Aditya Khosla

CONTACT
INFORMATION

32 Vassar St, D-475B
Massachusetts Institute of Technology
Cambridge, MA 02139
USA

Phone: (559) 737-3398
E-mail: khosla@csail.mit.edu
Website: http://mit.edu/khosla
Google Scholar profile (Citations: 2080, h-index: 17)

RESEARCH
INTERESTS

Computer Vision: large-scale computer vision, human visual perception, action recognition, fine-grained image categorization, object detection/recognition, vision for graphics
Machine Learning: deep learning, graphical models, random forest, max-margin learning

EDUCATION

Massachusetts Institute of Technology, Cambridge, MA, USA   08/2011 – Present
Ph.D. Student in Computer Science (GPA: 5.0/5.0)

- Advisor: Antonio Torralba
- Minor: Finance

Stanford University, Stanford, CA, USA   09/2009 – 08/2011
M.S. in Computer Science (GPA: 4.0/4.0)

- Advisor: Fei-Fei Li
- Specialization: Artificial Intelligence

California Institute of Technology, Pasadena, CA, USA   09/2005 – 06/2009
B.S. in Computer Science, Electrical Engineering and Economics with Honors (GPA: 3.8/4.0)

- Advisor: Tracey Ho

PUBLICATIONS

[28] R. M. Cichy, **A. Khosla**, D. Pantazis, and A. Oliva. Neural Dynamics of the Cortical Representation of Scenes: Evidence from Magnetoencephalography and Deep Neural Networks. *To appear in NeuroImage, 2016.* Website

[27] **A. Khosla***, K. Krafka*, P. Kellnhofer, S. Bhandarkar, W. Matusik and A. Torralba. Eye Tracking for Everyone. *To appear in the IEEE Conference on Computer Vision and Pattern Recognition (**CVPR**), 2016.* (* indicates equal contribution)

[26] B. Zhou, **A. Khosla**, A. Lapedriza, A. Oliva and A. Torralba. Learning Deep Features for Discriminative Localization. *To appear in the IEEE Conference on Computer Vision and Pattern Recognition (**CVPR**), 2016.* Preprint

[25] **A. Khosla***, A. Recasens*, C. Vondrick, and A. Torralba. Where Are They Looking? *Advances in Neural Information Processing Systems (**NIPS**), 2015.* (**Spotlight presentation**, * indicates equal contribution) Project website

[24] **A. Khosla**, A. Raju, A. Torralba, and A. Oliva. Understanding and Predicting Image Memorability at a Large Scale. *IEEE International Conference on Computer Vision (**ICCV**), 2015.* Project website

[23] R. Dubey, J. Peterson, **A. Khosla**, M.-H. Yang, and B. Ghanem. What Makes an Object Memorable? *IEEE International Conference on Computer Vision (**ICCV**), 2015.* Project website

[22] B. Zhou, **A. Khosla**, A. Lapedriza, A. Oliva and A. Torralba. Object Detectors Emerge in Deep Scene CNNs. *International Conference on Learning Representations (**ICLR**), 2015.* (**Oral presentation**) Paper

[21] C. Vondrick, **A. Khosla**, H. Pirsiavash, T. Malisiewicz, and A. Torralba. Visualizing Object Detection Features. *International Journal of Computer Vision (**IJCV**).* Paper

[20] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, **A. Khosla**, M. Bernstein, A. Berg and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (**IJCV**), 2015.* Paper

[19] Z. Wu, S. Song, **A. Khosla**, F. Yu, L. Zhang, X. Tang, and J. Xiao. 3D ShapeNets: A Deep Representation for Volumetric Shapes. *IEEE Conference on Computer Vision and Pattern Recognition (**CVPR**), 2015.* (**Oral presentation**) Project website

[18] J. Lim, **A. Khosla**, and A. Torralba. FPM: Fine pose Parts-based Model with 3D CAD models. *European Conference on Computer Vision (**ECCV**), 2014.*Paper

[17] **A. Khosla**, A. Das Sarma, and R. Hamid. What makes an image popular? *International World Wide Web Conference (**WWW**), 2014.* Project website

[16] **A. Khosla***, B. An*, J. Lim*, and A. Torralba. Looking Beyond the Visible Scene. *IEEE Conference on Computer Vision and Pattern Recognition (**CVPR**), 2014.* (* indicates equal contribution) Project website

[15] **A. Khosla**, B. Yao, and L. Fei-Fei. An Integration of Randomization and Discrimination for Classifying Human-Object Interaction Activities. *Invited book chapter in Springer book: Human-Centered Social Media Analytics, 2013*

[14] **A. Khosla**, W. Bainbridge, A. Torralba, and A. Oliva. Modifying the Memorability of Face Photographs. *IEEE International Conference on Computer Vision (**ICCV**), 2013.* Project website

[13] C. Vondrick, **A. Khosla**, T. Malisiewicz, and A. Torralba. HOGgles: Visualizing Object Detection Features. *IEEE International Conference on Computer Vision (**ICCV**), 2013.* (**Oral presentation**) Project website

[12] **A. Khosla**, R. Hamid, C.J. Lin, and N. Sundaresan. Large Scale Video Summarization using Web Image Priors. *IEEE Conference on Computer Vision and Pattern Recognition (**CVPR**), 2013.*

[11] A. Oliva, P. Isola, **A. Khosla**, W. Bainbridge. What makes a picture memorable? *SPIE Newsroom Article, 2013.*Newsroom

[10] **A. Khosla**, J. Xiao, A. Torralba, and A. Oliva. Memorability of Image Regions. *Advances in Neural Information Processing Systems (**NIPS**), 2012.* Project website

[9] **A. Khosla***, J. Xiao*, P. Isola, A. Torralba, and A. Oliva. Image Memorability and Visual Inception. ***SIGGRAPH Asia*** *Technique Briefs, 2012 (Invited paper)* (* indicates equal contribution) Project website

[8] **A. Khosla**, T. Zhou, T. Malisiewicz, A.A. Efros, and A. Torralba. Undoing the Damage of Dataset Bias. *European Conference on Computer Vision (**ECCV**), 2012.* Project website

[7] B. Yao, X. Jiang, **A. Khosla**, A.L. Lin, L.J. Guibas, and L. Fei-Fei. Human Action Recognition by Learning Bases of Action Attributes and Parts. *IEEE International Conference on Computer Vision (**ICCV**), 2011.* (**Oral presentation, PASCAL VOC 2011 Winner, comp10**) Dataset website

[6] B. Yao, **A. Khosla**, and L. Fei-Fei. Classifying Actions and Measuring Action Similarity by Modeling the Mutual Context of Objects and Human Poses. *International Conference on Machine Learning (**ICML**), 2011.* (**Oral presentation** in the cross-conference track) Paper

[5] J. Ngiam, **A. Khosla**, M. Kim, J. Nam, H. Lee, and A.Y. Ng. Multimodal Deep Learning. *International Conference on Machine Learning (**ICML**), 2011.* (**Oral presentation**) Paper

[4] **A. Khosla***, B. Yao* and L. Fei-Fei. Combining Randomization and Discrimination for Fine-Grained Image Categorization. *IEEE Conference on Computer Vision and Pattern Recognition (**CVPR**), 2011.* (* indicates equal contribution) (**PASCAL VOC 2011/2012 Winner, comp9**)
Project website

[3] **A. Khosla**, Y. Cao, C.C.Y. Lin, H.K. Chiu, J. Hu, and H. Lee. An Integrated Machine Learning Approach to Stroke Prediction. *16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (**KDD**), 2010.* (**Oral presentation**) Code

[2] S. Vyetrenko, **A. Khosla** and T. Ho. On combining information theoretic and cryptographic approaches to network coding security against the pollution attack. *Asilomar Conference on Signals, Systems and Computers (**Asilomar**), 2009.*

[1] A. Mahabal, S. G. Djorgovski, R. Williams, A. Drake, C. Donalek, M. Graham, B. Moghaddam, M. Turmon, J. Jewell, **A. Khosla** and B. Hensley. Towards Real-time Classification of Astronomical Transients. *International conference on Classification and Discovery in Large Astronomical Surveys (**Class**), 2008.*

SMALL CAPS: WORKSHOPS/ TECH REPORTS

[3] **A. Khosla**, B. Zhou, Y. Aytar, and A. Oliva. Mini Places Challenge: A Teaching Tool for Data-Driven Computer Vision. *Technical Report, 2015.* Website

[2] **A. Khosla**, D. Hoiem, and S. Belongie. Analysis of Reviews for CVPR2012. *Technical Report, 2012* Project website

[1] **A. Khosla**, N. Jayadevaprakash, B. Yao and L. Fei-Fei. Stanford Dogs: Novel Dataset for Fine-Grained Image Categorization. *First Workshop on Fine-Grained Visual Categorization (FGVC), IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2011* Dataset website

SMALL CAPS: SELECTED HONORS AND AWARDS

Winner, Facebook Fellowship 2013

Amazon Education Research Grant 2014 ($3,000)

NVidia Hardware Grant 2014 (∼$10,000)

Winner, PASCAL VOC Action Classification (comp9) Challenge 2012

Best Research Award, MIT CSAIL Student Workshop (CSW) 2012

Finalist, Qualcomm Innovation Fellowship 2012

Winner, PASCAL VOC Action Classication (comp9 & comp10) Challenge 2011

Silver Medal in the National Chemistry Olympiad

Silver Medal in the National Olympiad of Informatics

Flintridge Foundation Summer Undergraduate Research Fellow

Nominated for Associated Students of California Institute of Technology (ASCIT) Teaching Award

Television Coverage	Interview: Could this formula be the key to photo-sharing success?	BBC, UK, 2014
	Interview: Latest Trends in Social Media	Zee News, India, 2014
	Playing With How We Keep Faces Straight	New York Times Minute, 2013
Selected Press Coverage	Scientists have uncovered exactly what makes a photo memorable	The Washington Post, 2015
	What Makes Unforgettable Images So Memorable?	The Atlantic, 2015
	MIT Researchers Train An Algorithm To Predict How Boring Your Selfie Is	TechCrunch, 2015
	See How Memorable Your Profile Pictures Are With MIT's MemNet Algorithm	Lifehacker, 2015
	Are Your Photos Memorable? This Algorithm Will Tell You	Discover Magazine, 2015
	Brain-Like Computer System Identifies, Predicts Most 'Memorable' Photos	NBC News, 2015
	MIT's AI Algorithm Predicts how Memorable your Photo Is	The TODAY Show, 2015
	An algorithm can tell if your face is forgettable	Engadget, 2015
	What Makes An Image Memorable? This MIT AI Might Hold The Key	Fast Company, 2015
	New Deep Learning Tool Shows You the Most Memorable Parts of Your Photos	Gizmodo, 2015
	AI Project Designed To Recognize Scenes Surprises By Identifying Objects, Too	TechCrunch, 2015
	MIT research analyzing faces may explain why these CEOs are successful	MarketWatch, 2015
	Algorithm Predicts A Neighborhood's Crime Rate Using Google StreetView	Fast Company, 2014
	Computers are learning to size up neighborhoods using photos	Engadget, 2014
	Can you out-race a computer?	MIT News, 2014
	Computer Eyesight Gets a Lot More Accurate	The New York Times, 2014
	Revolutionary Technique That Quietly Changed Machine Vision Forever	MIT Tech Review, 2014
	Now You Can Craft Your Selfie For Maximum Likes – With Science	TechCrunch, 2014
	This Site Can Tell You How Popular Your Instagram Will Be	TIME, 2014
	At Last, A Formula For What Photos The Internet Likes	Huffington Post, 2014
	MIT algorithm predicts how popular your Instagram photo will be	The Verge, 2014
	Algorithm Can Predict Your Photo's Popularity Before You Post It	Entrepreneur, 2014
	How guns and bikinis make your online photos more popular	The Guardian, UK, 2014
	Here's How to Craft the Perfect Instagram, According to Science	Cosmopolitan, 2014
	Pourquoi votre photo va marcher sur les réseaux sociaux	Le Monde, France, 2014
	Zelite Sto Vise Lajkova Na Instagramu?	Telegraf, Russia, 2014
	Algoritmo ajuda a prever popularidade de fotos no Instagram	Exame, Brazil, 2014
	How To Make Your Face (Digitally) Unforgettable	NPR, 2014
	How to make your face unforgettable	The Telegraph, UK, 2014
	New software makes your pictures more memorable	The Independent, UK, 2014
	Forget Photoshop this software will make your face UNFORGETTABLE	DailyMail, UK, 2014
	Cómo hacer que nadie olvide tu cara	ABC.es, Spain, 2014
	Playing With How We Keep Faces Straight	The New York Times, 2013
	You Will Remember 'Me'	New Scientist, 2013
	MIT researchers develop algorithm to make your face more memorable	Wired, UK, 2013
	This Algorithm Can Make Pictures of Your Face More Memorable	Gizmodo, 2013
	New algorithm makes your face more memorable	Times Of India, India, 2013
	Never Forget a Face	MIT Tech Review, 2013
	Teaching computers to see – by learning to see like computers	MIT News, 2013

SELECTED RESEARCH EXPERIENCE	**Massachusetts Institute of Technology, CSAIL**	08/2011 - Present

Massachusetts Institute of Technology, CSAIL   08/2011 - Present
Graduate Research Assistant with Antonio Torralba
- Developing algorithms to modify the memorability of images automatically
- Developing algorithms to determine dataset bias in various computer vision datasets
- Developing algorithms to invert and visualize features such as Histogram of Oriented Gradients

Microsoft Research, Interactive Visual Media (IVM) Group   06/2013 - 08/2013
Summer Research Intern mentored by Larry Zitnick and Piotr Dollar

eBay Research Labs, Computer Vision Group   06/2012 - 08/2012
Summer Research Intern mentored by Raffay Hamid
- Developed algorithms for video summarization, published in CVPR 2013.

Stanford University, Computer Science Department   06/2010 - 08/2011
Graduate Research Assistant with Fei-Fei Li
- Developed algorithms for action recognition in still images and fine-grained image categorzation
- Proposed a novel dataset for fine-grained image categorization (Stanford Dogs)
- Won the PASCAL VOC Action Classification Challenge 2011 & 2012 (CVPR 2011 paper)

Stanford University, Computer Science Department   09/2009 - 02/2011
Research Project with Andrew Ng
- Developed deep-learning algorithms to combine information from multi-modal data effectively
- Specific task involved audio-visual fusion for automatic speech recognition

Stanford University, Computer Science Department   09/2009 - 06/2010
Research Project with Honglak Lee and Junling Hu
- Proposed a novel feature selection algorithm and a prediction algorithm better suited to disease prediction, and specifically tested on stroke prediction in elderly patients

TEACHING EXPERIENCE

Massachusetts Institute of Technology, CSAIL   01/2014
<u>Instructor</u> *for Visual Recognition through Machine Learning Competition (6.S093)* Website
- Designed and taught a short course for beginners in computer vision to provide hands-on experience in building an object recognition system using state-of-the-art technologies like deep learning. Competition prizes were sponsored by Microsoft and Google.

Massachusetts Institute of Technology, CSAIL   09/2015 - 12/2015
<u>Lecturer</u> *and* <u>Head Teaching Assistant</u> *for Advances in Computer Vision (6.869)*
- Primary organizer of the Mini Places Challenge. This challenge has been made available to the computer vision community as a teaching tool for data-driven computer vision.

Massachusetts Institute of Technology, CSAIL   09/2013 - 12/2013
Graduate Teaching Assistant for Advances in Computer Vision (6.869)

Stanford University, Computer Science Department   01/2011 - 03/2011
Graduate Teaching Assistant for Introduction to Computer Vision (CS223B)
- Developed the problem sets, exams, syllabus and structure of the course.

California Institute of Technology, Electrical Engineering Department 01/2008 - 06/2009
Teaching Assistant for Principles of Microprocessor Systems and Microprocessors Laboratory
- Held recitations/office hours for students and graded assignments. Ensured all the hardware was in working condition for the course

California Institute of Technology, Electrical Engineering Department 09/2008 - 06/2009
Teaching Assistant for Experimental Projects in Electronic Circuits (EE91ab and EE90)
- Provided students with assistance in their electronic circuit projects, and helped the professor with the course administration

California Institute of Technology, Computer Science Department   09/2007 - 12/2007
Teaching Assistant for Introduction to Sequential Programming (CS1)

STUDENTS SUPERVISED	Adria Recasens, *PhD, MIT, Where are they looking?*	06/2014 – Present
	Harini Kannan, *MEng, MIT, Webcam-based Eye Tracking*	11/2015 – Present
	Prafulla Dhariwal, *UROP, MIT, Scaling up RGB-D Data*	10/2015 – Present
	Zhirong Wu, *PhD, CUHK, Learning Image Manipulation*	09/2015 – Present
	Xavier Puig, *Visiting student, MIT, Frames in Places*	11/2014 – Present
	Kyle Krafka, *PhD, University of Georgia, Eye Tracking for Everyone*	01/2015 – Present
	Rachit Dubey, *PhD, UC Berkeley, Generating Split-Depth Images*	03/2015 – Present
	Akshat Dave, *MS, UCSD, Predicting Personality from Faces*	06/2015 – Present
	Greg Beams, *SuperUROP, MIT, myMemory: Application Development*	09/2015 – Present
	Dion Low, *SuperUROP, MIT, myMemory: Game Design*	09/2015 – Present
	Mikael Mengistu, *SuperUROP, MIT, Building Individual Models of Memory*	09/2015 – Present
	Tossaporn Sangja, *UROP, MIT, Large-scale Face Memorability*	10/2015 – Present
	Eric Lau, *Summer UROP, MIT, Identifying individual differences in Memory*	06/2015 – 09/2015
	Akhil Raju, *MEng, MIT, Modifying the Memorability of Photos*	03/2014 – 06/2015
	Will Grathwohl, *Undergraduate, MIT, Real-time Content Moderation*	09/2013 – 09/2014
	Sebastian Leon, *Undergraduate, MIT, Real-time Content Moderation*	09/2013 – 09/2014
	Amartya Biswas, *UROP, MIT, Fast Approximate Feature Extraction*	03/2014 – 06/2014

ACADEMIC SERVICE

Guest Editor, International Journal of Computer Vision (IJCV) Special Issue on Scene Understanding, with James Hays, Derek Hoiem and Jianxiong Xiao

Workshop Chair, Workshop on Visualization for Deep Learning at ICML 2016 with John Canny, Polo Chau and Biye Jiang Website

Organizer, Places2 Scene Recognition Challenge, in conjunction with ImageNet Large Scale Visual Recognition Challenge (ILSVRC), at ICCV 2015 Website

'Reviewing the Reviewers' Organizer, with Derek Hoiem and Serge Belongie at CVPR 2012 (and also at ECCV 2012). Published a technical report that contains our findings. Website

Workshop Chair, 4[th] Scene Understanding Workshop (SUNw'16) at CVPR 2016 with James Hays, Silvio Savarese, and Jianxiong Xiao Website

Workshop Chair, 3[rd] Scene Understanding Workshop (SUNw'15) at CVPR 2015 with James Hays, Silvio Savarese, and Jianxiong Xiao Website

Workshop Chair, 2[nd] Scene Understanding Workshop (SUNw'14) at CVPR 2014 with James Hays, Derek Hoiem, and Jianxiong Xiao Website

Workshop Chair, 1[st] Scene Understanding Workshop (SUNw'13) at CVPR 2013 with James Hays, Derek Hoiem, and Jianxiong Xiao Website

Co-organizer, Fine-Grained Challenge (FGComp) at ICCV 2013 Website

Co-organizer, ImageNet Large Scale Visual Recognition Challenge (ILSVRC) 2012 Website

Reviewer, MIT Graduate Student Admissions, 2014

Reviewer, IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)

Reviewer, International Journal of Computer Vision (IJCV)

Reviewer, IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2015, 2016

Reviewer, International Conference on Computer Vision (ICCV) 2015

Reviewer, European Conference on Computer Vision (ECCV) 2012, 2014, 2016

Reviewer, Neural Information Processing Systems (NIPS) 2013, 2014, 2015

Reviewer, ACM SIGCHI 2016

Reviewer, SIGGRAPH Asia 2015

Program Committee, Fine-Grained Visual Categorization (FGVC^2) Workshop at CVPR 2015

Program Committee, First Workshop on Storytelling with Images and Videos at ECCV 2014

Program Committee, Fine-Grained Visual Categorization (FGVC^2) Workshop at CVPR 2013

Reviewer, Winter conference on Applications and Computer Vision (WACV), 2014, 2015, 2016

Reviewer, Pattern Recognition Letters (PRLetters)

Reviewer, Neurocomputing

Reviewer, Asian Conference on Computer Vision (ACCV) 2012, 2014

Reviewer, Indian Conference on Computer Vision, Graphics and Image Processing (ICVGIP), 2014

Reviewer, Transactions on Image Processing (TIP)

Reviewer, Transactions on Knowledge and Data Engineering (TKDE)

Reviewer, SPIE Journal of Optical Engineering (OE)

Reviewer, SPIE Journal of Electronic Imaging (JEI)

Reviewer, ACM Journal on Computing and Cultural Heritage (JOCCH)

Reviewer, Springer International Journal of Machine Learning and Cybernetics (JMLC)

Reviewer, IEEE Transactions on Cybernetics

Reviewer, Journal of Biomedical Informatics (JBI)

Invited Talks   Understanding the Internal Representation of CNNs. *Google Research.* Mountain View, CA, 2016. March 29, 2016.

Predicting Human Visual Memory using Deep Learning. *Adobe Research.* San Francisco, CA, 2016. March 24, 2016.

Predicting Human Visual Memory using Deep Learning. *Carnegie Mellon University.* Pittsburgh, PA, 2016. March 22, 2016.

Predicting Human Visual Memory using Deep Learning. *MIT Media Lab.* Cambridge, MA, 2016. March 14, 2016.

Predicting Human Visual Memory using Deep Learning. *University of Texas, Austin.* Austin, TX, 2016. March 4, 2016.

Understanding and Improving the Internal Representation of CNNs. *Amazon Lab 126.* Sunnyvale, CA, USA. January 29, 2016.

Where are they looking? *Re.Work Deep Learning Summit San Francisco.* San Francisco, CA, USA. January 28, 2016.

Building a Semantic Understanding of the Internal Representation of CNNs. *Vision and Graphics Seminar, MIT.* Cambridge, MA, USA. December 1, 2015.

Understanding and Improving the Internal Representation of CNNs. *Machine Learning Tea, MIT.* Cambridge, MA, USA. November 30, 2015.

Applications of Deep Learning in Computer Vision. *Invited Lecture, 6.869, MIT.* Cambridge, MA, USA. November 12, 2015.

Introduction to Deep Learning. *Invited Lecture, 6.869, MIT.* Cambridge, MA, USA. October 27, 2015.

Understanding and Improving the Internal Representation of CNNs. *Boston University IVC Seminar.* Boston, MA, USA. October 20, 2015.

Understanding and Improving the Internal Representation of CNNs. *Learn AI with the Best.* Digital conference. October 17, 2015.

Understanding and Improving the Internal Representation of CNNs. *Lincoln Laboratory Seminar.* Lexington, MA, USA. October 6, 2015.

Predicting Human Visual Memory using Deep Learning. *Re.Work Deep Learning Summit Boston.* Boston, MA, 2015. May 27, 2015.

Crafting the Perfect Selfie Using Computer Vision. *Invited Lecture, CMPSCI 670, University of Massachusetts.* Amherst, MA, USA. November 19, 2014.

Visual Intelligence: Seeing beyond the Immediate Image. *Wolfram Data Summit.* Washington, DC, USA. September 4, 2014.

Visual Intelligence: Beyond the Immediate Scene. *Scene Understanding Workshop (SUNw) at CVPR 2014.* Columbus, OH, USA. June 23, 2014.

Image Memorability and Visual Inception. *SIGGRAPH Asia, invited session: Computer Vision Meets Computer Graphics.* Singapore. December 1, 2012

Image Memorability and Visual Inception. *Graphics Workshop at National University of Singapore (NUS).* Singapore. December 2, 2012

Image Memorability and Visual Inception. *Nanyang Technological University (NTU).* Singapore. November 30, 2012

How to make an image more memorable? *MIT CSAIL Student Workshop (CSW).* Dedham, MA, USA. September 22, 2012

Image Inception: How to make an image more memorable? *Qualcomm Inc.* San Jose, CA, USA. April 26, 2012

Action Classication: An Integration of Randomization and Discrimination in A Dense Feature Representation. *ICCV Workshop on PASCAL Visual Object Classes Challenge.* Barcelona, Spain. November 7, 2011

Combining Randomization and Discrimination for Fine-Grained Image Categorization. *CVPR Workshop on Fine-Grained Visual Categorization.* Colorado Springs, CO, USA. June 25, 2011.

Automated transient classification using a Bayesian Event Classifier. *Southern California Conference for Undergraduate Research.* Cal Poly Pomona, CA, USA. November 22, 2008.

PROFESSIONAL EXPERIENCE	**VocalBee** (www.vocalbee.com)	07/2013 - Present

Chief Scientific Advisor
VocalBee is a startup focused on photo monetization by crowdsourcing photos for advertising. As their primary scientific advisor, I helped to develop the memorability and popularity engines (based on my research) that provide novel insights from images.

Mirador (www.mirador.im)   12/2013 - 09/2014
Scientific Advisor
Mirador was a startup focused on real-time content moderation. As their scientific advisor, I helped to develop better, faster and more scalable deep learning algorithms for content moderation.

SmarterWrite (www.smarterwrite.com)   05/2011 - 08/2013
Co-Founder
SmarterWrite was a startup focused on increasing the use of data-driven teaching in pre-university

education. We received offers of funding from various sources but decided to go separate ways due to prior commitments.

BotSquare Inc (www.botsquare.com)   03/2010 - 06/2010
*Computer Vision Researcher/Software Engineer with **Navneet Dalal***
BotSquare (now called Flutter) was a computer vision startup focused on developing a user-interface using hand gesture recognition from a single webcam. I prototyped and optimized basic HOG code for use on a regular machine in real-time. Flutter was acquired by Google in 2013.

Oracle USA, Inc, Environment Management System Group   06/2007 - 09/2007
Summer Intern with Sreeji Das
Tested the Xen system of virtual machines for use with the current infrastructure of the code, and evaluated its pros and cons of implementation.

COMPUTER SKILLS
- Languages/Markup: C/C++, Matlab, CUDA, Java, HTML, CSS, Javascript, SQL, LaTeX
- Operating Systems: Linux, Windows, Mac OS X.

REFERENCES   Available upon request