.M. Siskind, ‘Learning to Describe Video with Weak Supervision by Exploiting Negative Sentential Information,’ *Proceedings of the Twenty-Ninth Conference on Artificial Intelligence (AAAI)*, pp. 3855–3863, Austin, TX, 25–30 January 2015, oral, 1 citation.

<http://engineering.purdue.edu/~qobi/papers/aaai2015b.pdf>

2016

J.M. Siskind and B.A. Pearlmuter, ‘Binomial Checkpointing for Arbitrary Programs with No User Annotation,’ *Extended abstract, presented at the International Conference on Algorithmic Differentiation (AD)*, Oxford, UK, 12–15 September 2016, oral.

<http://engineering.purdue.edu/~qobi/papers/ad2016a.pdf>

J.M. Siskind and B.A. Pearlmuter, ‘Efficient Implementation of a Higher-Order Language with Built-In AD,’ *Extended abstract, presented at the International Conference on Algorithmic Differentiation (AD)*, Oxford, UK, 12–15 September 2016, oral.

<http://engineering.purdue.edu/~qobi/papers/ad2016b.pdf>

A.G. Baydin, B.A. Pearlmuter, and J.M. Siskind, ‘DiffSharp: An AD Library for .NET Languages,’ *Extended abstract, presented at the International Conference on Algorithmic Differentiation (AD)*, Oxford, UK, 12–15 September 2016, oral.

<http://engineering.purdue.edu/~qobi/papers/ad2016c.pdf>

A.G. Baydin, B.A. Pearlmuter, and J.M. Siskind, ‘Tricks from Deep Learning,’ *Extended abstract, presented at the International Conference on Algorithmic Differentiation (AD)*, Oxford, UK, 12–15 September 2016, oral.

<http://engineering.purdue.edu/~qobi/papers/ad2016d.pdf>

R. Kelly, B.A. Pearlmuter, and J.M. Siskind, ‘Evolving the Incremental λ Calculus into a Model of Forward AD,’ *Extended abstract, presented at the International Conference on Algorithmic Differentiation (AD)*, Oxford, UK, 12–15 September 2016, poster.

<http://engineering.purdue.edu/~qobi/papers/ad2016e.pdf>

Patents Awarded

2004

J.M. Siskind, ‘Method of Determining the Stability of Two Dimensional Polygonal Scenes,’ US patent 6,693,630, 17 February 2004, 6 citations.

<http://google.com/patents/US6693630.pdf>

2005

J.M. Siskind, ‘Method for Computing all Occurrences of a Compound Event from Occurrences of Primitive Events,’ US patent 6,941,290, 6 September 2005, 6 citations.

<http://google.com/patents/US6941290.pdf>

2012

J.M. Siskind and B.A. Pearlmutter, ‘Map-Closure: A General Purpose Mechanism for Nonstandard Interpretation,’ US patent 8,281,299, 2 October 2012, 2 citations.

<http://google.com/patents/US8281299.pdf>

2014

J.M. Siskind and B.A. Pearlmutter, ‘Automatic Derivative Method for a Computer Programming Language,’ US patent 8,739,137, 27 May 2014, 7 citations.

<http://google.com/patents/US8739137.pdf>

2015

J.M. Siskind, A. Barbu, N. Siddharth, and H. Yu, ‘Correlating Videos and Sentences,’ US patent 9,183,466, 10 November 2015, 1 citation.

<https://patents.google.com/patent/US9183466B2/en>

Patent Applications Filed

2016

J.M. Siskind and H. Yu, ‘System and Method for Sentence Directed Video Object Codetection,’ US patent application 62/346459, filed 6 June 2016.

J.M. Siskind, D.P. Barrett, S.A. Bronikowski, and H. Yu, ‘System and Method for Controlling a Self-Guided Vehicle,’ US patent application 62/379245, filed 25 August 2016.

J.M. Siskind and B.A. Pearlmutter, ‘System and Method for Binomial Checkpointing,’ US patent application 62/393894, filed 13 September 2016.

Book Chapters

1997

J.M. Siskind, ‘Visual Event Perception,’ in *Symbolic Visual Learning*, K. Ikeuchi and M. Veloso ed., Oxford University Press, chapter 9, pp. 225–263, January 1997, 11 citations.

<http://engineering.purdue.edu/~qobi/papers/svl95.ps.Z>

2000

J.M. Siskind, ‘Learning Word-to-Meaning Mappings,’ in *Models of Language Acquisition: Inductive and Deductive Approaches*, P. Broeder and J. Murre ed., Oxford University Press, chapter 7, pp. 121–153, July 2000, 35 citations.

<http://engineering.purdue.edu/~qobi/papers/cmla96.ps.Z>

Workshop Publications

1991

J.M. Siskind, ‘Dispelling Myths about Language Bootstrapping,’ *Proceedings of the AAAI Spring Symposium Workshop on Machine Learning of Natural Language and Ontology*, pp. 157–164, March 1991, oral.

<http://engineering.purdue.edu/~qobi/papers/dispell.ps.Z>

J.M. Siskind, ‘Naive Physics, Event Perception, Lexical Semantics and Language Acquisition,’ *Proceedings of the AAAI Spring Symposium Workshop on Machine Learning of Natural Language and Ontology*, pp. 165–168, March 1991, oral.

<http://engineering.purdue.edu/~qobi/papers/aaai91ss.pdf>

1993

J.M. Siskind, ‘Solving a Lexical Acquisition Task via an Encoding as a Propositional Satisfiability Problem,’ *Proceedings of the CUNY Sentence Processing Conference*, p. 82, March 1993, poster.

<http://engineering.purdue.edu/~qobi/papers/cuny93.ps.Z>

J.M. Siskind, 'Grounding Language in Perception,' *Proceedings of the Annual Conference of the Society of Photo-Optical Instrumentation Engineers (SPIE)*, volume 2055, Intelligent Robots and Computer Vision XII, pp. 198–211, September 1993, oral, 7 citations.

<http://engineering.purdue.edu/~qobi/papers/spie93.ps.Z>

1994

J.M. Siskind, 'Axiomatic Support for Event Perception,' *Proceedings of the AAI Workshop on Integration of Natural Language and Vision Processing*, pp. 153–160, August 1994, oral.

<http://engineering.purdue.edu/~qobi/papers/aaai94ws.ps.Z>

1995

J.M. Siskind, 'Robust Lexical Acquisition Despite Extremely Noisy Input,' *Proceedings of the 19th Boston University Conference on Language Development*, MacLaughlin, D. and McEwen, S. ed., Cascadilla Press, pp. 587–598, March 1995, oral, 1 citation.

<http://engineering.purdue.edu/~qobi/papers/bucl95.ps.Z>

1996

J.M. Siskind, 'Unsupervised Learning of Visually-Observed Events,' *Proceedings of the AAI Fall Symposium Workshop on Learning Complex Behaviors in Adaptive Intelligent Systems*, pp. 82–83, November 1996, oral.

<http://engineering.purdue.edu/~qobi/papers/aaai96fs.ps.Z>

1998

J.M. Siskind, 'Visual Event Perception,' *Proceedings of the Ninth NEC Research Symposium*, August 1998, oral. Also available as Technical Report 99-033, NEC Research Institute, Inc., March 1999, 17 citations.

<http://engineering.purdue.edu/~qobi/papers/nara98.pdf>

2003

I. Pollak, J.M. Siskind, M.P. Harper, and C.A. Bouman, 'Multiscale Random Tree Models,' *IS&T/SPIE 15th Annual Symposium on Electronic Imaging*, Santa Clara, CA, 21–24 January 2003.

J.M. Siskind, I. Pollak, M.P. Harper, and C.A. Bouman, 'Stochastic Grammars for Images on Arbitrary Graphs,' *Proceedings of the IEEE Workshop on Statistical Signal Processing*, St. Louis, MO, 28 September–1 October 2003, 1 citation.

2004

I. Pollak, J.M. Siskind, M.P. Harper, and C.A. Bouman, 'Modeling of Images with Spatial Random Trees,' *SIAM Conference on Imaging Science*, Salt Lake City, UT, 3–5 May 2004.

2005

J.M. Siskind and B.A. Pearlmutter, 'Perturbation Confusion and Referential Transparency: Correct Functional Implementation of Forward-Mode AD,' *Draft Proceedings of the 17th International Workshop on Implementation and Application of Functional Languages (IFL2005)*, Dublin, Ireland, 19–21 September 2005, oral, 12 citations.

<http://engineering.purdue.edu/~qobi/papers/ifl2005.pdf>

2014

A. Barbu, N. Siddharth, and J.M. Siskind, 'Language-Driven Video Retrieval,' *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshop on Vision Meets Cognition (CVPR Workshop)*, Columbus, OH, 23 June 2014, poster.

<http://engineering.purdue.edu/~qobi/papers/fpic2014.pdf>

Technical Reports

1989

J.M. Siskind, 'The Culprit Pointer Method for Selective Backtracking,' S. M. thesis, Artificial Intelligence Laboratory, MIT, January 1989.

1991

D.A. McAllester and J.M. Siskind, 'Lifting Transformations,' AI Memo 1343, Artificial Intelligence Laboratory, MIT, December 1991, 4 citations.

<http://engineering.purdue.edu/~qobi/papers/AIM-1343.pdf>

1992

J.M. Siskind, 'Naive Physics, Event Perception, Lexical Semantics and Language Acquisition,' Ph.D. thesis, Artificial Intelligence Laboratory, MIT, January 1992, 93 citations.

<http://engineering.purdue.edu/~qobi/papers/phd.ps.Z>

1993

J.M. Siskind and D.A. McAllester, 'Screamer: A Portable Efficient Implementation of Nondeterministic Common Lisp,' Technical Report IRCS-93-03, Institute for Research in Cognitive Science, University of Pennsylvania, 1993, 44 citations.

<http://engineering.purdue.edu/~qobi/papers/ircs-93-03.ps.Z>

J.M. Siskind, 'Lexical Acquisition as Constraint Satisfaction,' Technical Report IRCS-93-41, Institute for Research in Cognitive Science, University of Pennsylvania, 1993, 10 citations.

<http://engineering.purdue.edu/~qobi/papers/ircs-93-41.ps.Z>

1999

J.M. Siskind, 'Flow-Directed Lightweight Closure Conversion,' Technical Report 99-105, NEC Research Institute, Inc., July 1999. Revised as Technical Report 99-190R, NEC Research Institute, Inc., December 1999, 30 citations.

<http://engineering.purdue.edu/~qobi/papers/fdlcc.pdf>

2002

S. Wang and J.M. Siskind, 'Image Segmentation with Ratio Cut—Extended Version,' Technical Report TR-ECE-02-07, School of Electrical and Computer Engineering, Purdue University, 2002, 3 citations.

<http://engineering.purdue.edu/~qobi/papers/TR-ECE-02-07.pdf>

2003

I. Pollak, J.M. Siskind, M.P. Harper, and C.A. Bouman, 'Spatial Random Trees and the Center-Surround Algorithm,' Technical Report TR-ECE-03-03, School of Electrical and Computer Engineering, Purdue University, 2003, 12 citations.

<http://engineering.purdue.edu/~qobi/papers/TR-ECE-03-03.pdf>

2008

J.M. Siskind and B.A. Pearlmutter, 'Using Polyvariant Union-Free Flow Analysis to Compile a Higher-Order Functional-Programming Language with a First-Class Derivative Operator to Efficient Fortran-like Code,' Technical Report TR-ECE-08-01, School of Electrical and Computer Engineering, Purdue University, January 2008, 8 citations.

<http://docs.lib.purdue.edu/ecetr/367>

J.M. Siskind and B.A. Pearlmutter, 'Putting the Automatic Back into AD: Part I, What's Wrong,' Technical Report TR-ECE-08-02, School of Electrical and Computer Engineering, Purdue University, January 2008.

<http://docs.lib.purdue.edu/ecetr/368>

B.A. Pearlmutter and J.M. Siskind, 'Putting the Automatic Back into AD: Part II, Dynamic, Automatic, Nestable, and Fast,' Technical Report TR-ECE-08-03, School of Electrical and Computer Engineering, Purdue University, January 2008, 1 citation.

<http://docs.lib.purdue.edu/ecetr/369>

J.M. Siskind and B.A. Pearlmutter, 'Nesting Forward-Mode AD in a Functional Framework,' Technical Report TR-ECE-08-09, School of Electrical and Computer Engineering, Purdue University, October 2008.

<http://docs.lib.purdue.edu/ecetr/377>

2011

J.M. Siskind, 'Grounding the Lexical Semantics of Verbs in Visual Perception Using Force Dynamics and Event Logic,' arXiv:1106.0256, 1 June 2011.

<http://arxiv.org/abs/1106.0256>

A.P. Fern, R.L. Givan, and J.M. Siskind, 'Specific-to-General Learning for Temporal Events with Application to Learning Event Definitions from Video,' arXiv:1106.4572, 22 June 2011.

<http://arxiv.org/abs/1106.4572>

2012

A. Radul, B.A. Pearlmutter, and J.M. Siskind, 'AD in Fortran, Part 1: Design,' arXiv:1203.1448, 7 March 2012, 3 citations.

<http://arxiv.org/abs/1203.1448>

A. Radul, B.A. Pearlmutter, and J.M. Siskind, 'AD in Fortran, Part 2: Implementation via Preprocessor,' arXiv:1203.1450, 7 March 2012. <http://arxiv.org/abs/1203.1450>

A. Barbu, A. Michaux, N. Siddharth, and J.M. Siskind, 'Simultaneous Object Detection, Tracking, and Event Recognition,' arXiv:1204.2741, 12 April 2012.

<http://arxiv.org/abs/1204.2741>

A. Barbu, A. Bridge, Z. Burchill, D. Coroian, S. Dickinson, S. Fidler, A. Michaux, S. Mussman, N. Siddharth, D. Salvi, L. Schmidt, J. Shangquan, J.M. Siskind, J. Waggoner, S. Wang, J. Wei, Y. Yin, and Z. Zhang, 'Video In Sentences Out,' arXiv:1204.2742, 12 April 2012.

<http://arxiv.org/abs/1204.2742>

N. Siddharth, A. Barbu, and J.M. Siskind, 'Seeing Unseeability to See the Unseeable,' arXiv:1204.2801, 12 April 2012.

<http://arxiv.org/abs/1204.2801>

A. Barbu, A. Bridge, D. Coroian, S. Dickinson, S. Mussman, N. Siddharth, D. Salvi, L. Schmidt, J. Shangquan, J.M. Siskind, J. Waggoner, S. Wang, J. Wei, Y. Yin, and Z. Zhang, 'Large-Scale Automatic Labeling of Video Events with Verbs Based on Event-Participant Interaction,' arXiv:1204.3616, 16 April 2012, 3 citations.

<http://arxiv.org/abs/1204.3616>

O. Manzyuk, B.A. Pearlmutter, A.A. Radul, D.R. Rush, and J.M. Siskind, 'Confusion of Tagged Perturbations in Forward Automatic Differentiation of Higher-Order Functions,' arXiv:1211.4892, 20 November 2012, 3 citations.

<http://arxiv.org/abs/1211.4892>

2013

A. Barbu, N. Siddharth, C. Xiong, J.J. Corso, C.D. Fellbaum, C. Hanson, S.J. Hanson, S. Hélie, E. Malaia, B.A. Pearlmutter, J.M. Siskind, T.M. Talavage, and R.B. Wilbur, 'The Compositional Nature of Verb and Argument Representations in the Human Brain,' arXiv:1306.2293, 10 June 2013.

<http://arxiv.org/abs/1306.2293>

D.P. Barrett and J.M. Siskind, 'Felzenszwalb-Baum-Welch: Event Detection by Changing Appearance,' arXiv:1306.4746, 20 June 2013.

<http://arxiv.org/abs/1306.4746>

H. Yu and J.M. Siskind, 'Discriminative Training: Learning to Describe Video with Sentences, from Video Described with Sentences,' arXiv:1306.5263, 21 June 2013.

<http://arxiv.org/abs/1306.5263>

N. Siddharth, A. Barbu, and J.M. Siskind, 'Seeing What You're Told: Sentence-Guided Activity Recognition in Video,' arXiv:1308.4189, 19 August 2013.

<http://arxiv.org/abs/1308.4189>

A. Barbu, N. Siddharth, and J.M. Siskind, ‘Saying What You’re Looking For: Linguistics Meets Video Search,’ arXiv:1309.5174, 20 September 2013, 2 citations.

<http://arxiv.org/abs/1309.5174>

2014

A. Barbu, A. Bridge, Z. Burchill, D. Coroian, S. Dickinson, S. Fidler, A. Michaux, S. Mussman, N. Siddharth, D. Salvi, L. Schmidt, J. Shangguan, J.M. Siskind, J. Waggoner, S. Wang, J. Wei, Y. Yin, and Z. Zhang, ‘Video in sentences out,’ arXiv:1408.6418, 9 August 2014.

<http://arxiv.org/abs/1408.6418>

H. Yu, D.P. Barrett, and J.M. Siskind, ‘A Faster Method for Tracking and Scoring Videos Corresponding to Sentences,’ arXiv:1411.4064, 14 November 2014.

<http://arxiv.org/abs/1411.4064>

2015

A. Barbu, N. Siddharth, C. Xiong, J.J. Corso, C.D. Fellbaum, C. Hanson, S.J. Hanson, S. Hélie, E. Malaia, B.A. Pearlmutter, J.M. Siskind, T.M. Talavage, and R.B. Wilbur, ‘The Compositional Nature of Event Representations in the Human Brain,’ arXiv:1505.06670, 25 May 2015.

<http://arxiv.org/abs/1505.06670>

H. Yu and J.M. Siskind, ‘Sentence Directed Video Object Codetection,’ arXiv:1506.02059, 5 June 2015, 2 citations.

<http://arxiv.org/abs/1506.02059>

D.P. Barrett, S.A. Bronikowski, H. Yu, and J.M. Siskind, ‘Robot Language Learning, Generation, and Comprehension,’ arXiv:1508.06161, 25 August 2015, 3 citations.

<http://arxiv.org/abs/1508.06161>

D.P. Barrett, R. Xu, H. Yu, and J.M. Siskind, ‘Collecting and Annotating the Large Continuous Action Dataset,’ arXiv:1511.05914, 18 November 2015.

<http://arxiv.org/abs/1511.05914>

A.G. Baydin, B.A. Pearlmutter, A.A. Radul, and J.M. Siskind, ‘Automatic differentiation in machine learning: a survey,’ arXiv:1502.05767, 19 April 2015, 22 citations.

<http://arxiv.org/abs/1502.05767>

A.G. Baydin, B.A. Pearlmutter, and J.M. Siskind, ‘DiffSharp: Automatic Differentiation Library,’ arXiv:1511.07727, 26 November 2015.

<http://arxiv.org/abs/1511.07727>

Invited Presentations

1982

IEEE Asilomar Microcomputer Workshop, April 1982.

Edinburgh Conference on Silicon Compilation, July 1982.

1984

IEEE Catalina Workshop on Logic Synthesis and Silicon Compilation, April 1984.

Panel discussion on Silicon Compilation, *Design Automation Conference (DAC)*, June 1984.

1985

Panel discussion on Silicon Compilation, *Custom Integrated Circuits Conference (CICC)*, May 1985.

1991

‘Dispelling Myths about Language Bootstrapping,’ *Boston University Conference on Language Development*, October 1991.

1994

‘Computational Models of Lexical Acquisition and Event Perception in Support of Lexical Acquisition,’ *Workshop on Cognitive Models of Language Acquisition*, April 1994.

1999

‘Visual Event Perception,’ Invited presentation at *ICCV workshop on Integration of Speech and Image Understanding*, Corfu, Greece, Tuesday 21 September 1999, declined.

2002

Commonsense Symposium, April 2002.

2003

‘Computational Modeling of Child Language Acquisition: A retrospective on Fifteen Years of Research,’ CNRS, June 2003.

2004

‘Stochastic Spatio-Temporal Grammars for Images and Video,’ *AAAI Workshop on Anchoring*, July 2004.

‘Stochastic Spatio-Temporal Grammars for Images and Video,’ *AI in the Wild, Belgian-Dutch Conference on Artificial Intelligence (BNAIC)*, October 2004.

2007

‘Learning to Represent the Lexical Semantics of Verbs with Force Dynamics from Visual Input,’ Trinity College, Dublin, Tuesday 16 January 2007.

‘Stochastic Spatio-Temporal Grammars for Images and Video,’ *NIPS 2007 workshop on The Grammar of Vision: Probabilistic Grammar-Based Models for Visual Scene Understanding and Object Categorization*, Saturday 8 December 2007.

2008

Invited to participate and present at the *Dagstuhl Seminar on Logic and Probability for Scene Interpretation*, Sunday 17–Friday 22 February 2008, declined due to family illness.

‘AD for Probabilistic Programming,’ *NIPS 2008 workshop on Probabilistic Programming: Universal Languages and Inference; systems; and applications*, Saturday 13 December 2008.

2011

‘Generating English Utterances to Describe Events in Video,’ *CVPR 2011 workshop on Activity Recognition Competition*, Monday 20 June 2011.

‘Mediating Cross-Modal Perception, Motor Control, Language, and Reasoning with Common and Deep Semantic Representations,’ *AAAI 2011 workshop on Language-Action Tools for Cognitive Artificial Agents: Integrating Vision, Action, and Language*, Sunday 7 August 2011.

‘Grounding Natural Language in Computer Vision and Robotics,’ *NIPS 2011 workshop on Integrating Language and Vision*, Friday 16 December 2011.

2012

‘Mediating Cross-Modal Perception, Motor Control, Language, and Reasoning with Common and Deep Semantic Representations,’ University of Buffalo (SUNY), Wednesday 7 March 2012.

2013

‘Learning to Say What You See and See What You Say,’ *Joint NAACL/ICML Symposium on Natural Language Processing (NLP) and Machine Learning (ML)*, Saturday 15 June 2013.

‘Seeing and Saying: Compositionally Linking Vision and Language,’ *UW MSR Summer Institute: Understanding Situated Language in Everyday Life*, Wednesday 24 July 2013.

2014

‘Seeing, Saying, Doing, and Thinking: The Compositional Structure of Perception, Language, Action, and Thought,’ *AAAI Spring Symposia: Knowledge Representation and Reasoning in Robotics and Qualitative Representations for Robots*, Monday 24 March 2014.

‘Learning to Ground Sentences in Video,’ *NIPS 2014 workshop on Learning Semantics*, Friday 12 December 2014.

2015

‘What every machine-learning researcher should know about AD,’ *DALI 2015 workshop on Probabilistic Programming*, Friday 10 April 2015.

‘Computers XOR Vision XOR Language,’ *CVPR 2015 workshop on Vision and Language*, Thursday 11 June 2015.

‘Common Sense Through Language Grounded in Vision and Motor Control,’ *CogSci 2015 workshop on Language & Common sense: Integrating across psychology, linguistics, and computer science*, Wednesday 22 July 2015.

‘Sentence-directed video object codetection (from egocentric mobile-robot video),’ *ICCV 2015 workshop on Describing and Understanding Video & The Large Scale Movie Description Challenge (LSMDC)*, Saturday 12 December 2015.

Teaching Experience

University of Pennsylvania

September–December 1992

Adjunct Professor, Department of Computer and Information Science

CIS520, the graduate introductory core course on artificial intelligence, Fall 1992.

Philadelphia, Pennsylvania

University of Toronto

August 1993–December 1995

Assistant Professor, Department of Computer Science

CSC324, the undergraduate programming languages course, Fall 1995.

CSC484, the undergraduate introductory course on artificial intelligence, Spring 1994, Spring 1995.

CSC485/CSC2501, the graduate introductory course on natural-language understanding, Fall 1993, Fall 1994, Fall 1995.

CSC2528, the graduate reading seminar on computational linguistics, Spring 1994, Spring 1995.

Toronto, Ontario

Technion, Israel Institute of Technology

February–August 1996

Senior Lecturer, Department of Electrical Engineering

048864, a graduate seminar on visual event perception, Spring 1996.

Haifa, Israel

University of Vermont

September 1996–May 1997

Visiting Assistant Professor, Department of Computer Science and Electrical Engineering

CS103, the undergraduate programming languages course, Fall 1996, Spring 1997.

CS251, the undergraduate introductory course on artificial intelligence, Fall 1996.

CS295, the undergraduate introductory course on computational linguistics, Spring 1997.

Burlington, Vermont

Princeton University

September 1998–February 1999

Visiting Lecturer, Department of Computer Science

COS302, the undergraduate introductory course on artificial intelligence, Fall 1998.

Princeton, New Jersey

Purdue University**West Lafayette, Indiana**

January 2002–present

Associate Professor in the School of Electrical and Computer Engineering

ECE47300, the undergraduate introductory course on artificial intelligence, Spring 2002, Spring 2003, Spring 2004, Spring 2005, Spring 2006, Spring 2007, Spring 2008, Spring 2009, Spring 2010, Spring 2011, Spring 2012, Spring 2013, Spring 2014, Spring 2015.

ECE57000, the graduate introductory course on artificial intelligence, Fall 2002, Fall 2003, Fall 2004, Fall 2005, Fall 2006, Fall 2007, Fall 2008, Fall 2009, Fall 2010, Fall 2011, Fall 2012, Fall 2013, Fall 2014, Fall 2015.

Current Students

Hamad Ahmed, Ph.D. student.

Thomas Ilyevsky, Ph.D. student.

Jared Johansen, Ph.D. student.

Chengxi Li, Ph.D. student.

Ren Li, Ph.D. student.

Former Students**2005**Vamsi Vytla, M.S., May 2005, *Recovering Unobservable Properties of a Visual Scene Using Perceiver Framework*, Purdue University.Pranay Gupta, M.S., December 2005, *Visual Perception and Imitation Grounded in Action*, Purdue University.**2013**Andrei Barbu, Ph.D., December 2013, *Reasoning Across Language and Vision in Machines and Humans*, Purdue University.**2014**Siddharth Narayanaswamy, Ph.D., May 2014, *Compositionality in Vision and Language*, Purdue University.**2016**Daniel Paul Barrett, Ph.D., May 2016, *Learning in Vision and Robotics*, Purdue University.Scott Alan Bronikowski, Ph.D., May 2016, *Grounding Robot Motion in Natural Language and Visual Perception*, Purdue University.Haonan Yu, Ph.D., May 2016, *Grounding Language in Video*, Purdue University.External Service**1990**Session chair, Phonology and Syntax, *Annual Conference of the Cognitive Science Society (CogSci)*, 1990.**1991**Member student session program committee, *Annual Meeting of the Association for Computation Linguistics (ACL)*, 1991.**1995**Member program committee, *International Workshop on Constraint-Based Reasoning, International Conference of the Florida Artificial Intelligence Research Society (FLAIRS)*, 1995.Member program committee, *Computational Models for Integrating Language and Vision, AAAI Fall Symposium Series*, 1995.

1996

Member program committee, *National Conference on Artificial Intelligence (AAAI)*, 1996.

Member program committee, *Second International Workshop on Constraint Based Reasoning, International Conference of the Florida Artificial Intelligence Research Society (FLAIRS)*, 1996.

Member program committee, *Workshop on the Representations and Processes between Vision and Natural Language, European Conference on Artificial Intelligence (ECAI)*, 1996.

Member international advisory committee, *Association for Computational Linguistics Special Interest Group on Natural Language Learning (SIGNLL)*, 1996–1997.

1997

Member program committee, *Conference on Computational Psycholinguistics (CPL)*, 1997.

1999

Member program committee, *Workshop on Computational Natural Language Learning, Conference of the European Chapter of the Association for Computational Linguistics (EACL)*, 1999.

2001

Member program committee, *Musical Constraints Workshop, Seventh International Conference on Principles and Practice of Constraint Programming*, 2001.

2003

Member program committee, *Scheme Workshop*, 2003.

Co-chair and co-organizer, *Workshop on Learning Word Meaning from Non-Linguistic Data, Human Language Technology and North American Association for Computational Linguistics (HLT-NAACL)*, with Regina Barzilay and Ehud Reiter, 2003.

2004

Member program committee, *Workshop on Psycho-computational Models of Human Language Acquisition (PsychoComp), International Conference on Computational Linguistics (COLING)*, 2004.

Member program committee, *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2004.

Member program committee, *IEEE Workshop on Perceptual Organization in Computer Vision (POCV), IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2004.

2005

Member program committee, *Annual Conference of the Cognitive Science Society (CogSci)*, 2005.

Member program committee, *International Joint Conference on Artificial Intelligence (IJCAI)*, 2005.

2006

Member program committee, *Annual Conference of the Cognitive Science Society (CogSci)*, 2006.

Member program committee, *International Conference on Development and Learning (ICDL)*, 2006.

Co-chair and co-organizer, *IEEE Workshop on Perceptual Organization in Computer Vision (POCV), IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, with James Elder, 2006.

2007

Member program committee, *National Conference on Artificial Intelligence (AAAI)*, 2007.

Member program committee, *International Conference on Development and Learning (ICDL)*, 2007.

External hiring-panel member (faculty search and interview committee), *National University of Ireland Maynooth (NUIM)*, 2007.

2008

Member program committee, *National Conference on Artificial Intelligence (AAAI)*, 2008.

Member program committee, *European Lisp Symposium (ELS)*, 2008.

2010

Member program committee, *National Conference on Artificial Intelligence (AAAI)*, 2010.

Member program committee, *International Conference on Development and Learning (ICDL)*, 2010.

Session chair, *International Conference on Robotics and Automation (ICRA)*, 2010.

2011

Member program committee, *International Conference on Development and Learning (ICDL)*, 2011.

2012

Member program committee, *National Conference on Artificial Intelligence (AAAI)*, 2012.

Member program committee, *Sixth International Conference on Automatic Differentiation (AD)*, 2012.

External reviewer, *Programming Language Design and Implementation (PLDI)*, 2012.

Member program committee, *First Annual Conference for Advance in Cognitive Systems (ACS)*, 2012.

2013

Member program committee, *National Conference on Artificial Intelligence (AAAI)*, 2013.

Member program committee, *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013.

Associate editor and member editorial board, *Artificial Intelligence (AIJ)*, 2013–present.

Session chair, *Annual Meeting of the Association for Computation Linguistics (ACL)*, 2013.

2014

Member program committee, *Annual Meeting of the Association for Computation Linguistics (ACL)*, 2014.

2016

Program co-chair, senior member track, *Conference on Artificial Intelligence (AAAI)*, with Sarit Kraus, 2016.

Professional Society Membership

Senior member IEEE, member Computer Society, member ACM, member AAAI, member ACL, member SIGNLL, member Cognitive Science Society, member AAUP.

Internal Service

1993

Richard Mann, graduate research, 1993–1996, University of Toronto, postdoc, January 1998–July 1999, NEC Research Institute. Advisor: Allen Jepson.

Paul Glenn, undergraduate thesis, University of Toronto.

James Rootham, undergraduate research, University of Toronto.

Patrick Tierney, undergraduate research, University of Toronto.

Quaid Morris, undergraduate research, University of Toronto.

Alexander Budanitsky, undergraduate research, University of Toronto.

Philip Edmonds, M.S. reader, August 1993, University of Toronto. Advisor: Graeme Hirst.

Igor Jurisica, depth oral, October 1993, Ph.D. proposal, September 1994, University of Toronto.

Manfred Stede, Ph.D. proposal, May 1994, senate oral, January 1995, University of Toronto. Advisor: Graeme Hirst.

1994

Tony Plate, senate oral, July 1994, University of Toronto. Advisor: Geoffrey Hinton.

Carl Rasmussen, depth oral, November 1994, University of Toronto. Advisor: Geoffrey Hinton.

1995

David St. Onge, M.S. reader, May 1995, University of Toronto. Advisor: Graeme Hirst.

Daniel Ansari, M.S. reader, May 1995, University of Toronto. Advisor: Graeme Hirst.

Xiaoyuan Tu, depth oral, May 1995, University of Toronto. Advisor: Demetri Terzopolous.

1996

Yair Yehuda, Ph.D. defence, June 1996, Technion. Advisor: Daniel Berry.

1998

John Hale, Brown University, summer intern, 1998, NEC Research Institute.

1999

Amit Roy Chowdhury, University of Maryland, summer intern, 1999, NEC Research Institute.

2000

Rama Bindiganavale, Ph.D. defense, July 2000, University of Pennsylvania. Advisor: Norman Badler.

Song Wang, UIUC, summer intern, 2000, 2001, NEC Research Institute. Advisor: Zhi-Pei Liang.

2001

Curran Nachbar, Brown University, summer intern, 2001, NEC Research Institute.

2002

Min Lu, graduate research, Spring–Fall 2002, Ph.D. advisory committee, member, Purdue University.

Abhilasha Bhargav, summer intern, 2002, ECE49600, Fall 2002, Purdue University.

Blake Matheny, summer intern, 2002, 2003, undergraduate research, Fall 2002–Spring 2003, Purdue University.

James J. Sherman, Jr., summer intern, 2002, undergraduate research, Fall 2002–Spring 2003, SURF, 2003, Purdue University.

Vamsi Vytla, summer intern, 2002, undergraduate research, Fall 2002–Spring 2003, SURF, 2003, Purdue University.

Paul Boxer, Ph.D. reader, July 2002, Royal Melbourne Institute of Technology, Melbourne, Australia. Advisor: Victor B. Ciesielski.

Alan P. Fern, preliminary examination, November 2002, final examination, June 2004, Purdue University. Advisor: Robert L. Givan.

Keith Henderson, ECE49600, Fall 2002, Purdue University.

2003

CE-3 QE question, January 2003, August 2003, January 2004, August 2004, August 2005, January 2006, August 2006, January 2007, August 2007, August 2009, August 2010, August 2011, August 2012, August 2013, August 2014, Purdue University.

Matthew Charlton Greig, preliminary examination, April 2003, Purdue University. Advisor: Robert L. Givan.

Sungwook Yoon, preliminary examination, April 2003, final examination, July 2006, Purdue University. Advisor: Robert L. Givan.

Bingrui Foo, summer intern, 2003, undergraduate research, Fall 2003, Purdue University.

Pranay Gupta, summer intern, 2003, Purdue University.

Eri Haneda, M.S. defense, July 2003, Purdue University. Advisor: Thomas Michael Talavage.

2004

Kai-yuh Hsiao, Ph.D. qualifying exam, April 2004, Media Laboratory, MIT. Advisor: Deb Roy.

2005

Bharath Seshadri, preliminary examination, April 2005, final examination, December 2006, Purdue University. Advisor: Irith Pomeranz.

Jeremy Tryba, undergraduate research, Fall 2005–Spring 2006, Purdue University.

Lin Zhu, preliminary examination, December 2005, Purdue University. Advisor: Robert L. Givan.

Nuha Alkatib, advisory committee, member, Purdue University.

Huaili Wang, advisory committee, member, Purdue University.

Tracey Brown, non-thesis M.S. advisory committee, member, Purdue University.

Zhanat Abylkassym, non-thesis M.S. advisory committee, chair, Purdue University.

Joshua Vincent Dillon, Ph.D. advisory committee, member, Purdue University.

2006

Rajat Agarwal, undergraduate research, January–November 2006, Purdue University.

Andrei Barbu, University of Waterloo, SURF, 2006, Purdue University.

Sarah Fortune, Trinity College Dublin, SURF, 2006, Purdue University.

Shammi Didla, SURF, 2006, Purdue University.

Bingrui Foo, Ph.D. advisory committee, member, Purdue University.

2007

Iheukwumere Onwuka, University College London, undergraduate research, Fall 2007–Spring 2008, ECE49600, Fall 2007, Purdue University.

Ross Beranek, ECE49600, Fall 2007, Purdue University.

Adaku Daisy Egeolu, non-thesis M.S. advisory committee, member, Purdue University.

Shawn Alan Brownfield, thesis M.S. advisory committee, chair, Purdue University.

Rezwanuzzaman Chowdhury, non-thesis M.S. advisory committee, co-chair, Purdue University.

Johannes Jacobus Cilliers, non-thesis M.S. advisory committee, member, Purdue University.

Adam Michael Horton, non-thesis M.S. advisory committee, chair, Purdue University.

Stephen Robert Rudolph, non-thesis M.S. advisory committee, chair, Purdue University.

Wadzanai Duiton Lunga, Ph.D. advisory committee, member, Purdue University.

2008

Alejandrina Cristià, LING59000, Spring 2008, Purdue University.

Anchal Dube, SURF, 2008, undergraduate research, Fall 2008–Spring 2009, Purdue University.

Padmini Jaikumar graduate research, Fall 2008, Purdue University.

2009

Karl P. Ostmo, M.S. defense, April 2009, Purdue University. Advisor: David Ebert.

Huanren Zhang, non-thesis M.S. advisory committee, member, Purdue University.

Brian Jay Thomas, SURF, 2009, undergraduate research, Fall 2009–Spring 2010, Purdue University.

Isaac P. Jones, SURF, 2009, Purdue University.

Jeffrey Ryan Mizek, non-thesis M.S. advisory committee, member, Purdue University.

Chang Xu, non-thesis M.S. advisory committee, member, Purdue University.

Yixi Yang, thesis M.S. advisory committee, member, Purdue University.

2010

Seongwoon Ko, ECE49600, Fall 2010, Purdue University.

Alexander Bridge, undergraduate research, December 2010–January 2011, Purdue University.

Dan Coroian, undergraduate research, December 2010–January 2011, Purdue University.

Sam Mussman, undergraduate research, December 2010–January 2011, Purdue University.

Lara Schmidt, undergraduate research, December 2010–January 2011, Purdue University.

Jiangnan Shangguan, undergraduate research, December 2010–January 2011, Purdue University.

Jinlian Wei, undergraduate research, December 2010–January 2011, Purdue University.

Yifan Yin, undergraduate research, December 2010–January 2011, Purdue University.

2011

Aaron Michaux, graduate research, May 2011–March 2013, Purdue University.

Zachary J. Burchill, University of Pennsylvania, summer intern, 2011, 2012, 2013, 2014, Purdue University.

Ryan Buffington, graduate research, August 2011–May 2012, Purdue University.

Tommy Y. Chang, graduate research, August 2011–January 2013, Purdue University.

2012

David Schwartzman Cohenca, graduate research, January–June 2012, Purdue University.

2013

Jordan R Glassley, ECE47300, Spring 2013, honors contract, Purdue University.

Tsz Kwan Lam, LING69000, Spring 2013, preliminary examination, November 2013, final examination, June 2015, Purdue University. Advisor: Ronnie B. Wilbur.

Anantha P Raghuraman, graduate research, Fall 2013, non-thesis M.S. advisory committee, chair, Purdue University.

Xiran Wang, graduate research, August 2013–June 2014, Purdue University.

Rajesh Kalyanam, preliminary examination, September 2013, final examination, July 2016, Purdue University. Advisor: Robert L. Givan.

Scott Benjamin Pillow, non-thesis M.S. advisory committee, member, Purdue University.

Maxwell Allen Reeder, non-thesis M.S. advisory committee, member, Purdue University.

Xiaolin Wang, non-thesis M.S. advisory committee, member, Purdue University.

Joon Hee Choi, non-thesis M.S. advisory committee, chair, Purdue University.

Lin Yuan, preliminary examination, November 2013, Purdue University. Advisor: Robert L. Givan.

Hongbo Yang, non-thesis M.S. advisory committee, member, Purdue University.

Abdullah Bader Alshaibani, thesis M.S. advisory committee, member, Purdue University.

Yihan Ding, non-thesis M.S. advisory committee, member, Purdue University.

2014

Colin Graber, ECE47300, Spring 2014, honors contract, ECE49600, Spring 2015, Purdue University.