

DHRUV BATRA

CURRENT POSITION

Assistant Professor Fall 2016 – Present
School of Interactive Computing
Georgia Institute of Technology

Research Scientist Aug 2017 – Present
Facebook AI Research (FAIR)

Research Assistant Professor Fall 2016 – Present
Bradley Department of Electrical and Computer Engineering
Virginia Tech

Lead: Machine Learning & Perception Group

Research Interests: Machine Learning, Computer Vision, Artificial Intelligence

EDUCATION

Ph.D., Carnegie Mellon University 2010

Thesis:

“On Graph-Structured Discrete Labelling Problems in Computer Vision: Learning, Inference, and Applications.”

Committee: Tsuhan Chen (Chair), Carlos Guestrin, Martial Hebert, Jose Moura, Rahul Sukthankar, Ramin Zabih

M.S., Carnegie Mellon University 2006

Project Supervisor: Martial Hebert

B.Tech, Institute of Technology, Banaras Hindu University (Now called IIT-BHU), India 2005

SELECT HONORS / AWARDS

- **Outstanding Junior Faculty Research award**
College of Computing, Georgia Tech, 2018

Awarded the by Dean and Awards Committee of College of Computing at Georgia Tech to typically 1 Assistant Professor across all departments in the college for the quality of publications and impact of research.

- **AWS Machine Learning Research Award, 2018**

- **Best Paper Award**

Natural Language Does Not Emerge 'Naturally' in Multi-Agent Dialog.

Conference on Empirical Methods in Natural Language Processing (EMNLP), 2017.

One of 4 best papers (or top 0.26%) out of 1500 submissions (1466 reviewed, 323 accepted) to EMNLP 2017, one of the premier publication venues for research on AI with natural language capabilities.

- **Office of Naval Research (ONR) Young Investigator Program (YIP) award, 2017**

“Introduced in 1985, the ONR YIP is one of the nation’s oldest and most selective scientific research advancement programs. Its purpose is to fund early-career academic researchers—called investigators—whose scientific pursuits show outstanding promise for supporting the Department of Defense, while also

promoting their professional development.

In 2017, the awards were made to 33 scientists whose research holds strong promise across several naval-relevant science and technology areas ... selected from over 360 highly qualified applicants."

- **Amazon Academic Research Award, 2016**

- **Best Student Paper Award**

Yash Goyal, Akrit Mohapatra, Devi Parikh, Dhruv Batra.

Towards Transparent AI Systems: Interpreting Visual Question Answering Models.

International Conference on Machine Learning (ICML) Workshop on Visualization for Deep Learning, 2016.

- **Best Student Paper Award**

Abhishek Das, Harsh Agrawal, Larry Zitnick, Devi Parikh, Dhruv Batra.

Human Attention in Visual Question Answering: Do Humans and Deep Networks Look at the Same Regions?

International Conference on Machine Learning (ICML) Workshop on Visualization for Deep Learning, 2016.

- **Google Faculty Research Awards, 2013, 2015**

"Google Faculty Research Awards Program aims to recognize and support world-class, permanent faculty pursuing cutting-edge research in areas of mutual interest. ...The award is highly competitive - only 15% of applicants receive funding - and each proposal goes through a rigorous Google-wide review process."

- **Best Poster Paper Award**

Visual Question Answering, Antol et al., Workshop on Object Understanding for Interaction.

International Conference on Computer Vision (ICCV), 2015.

- **Outstanding Reviewer Award**

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

- **Outstanding New Assistant Professor award**

College of Engineering, Virginia Tech, 2015

Awarded the by Dean of College of Engineering at Virginia Tech to 5 Assistant Professors across all 12 departments in the college (~350 total faculty), selected for their research impact, teaching innovation, service, and outreach.

- **National Science Foundation (NSF) CAREER award, 2014**

"The Faculty Early Career Development (CAREER) Program is a Foundation-wide activity that offers the National Science Foundation's most prestigious awards in support of early-career faculty who have the potential to serve as academic role models in research and education and to lead advances in the mission of their department or organization."

- **Army Research Office (ARO) Young Investigator Program (YIP) award, 2014**

"YIP awards are one of the most prestigious honors bestowed by the Army on outstanding scientists beginning their independent careers. The objective of the YIP is to attract outstanding young university faculty Members to pursue fundamental research in areas relevant to the Army, to support their research in these areas, and to encourage their teaching and research careers."

- Virginia Tech Office of Vice President for Research
Scholar of the Week, 2014

- Virginia Tech Center for Instructional Development and Education Research
Teacher of the Week, 2013

- Carnegie Mellon University College of Engineering Dean's Fellowship, 2007
- Chairman, IEEE Student Chapter, IT-BHU (2004-05)

PAST APPOINTMENTS

Visiting Researcher Facebook AI Research (FAIR)	Fall 2016 – Summer 2017
Assistant Professor Bradley Department of Electrical and Computer Engineering, Virginia Tech Lead: Machine Learning & Perception Group @ VT Core Member: Virginia Center for Autonomous Systems (VaCAS) Faculty Member: VT Discovery Analytic Center (DAC)	Jan 2013 – Aug 2016
Visiting Researcher Microsoft Research, Redmond, WA Hosts: Larry Zitnick, Meg Mitchell	May – Aug 2015
Research Assistant Professor Toyota Technological Institute at Chicago	Sep 2010 – Dec 2012
Visiting Faculty, Carnegie Mellon University Host: Carlos Guestrin	May – Aug 2012
Visiting Scientist, CSAIL Massachusetts Institute of Technology Host: Bill Freeman	Mar – May 2011
Intern, Microsoft Research Lab Cambridge (UK) Mentors: Sebastian Nowozin and Pushmeet Kohli	Summer 2010
Visiting Student, Cornell University Mentors: Tsuhan Chen, Cornell University	Spring 2009 – Spring 2010
Intern, Kodak Research, Rochester, NY Mentor: Jiebo Luo, Eastman Kodak (now Professor at University of Rochester)	Summer 2008
Intern, Intel Research Lab, Pittsburgh, PA Mentor: Rahul Sukthankar, Intel; (Adjunct) Faculty, Robotics Institute, CMU	Summer 2007
Intern, Intel Research Lab, Pittsburgh, PA Mentor: Bart Nabbe, Intel; (Adjunct) Faculty, Robotics Institute, CMU	Summer 2006
Research Assistant, Vision and Mobile Robotics Lab, Carnegie Mellon University Mentor: Martial Hebert, Professor, Robotics Institute, CMU	Fall 2005 - Fall 2006
Intern, R&D Group, Electronics Division, Minda Huf Technologies Lmt., Delhi, India Project: Developing a Fingerprint Lock for Automobile Security	Summer 2005
Intern, Multimedia and Internet Technology Lab, Indian Institute of Technology (IIT) Delhi Mentor: Santanu Chaudhury, Professor, Indian Institute of Technology (IIT) Delhi	Summer 2004

TEACHING

Fall 2017		GT CS 7643 Deep Learning
Spring 2016	VT ECE 2574 Data-Structures and Algorithms	
Fall 2015	VT ECE 6504 Deep Learning for Perception	
Spring 2015	VT ECE 5984 Intro to Machine Learning	
Fall 2014	VT ECE 2574 Data-Structures and Algorithms	
Spring 2014	VT ECE 6504 Probabilistic Graphical Models and Structured Prediction	
Fall 2013	VT ECE 4984/5984 Intro to Machine Learning & Perception	
Spring 2013	VT ECE 4984/5984 Intro to Machine Learning & Perception	
Spring 2012	Co-Instructor (Instructor: Nati Srebro): TTIC 31070 Convex Optimization	
Fall 2010	Guest Lecturer: TTIC 31020 Intro to Statistical Machine Learning	
Fall 2008	Graduate Teaching Assistant (Instructor: Carlos Guestrin): CMU 10708 Probabilistic Graphical Models	

FUNDING & GIFTS

Research Grants & Gifts

- Gift funding from Mitsubishi Electric Research Lab (MERL)

2018
- AWS Machine Learning Research Award
"Towards AI Agents that can See, Talk, and Act"

2018
- DARPA Explainable AI (XAI)
"EQUAS: Explainable Question Answering System"
Lead Academic PI: Dhruv Batra. Co-PIs: Profs. Ray Mooney, Antonio Torralba, Devi Parikh. Prime: Raytheon.

2017 – 2021
- Office of Naval Research (ONR) Young Investigator Program (YIP) award
"Explainable and Trustworthy Intelligent Systems"

2017 – 2020
- Amazon Academic Research Award
"Visual Dialog"

2016
- Army Research Lab (ARL) Grant W911NF-15-2-0080
"Answering Binary Questions about Images"
Co-PI: Prof. Devi Parikh

2015 – 2018
- Defense Advanced Research Projects Agency (DARPA) Grant HR0011-16-1-0002
"Deep Cognitive Radio System Study"
Co-PI: Tim O'Shea (Hume Center, Virginia Tech)

2015 – 2016
- Google Faculty Research Award
"Visual Question Answering (VQA)"
Co-PI: Prof. Devi Parikh

2015
- Office of Naval Research (ONR) Grant N00014-14-1-0679
"Feedback-Enabled Joint Reasoning over Uncertain Sub-components of Perception"

2014 – 2017
- Army Research Office (ARO) Young Investigator Program (YIP) Award W911NF-14-1-0180
"Building Reflective, Transparent, and Integrated Intelligent Systems"

2014 – 2017
- NSF CAREER Award IIS-1350553
"Holistic Scene Understanding with Multiple Hypotheses from Vision Modules"

2014 – 2019

- Virginia Tech Institute for Critical Technology and Applied Science (ICTAS) Junior Faculty Award
"CloudCV: Large-Scale Distributed Computer Vision As A Cloud Service" 2014 – 2016
- L-3 Communications (sub-award from VT Hume Center)
"Detection of Naval Objects in Images" 2014 – 2015
- NSF Award IIS-1353694
"EAGER: Diverse M-Best Predictions from Probabilistic Models" 2013 – 2015

Equipment, Educational, Organizational Grants & Gifts

- Google Cloud Computing Support for CS 7643 Deep Learning 2017
- NRT-DESE: UrbComp: Data Science for Modeling, "Understanding, and Advancing Urban Populations"
PI: Naren Ramakrishnan (CS, VT) 2015 – 2020
- State Council of Higher Education for Virginia (SCHEV) Equipment Trust Grant 2014
- Oak Ridge Associated University (ORAU) Event Organization Grant
1st Mid Atlantic Computer Vision workshop at Virginia Tech; Co-PI: Devi Parikh. 2014
- Amazon Web Services (AWS) in Education: Machine Learning Research Grant 2013 – 2014
- Windows Azure for Research Award 2014
- NVIDIA Academic Hardware Gift 2014, 2015, 2016

RESEARCH GROUP AND STUDENT ADVISING

Staff

- Dr. Stefan Lee Research Scientist (2017 – Present)
- Dr. Peter Anderson Research Scientist (2018 – Present)

Students

- Aishwarya Agrawal PhD student (2014 – Present)
- Yash Goyal PhD student (2014 – Present)
- Michael Cogswell PhD student (2015 – Present)
- Abhishek Das PhD student (2016 – Present)
- Ashwin Kalyan PhD student (2016 – Present)
- Nirbhay Modhe PhD student (2017 – Present)
- Erik Wijmans PhD student (2017 – Present), Co-advised with Irfan Essa
- Harsh Agrawal PhD student (2018 – Present)

- Deshraj Yadav MS student (2017 – Present)

Visiting Students

- Rishabh Jain Intern, Fall 2017 – Spring 2018

Batra MLP Group Alumni (Graduated Students & Former Interns)

- Qing Sun PhD, Nov 2017 (first employment: Facebook)
- Gordon Christie PhD, Nov 2016 (first employment: Army Research Lab)

- Akrit Mohapatra MS, Apr 2018 (first employment: Ebay)
- Aroma Mahendru MS, Apr 2017 (first employment: Bloomberg)
- Latha Pemula MS, Sep 2016 (first employment: Amazon)
- Harsh Agrawal MS, May 2016 (first employment: Snapchat Research)
- Prakriti Banik MS, Jul 2015 (first employment: Bloomberg)
- Neelima Chavali MS, Jul 2015 (first employment: Many Trees Inc)
- Clint Solomon MS, May 2015 (first employment: PhotoKharma)

- Deshraj Yadav Intern, Fall 2016 – Spring 2017
 - BTech from JSS Noida; next position: MS student at Georgia Tech
- Tejas Khot Intern, Fall 2016 – Spring 2017
 - BTech from University of Mumbai; next position: MS student at RI, CMU
- Viraj Prabhu Intern, Fall 2016 – Spring 2017
 - BTech from BITS Pilani; next position: MS student at Georgia Tech
- Avi Singh Intern, Summer 2016
 - BTech from IIT-Kanpur; next position; PhD student at UC Berkeley
- Abhishek Das Intern, Fall 2015 – Spring 2016
 - BTech from IIT-Roorkee; next position: PhD student at VT
- Ashwin Kalyan Intern, Fall 2015 – Spring 2016
 - BTech from NIT Surathkal; next position: PhD student at VT
- Khushi Gupta Intern, Spring 2016
 - BTech from IIT-Gowahati; next position: MS student at RI, CMU
- Ahmed Osman Intern, Summer 2015
 - next position: MS student at UCL
- Faruk Ahmed Intern, Fall 2013 – Spring 2015
 - BTech from IEM; next position: PhD student at University of Montreal
- Senthil PS Intern, Summer 2013, Fall 2014 – Spring 2014
 - BTech student from IIT-G; next position: MS student at RI, CMU
- Aroma Mahendru Intern, Summer 2013, Fall 2014 – Spring 2014
 - BTech student from IIT-BHU; next position: MS student at VT
- Ratnesh Kumar Intern, Fall 2014
 - PhD student from INRIA; next position: MERL
- Qi Lou Intern, Fall 2013 – Spring 2014
 - MS student from OSU; next position: PhD student at UC Irvine
- Ankit Laddha Intern, Fall 2013 – Spring 2014
 - BTech student from IIT-Delhi; next position: MS student at RI, CMU
- Adarsh Prasad Intern, Summer 2013 – Fall 2013
 - BTech student IIT-Delhi; next position: PhD student at UT-Austin
- Harsh Agrawal Intern, Summer 2013
 - BTech student DTU; next position: MS student at VT
- Abhimanyu Dubey Intern, Summer 2013
 - BTech student IIT-Delhi

Former Student Collaborators & Interns

- Abner Guzmán Rivera, PhD student, UIUC (Advisor: Rob Rutenbar), Spring 2012 – Spring 2014
- Payman Yadollahpour, PhD student, TTIC (Advisor: Greg Shakhnarovich), Fall 2011 – Fall 2013
- Daozheng Chen, PhD University of Maryland (Advisor: David Jacobs), Fall 2011 – Spring 2013
- Kun Duan, PhD student, Indiana University (Advisor: David Crandall), Summer 2011
- Jason Lew, MS thesis, Cornell University (Advisor: Tsuhan Chen), Spring 2010
- Andrew Mui, MS thesis, Cornell University (Advisor: Tsuhan Chen), Spring 2010

- Tanmay Verma, BTech, IIIT-Delhi, Summer 2011
- Charvi Puri, BTech, IIIT-Delhi, Summer 2011
- Kevin Tang, BS, Cornell University, Spring 2009 (next position: PhD student at Stanford)

PUBLICATIONS / PATENTS / DEMOS

Book

1. Dhruv Batra, Adarsh Kowdle, Devi Parikh, Jiebo Luo, and Tsuhan Chen. Interactive Co-segmentation of Objects in Image Collections. *SpringerBriefs in Computer Science* 2011.

Book Chapter

2. Harsh Agrawal, Clint Solomon Mathialagan, Yash Goyal, Neelima Chavali, Prakriti Banik, Akrit Mohapatra, Ahmed Osman, Dhruv Batra. CloudCV: Large Scale Distributed Computer Vision as a Cloud Service. *Book Chapter, Mobile Cloud Visual Media Computing*. Editors: Gang Hua, Xian-Sheng Hua. Springer, 2015.

Thesis

3. Dhruv Batra. On Graph-Structured Discrete Labelling Problems in Computer Vision: Learning, Inference and Applications. *Ph.D. Thesis. Carnegie Mellon University*, August 2010.

Journal Articles

4. Abhishek Das, Satwik Kottur, Khushi Gupta, Avi Singh, Deshraj Yadav, Stefan Lee, José M. F. Moura, Devi Parikh, Dhruv Batra. Visual Dialog. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2018.
5. Abhishek Das*, Harsh Agrawal*, Larry Zitnick, Devi Parikh, Dhruv Batra. Human Attention in Visual Question Answering: Do Humans and Deep Networks Look at the Same Regions? *Computer Vision and Image Understanding (CVIU)*, 2017.
6. Gordon Christie, Ankit Laddha, Aishwarya Agrawal, Stanislaw Antol, Yash Goyal, Kevin Kochersberger, Dhruv Batra. Resolving Vision and Language Ambiguities Together: Joint Segmentation & Prepositional Attachment Resolution in Captioned Scenes. *Elsevier Journal on Computer Vision and Image Understanding (CVIU)*, 2017.
7. Aishwarya Agrawal, Jiasen Lu, Stanislaw Antol, Margaret Mitchell, C. Lawrence Zitnick, Devi Parikh, Dhruv Batra. VQA: Visual Question Answering. *International Journal of Computer Vision (IJCV), Special Issue on Combined Image and Language Understanding, Vol 123(1):4-31*, 2017.
8. C. Lawrence Zitnick, Aishwarya Agrawal, Stanislaw Antol, Margaret Mitchell, Dhruv Batra, Devi Parikh. Measuring Machine Intelligence Through Visual Question Answering. *AI Magazine, Vol 27, No 1*, 2016.
9. Vittal Premachandran, Daniel Tarlow, Alan Yuille, Dhruv Batra. Empirical Minimum Bayes Risk Prediction. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2016.
10. Dustin Bales, Pablo Tarazaga, Mary Kasarda, Dhruv Batra, Americo Woolard, Jeff Poston, Sriram Malladi. Gender Classification of Walkers Via Underfloor Accelerometer Measurements. *IEEE Internet of Things Journal, Vol 3 (6)*, 2016.
11. Delasa Aghamirzaie, Dhruv Batra, Lenwood S. Heath, Andrew Schneider, Ruth Grene, Eva Collakova. Transcriptome-wide functional characterization reveals novel relationships among differentially expressed transcripts in developing soybean embryos. *BMC Genomics*, 2015.

12. Joerg Kappes, Bjoern Andres, Christoph Schnoerr, Fred Hamprecht, Sebastian Nowozin, Dhruv Batra, Jan Lellmann, Nikos Komodakis, Sungwoong Kim, Bernhard Kausler, Carsten Rother. A Comparative Study of Modern Inference Techniques for Structured Discrete Energy Minimization Problems. *International Journal of Computer Vision (IJCV)*, 2015.
13. Kun Duan, Dhruv Batra, David Crandall. Human Pose Estimation via Multi-layer Composite Models. *Signal Processing*, 2014.
14. Adarsh Kowdle, Yao-Jen Chang, Andrew Gallagher, Dhruv Batra and Tsuhan Chen. Putting the User in the Loop for Image-Based Modeling. *International Journal of Computer Vision (IJCV)*, 2014.
15. Dhruv Batra, Adarsh Kowdle, Devi Parikh, Jiebo Luo, Tsuhan Chen. Interactively Co-segmenting Topically Related Images with Intelligent Scribble Guidance. *International Journal of Computer Vision (IJCV)*, 2011.

Peer-Reviewed Conference Papers (acceptance rates typically 2%-25%)

16. Satwik Kottur, José M. F. Moura, Devi Parikh, Dhruv Batra, Marcus Rohrbach. Visual Coreference Resolution in Visual Dialog using Neural Module Networks. European Conference on Computer Vision (ECCV), 2018.
17. Ramprasaath R. Selvaraju, Prithvijit Chattopadhyay, Mohamed Elhoseiny, Tilak Sharma, Dhruv Batra, Devi Parikh, Stefan Lee. Choose Your Neuron: Incorporating Domain Knowledge through Neuron Importance. European Conference on Computer Vision (ECCV), 2018.
18. Jianwei Yang, Jiasen Lu, Stefan Lee, Dhruv Batra, Devi Parikh. Graph R-CNN for Scene Graph Generation. European Conference on Computer Vision (ECCV), 2018.
19. Ashwin K Vijayakumar, Stefan Lee, Anitha Kannan, Dhruv Batra. Learn From Your Neighbor: Learning Multi-Modal Distributions from Sparse Annotation. International Conference on Machine Learning (ICML), 2018.
(Oral; Long Talk)
20. Abhishek Das, Samyak Datta, Georgia Gkioxari, Stefan Lee, Devi Parikh, Dhruv Batra. Embodied Question Answering. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
(Oral; Selection Rate: 70/3309 = 2.1%)
21. Aishwarya Agrawal, Dhruv Batra, Devi Parikh, Aniruddha Kembhavi. Don't Just Assume; Look and Answer: Overcoming Priors for Visual Question Answering. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
22. Jiasen Lu, Jianwei Yang, Dhruv Batra, Devi Parikh. Neural Baby Talk. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
(Spotlight)
23. Abhishek Mohta*, Ashwin K Vijayakumar*, Oleksandr Polozov, Dhruv Batra, Sumit Gulwani, Prateek Jain. Neural Guided Deductive Search for Real-time Program Synthesis. *International Conference on Learning Representations (ICLR)*, 2018.
24. Ashwin Vijayakumar, Michael Cogswell, Ramprasaath Selvaraju, Qing Sun, Stefan Lee, David Crandall, Dhruv Batra. Diverse Beam Search: Decoding Diverse Solutions from Neural Sequence Models. *AAAI Conference on Artificial Intelligence (AAAI)*, 2018.
25. Jiasen Lu, Anitha Kannan, Jianwei Yang, Devi Parikh, Dhruv Batra. Best of Both Worlds: Transferring Knowledge from Discriminative Learning to a Generative Visual Dialog Model. *Neural Information Processing Systems (NIPS)*, 2017.
26. Abhishek Das*, Satwik Kottur*, José M. F. Moura, Stefan Lee, Dhruv Batra. Learning Cooperative Visual Dialog Agents with Deep Reinforcement Learning. *International Conference on Computer Vision (ICCV)*, 2017.

(Oral; Selection Rate: 45/3220 = 1.4%)

27. Ramprasaath Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, Dhruv Batra. Grad-CAM: Why did you say that? Visual Explanations from Deep Networks via Gradient-based Localization. *International Conference on Computer Vision (ICCV)*, 2017.
28. Satwik Kottur, José M.F. Moura, Stefan Lee, Dhruv Batra. Natural Language Does Not Emerge 'Naturally' in Multi-Agent Dialog. *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2017.
(Best Paper Award; Selection Rate: 4/1500 = 0.26%)
29. Mike Lewis, Denis Yarats, Yann N. Dauphin, Devi Parikh, Dhruv Batra. Deal or No Deal? End-to-End Learning for Negotiation Dialogues. *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2017.
(Oral; Selection Rate: 113/1500 = 7.5%)
30. Aroma Mahendru, Viraj Prabhu, Akrit Mohapatra, Dhruv Batra, Stefan Lee. The Promise of Premise: Harnessing Question Premises in Visual Question Answering. *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2017.
31. Alexander H. Miller, Will Feng, Adam Fisch, Jiasen Lu, Dhruv Batra, Antoine Bordes, Devi Parikh, Jason Weston. ParlAI: A Dialog Research Software Platform. *Conference on Empirical Methods in Natural Language Processing (EMNLP) System Demonstrations Track*, 2017.
32. Viraj Prabhu, Prithvijit Chattopadhyay, Deshraj Yadav, Arjun Chandrasekaran, Abhishek Das, Stefan Lee, Dhruv Batra and Devi Parikh. *Evaluating Visual Dialog Agents via Cooperative Human-AI Games*. AAAI *Conference on Human Computation and Crowdsourcing (HCOMP)*, 2017.
33. Abhishek Das, Satwik Kottur, Khushi Gupta, Avi Singh, Deshraj Yadav, José M. F. Moura, Devi Parikh, Dhruv Batra. Visual Dialog. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
(Spotlight; Selection Rate: top-8% of 2680 submissions)
34. Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, Devi Parikh. Making the V in VQA Matter: Elevating the Role of Image Understanding in Visual Question Answering. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
35. Prithvijit Chattopadhyay, Ramakrishna Vedantam, Ramprasaath R. Selvaraju, Dhruv Batra, Devi Parikh. Counting Everyday Objects in Everyday Scenes. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
(Spotlight; Selection Rate: top-8% of 2680 submissions)
36. Qing Sun, Stefan Lee, Dhruv Batra. Bidirectional Beam Search: Forward-Backward Inference in Neural Sequence Models for Fill-in-the-Blank Image Captioning. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
37. Jianwei Yang, Anitha Kannan, Dhruv Batra, Devi Parikh. LR-GAN: Layered Recursive Generative Adversarial Networks for Image Generation. *International Conference on Learning Representations (ICLR)*, 2017.
38. Stefan Lee, Senthil Purushwalkam, Michael Cogswell, Viresh Ranjan, David Crandall, Dhruv Batra. Stochastic Multiple Choice Learning for Training Diverse Deep Ensembles. *Neural Information Processing Systems (NIPS)*, 2016.
39. Jiasen Lu, Jianwei Yang, Dhruv Batra, Devi Parikh. Hierarchical Question-Image Co-Attention for Visual Question Answering. *Neural Information Processing Systems (NIPS)*, 2016.

40. Harsh Agrawal*, Arjun Chandrasekaran*, Dhruv Batra, Devi Parikh, Mohit Bansal. Sort Story: Sorting Jumbled Images and Captions into Stories. *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2016.
41. Aishwarya Agrawal, Dhruv Batra, Devi Parikh. Analyzing the Behavior of Visual Question Answering Models. *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2016.
42. Arijit Ray, Gordon Christie, Mohit Bansal, Dhruv Batra, Devi Parikh. Question Relevance in VQA: Identifying Non-Visual And False-Premise Questions. *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2016.
43. Abhishek Das*, Harsh Agrawal*, C. Lawrence Zitnick, Devi Parikh, Dhruv Batra. Human Attention in Visual Question Answering: Do Humans and Deep Networks Look at the Same Regions? *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2016.
44. Gordon Christie*, Ankit Laddha*, Aishwarya Agrawal, Stanislaw Antol, Yash Goyal, Kevin Kochersberger, Dhruv Batra. Resolving Language and Vision Ambiguities Together: Joint Segmentation & Prepositional Attachment Resolution in Captioned Scenes. *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2016.
45. Timothy J O'Shea, Latha Pemula, Dhruv Batra, T. Charles Clancy. Radio Transformer Networks: Attention Models for Learning to Synchronize in Wireless Systems. *Asilomar Conference on Signals, Systems and Computers*, 2016.
46. Neelima Chavali*, Harsh Agrawal*, Aroma Mahendru*, Dhruv Batra. Object-Proposal Evaluation Protocol is 'Gameable'. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
(Spotlight; Acceptance Rate: 206/2123 = 9.7%)
47. Arjun Chandrasekaran, Ashwin K Vijayakumar, Stanislaw Antol, Mohit Bansal, Dhruv Batra, C. Lawrence Zitnick, Devi Parikh. We Are Humor Beings: Understanding and Predicting Visual Humor. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
(Spotlight; Acceptance Rate: 206/2123 = 9.7%)
48. Peng Zhang*, Yash Goyal*, Douglas Summers-Stay, Dhruv Batra, Devi Parikh. Yin and Yang: Balancing and Answering Binary Visual Questions. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
49. Jianwei Yang, Devi Parikh, Dhruv Batra. Simultaneously Discovering Image Clusters and Deep Representations. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
50. Ting-Hao Huang, Francis Ferraro, Nasrin Mostafazadeh, Ishan Misra, Jacob Devlin, Aishwarya Agrawal, Ross Girshick, Xiaodong He, Pushmeet Kohli, Dhruv Batra, Larry Zitnick, Devi Parikh, Lucy Vanderwende, Michel Galley, Margaret Mitchell. Visual Storytelling. *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*, 2016.
51. Nasrin Mostafazadeh, Nathanael Chambers, Xiadong He, Devi Parikh, Dhruv Batra, Lucy Vanderwende, Pushmeet Kohli, James Allen. A Corpus and Evaluation Framework for Deeper Understanding of Commonsense Stories. *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*, 2016.
(Oral)
52. Michael Cogswell, Faruk Ahmed, Ross Girshick, Larry Zitnick, Dhruv Batra. Reducing Overfitting in Deep Networks by Decorrelating Representations. *International Conference on Learning Representations (ICLR)*, 2016.
53. Ratnesh Kumar, Dhruv Batra. Pose Tracking by Efficiently Exploiting Global Features. *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2016.

-
54. Dustin Bales, Pablo Tarazaga, Mary Kasarda, Dhruv Batra. Gender Classification using Vibration Measurements. *IMAC Conference on Human Induced Vibrations*, 2016. Society for Experimental Mechanics.
 55. Qing Sun, Dhruv Batra. SubmodBoxes: Near-Optimal Search for a Set of Diverse Object Proposals. *Neural Information Processing Systems (NIPS)*, 2015.
 56. Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, Devi Parikh. VQA: Visual Question Answering. *International Conference on Computer Vision (ICCV)*, 2015.
 57. Faruk Ahmed, Daniel Tarlow, Dhruv Batra. Optimizing expected Intersection-over-Union with Candidate-Constrained CRFs. *International Conference on Computer Vision (ICCV)*, 2015.
 58. Qing Sun, Ankit Laddha, Dhruv Batra. Active Learning for Structured Probabilistic Models with Histogram Approximation. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
(Oral; Acceptance Rate: 71/2151 = 3.3%)
 59. Clint Solomon Mathialagan, Andrew Gallagher, Dhruv Batra. VIP: Finding Important People in Group Photographs. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
 60. Adarsh Prasad, Stefanie Jegelka, Dhruv Batra. Submodular meets Structured: Finding Diverse Subsets in Exponentially-Large Structured Item Sets. *Neural Information Processing Systems (NIPS)*, 2014.
(Spotlight; Acceptance Rate: 82/1678 = 4.9%)
 61. Vittal Premachandran, Daniel Tarlow, Dhruv Batra. Empirical Minimum Bayes Risk Prediction: How to extract an extra few% performance from vision models with just three more parameters. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.
 62. Kun Duan, David Crandall, Dhruv Batra. Multi-modal Learning in Loosely-Organized Web Images. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.
 63. Abner Guzman-Rivera, Pushmeet Kohli, Dhruv Batra, Rob Rutenbar. Efficiently Enforcing Diversity in Multi-Output Structured Prediction. *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2014.
(Oral; Acceptance Rate 7.7%)
 64. Daozheng Chen, Dhruv Batra, William T. Freeman. Group Norm for Learning Structured SVMs with Unstructured Latent Variables. *International Conference on Computer Vision (ICCV)*, 2013.
 65. Kevin Gimpel, Dhruv Batra, Greg Shakhnarovich, Chris Dyer. A Systematic Exploration of Diversity in Machine Translation. *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2013.
 66. Payman Yadollahpour, Dhruv Batra, Greg Shakhnarovich. Discriminative Re-ranking of Diverse Segmentations. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013.
 67. Joerg Kappes, Bjoern Andres, Christoph Schnoerr, Fred Hamprecht, Sebastian Nowozin, Dhruv Batra, Jan Lellmann, Nikos Komodakis, Sungwoong Kim, Bernhard Kausler, Carsten Rother. A Comparative Study of Modern Inference Techniques for Discrete Energy Minimization Problems. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013.
(Oral; Acceptance Rate 4%)
 68. Abner Guzman-Rivera, Pushmeet Kohli, Dhruv Batra. Faster Training of Structural SVMs with Diverse M-Best Cutting-Planes. *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2013.
(Oral; Acceptance Rate 11%)
 69. Abner Guzman-Rivera, Dhruv Batra, Pushmeet Kohli. Multiple Choice Learning: Learning to Produce Multiple Structured Outputs. *Neural Information Processing Systems (NIPS)*, 2012.

70. Dhruv Batra, Payman Yadollahpour, Abner Guzman-Rivera, Greg Shakhnarovich. Diverse M-Best Solutions in Markov Random Fields. *European Conference on Computer Vision (ECCV)*, 2012.
(Oral; Acceptance Rate 2.6%)
71. Dhruv Batra. An Efficient Message-Passing Algorithm for the M-Best MAP Problem. *The Conference on Uncertainty in Artificial Intelligence (UAI)*, 2012.
(Oral; Acceptance Rate 7.8%)
72. Tanmay Verma, Dhruv Batra. MaxFlow Revisited: An Empirical Comparison of Maxflow Algorithms for Dense Vision Problems. *British Machine Vision Conference (BMVC)*, 2012.
73. Kun Duan, Dhruv Batra, David J. Crandall. A Multi-layer Composite Model for Human Pose Estimation. *British Machine Vision Conference (BMVC)*, 2012.
74. Dhruv Batra, Ashutosh Saxena. Learning the Right Model: Efficient Max-Margin Learning in Laplacian CRFs. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
75. Daniel Tarlow, Dhruv Batra, Pushmeet Kohli, Vladimir Kolmogorov. Dynamic Tree Block Coordinate Descent. *International Conference on Machine Learning (ICML)*, 2011.
(Oral)
76. Dhruv Batra, Pushmeet Kohli. Making the Right Moves: Guiding Alpha-Expansion using Local Primal-Dual Gaps. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011.
77. Andrew Gallagher, Dhruv Batra, Devi Parikh. Inference for Order Reduction in Markov Random Fields. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011.
78. Dhruv Batra, Sebastian Nowozin, Pushmeet Kohli. Tighter Relaxations for MAP-MRF Inference: A Local Primal-Dual Gap based Separation Algorithm. *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2011.
79. Adarsh Kowdle, Yao-Jen Chen, Dhruv Batra, Tsuhan Chen. Scribble Based Interactive 3D Reconstruction via Scene Co-segmentation. *IEEE International Conference on Image Processing (ICIP)*, 2011.
(Oral; Acceptance Rate 12.8%)
80. Dhruv Batra, Andrew Gallagher, Devi Parikh, and Tsuhan Chen. Beyond trees: MRF Inference via Outer-Planar Decomposition. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010.
81. Dhruv Batra, Adarsh Kowdle, Devi Parikh, Jiebo Luo, and Tsuhan Chen. iCoseg: Interactive Co-segmentation with Intelligent Scribble Guidance. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010.
82. Dhruv Batra, Devi Parikh, Adarsh Kowdle, Tsuhan Chen. Seed Image Selection in Interactive Cosegmentation. *IEEE International Conference on Image Processing (ICIP)*, 2009.
83. Dhruv Batra, Rahul Sukthankar and Tsuhan Chen. Semi-Supervised Clustering via Learnt Codeword Distances. *British Machine Vision Conference (BMVC)*, 2008.
84. Dhruv Batra, Rahul Sukthankar and Tsuhan Chen. Learning Class-Specific Affinities for Image Labelling. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2008.
85. Dhruv Batra, Girish Singhal, and Santanu Chaudhury. Gabor Filter based Fingerprint Classification using Support Vector Machines. *IEEE India Annual Conference (INDICON)*, 2004.

Workshop Papers

86. Yash Goyal, Devi Parikh, Dhruv Batra. Towards Transparent AI Systems: Interpreting Visual Question Answering Models. *Scene Understanding Workshop; Visual Question Answering Challenge workshop. Conference on Computer Vision and Pattern Recognition (CVPR), 2016.*
- Workshop on Visualization for Deep Learning. International Conference on Machine Learning (ICML), 2016.*
(Best Student Paper Award)
87. Abhishek Das*, Harsh Agrawal*, C. Lawrence Zitnick, Devi Parikh, Dhruv Batra. Human Attention in Visual Question Answering. Do Humans and Deep Networks look at the same regions? *Scene Understanding Workshop; Visual Question Answering Challenge workshop. Conference on Computer Vision and Pattern Recognition (CVPR), 2016.*
- Workshop on Human Interpretability in Machine Learning; Workshop on Visualization for Deep Learning. International Conference on Machine Learning (ICML), 2016.*
(Best Student Paper Award)
88. Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, Devi Parikh. VQA: Visual Question Answering. *Workshop on Object Understanding for Interaction. International Conference on Computer Vision (ICCV), 2015.*
(Best Poster Award)
- Language & Vision Workshop; Future of Datasets in Computer Vision; Scene Understanding Workshop. Conference on Computer Vision and Pattern Recognition (CVPR), 2015.*
89. Gordon Christie, Ankit Laddha, Aishwarya Agrawal, Stan Antol, Yash Goyal, Dhruv Batra. Holistic Scene Understanding via Multiple Structured Hypotheses from Perception Modules. *Language & Vision Workshop; SUNw: Scene Understanding Workshop. Conference on Computer Vision and Pattern Recognition (CVPR), 2015.*
90. Xiao Lin, Michael Cogswell, Devi Parikh, Dhruv Batra. Propose and Re-rank Semantic Segmentation via Deep Image Classification. *Big Vision: International Workshop on Large Scale Visual Recognition and Retrieval, IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2014.*
91. Neelima Chavali, Dhruv Batra. Object Proposals using Nonparametric Bounding Box Transfer. *SUNw: Scene Understanding Workshop, IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2014.*
92. Qing Sun, Ankit Laddha, Dhruv Batra. Active Learning for Structured Probabilistic Models. *SUNw: Scene Understanding Workshop, IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2014.*
93. Michael Cogswell, Dhruv Batra. Semantic Segmentation with Deep Learning. *SUNw: Scene Understanding Workshop, IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2014.*
94. Adarsh Prasad, Stefanie Jegelka, Dhruv Batra. Submodular Maximization and Diversity in Structured Output Spaces. *Workshop on Discrete Optimization in Machine Learning (DISCML), Neural Information Processing Systems (NIPS), 2013.*
95. Abner Guzman-Rivera, Pushmeet Kohli, Dhruv Batra. Faster Training of Structural SVMs with Diverse M-Best Cutting-Planes. *Workshop on Discrete Optimization in Machine Learning (DISCML), Neural Information Processing Systems (NIPS), 2012.*
96. Varun Ramakrishna, Dhruv Batra. Mode-Marginals: Expressing Uncertainty via Diverse M-Best Solutions. *Workshop on Perturbations, Optimization, and Statistics, Neural Information Processing Systems (NIPS), 2012.*

97. Payman Yadollahpour, Dhruv Batra, Greg Shakhnarovich. M-Best Modes: Diverse M-Best Solutions in MRFs. *Workshop on Discrete Optimization in Machine Learning (DISCML), Neural Information Processing Systems (NIPS), 2011.*
98. Daozheng Chen, Dhruv Batra, William T. Freeman, Micah K. Johnson. Group Norm for Learning Latent Structural SVMs. *Workshop on Optimization in Machine Learning (OPT), Neural Information Processing Systems (NIPS), 2011.*
99. Dhruv Batra, Greg Shakhnarovich. Similarity Sensitive Nonlinear Embeddings. *Workshop on Kernels and Distances for Computer Vision, International Conference on Computer Vision (ICCV), 2011.*
100. Adarsh Kowdle, Dhruv Batra, Wen-Chao Chen, Tsuhan Chen. iModel: Interactive Co-segmentation for Object of Interest 3D Modeling. *Workshop on Reconstruction and Modeling of Large-Scale 3D Virtual Environments, European Conference on Computer Vision (ECCV), 2010.*
101. Dhruv Batra, Tsuhan Chen. Dynamic Planar-Cuts: Efficient Computation of Min-Marginals for Outer-Planar Models. *Workshop on Discrete Optimization in Machine Learning (DISCML), Neural Information Processing Systems (NIPS), 2009.*
102. Dhruv Batra, Adarsh Kowdle, Devi Parikh, Tsuhan Chen. Cutout-Search: Putting a name to the Picture. *Workshop on Internet Vision, IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2009.*
103. Dhruv Batra, Tsuhan Chen, and Rahul Sukthankar. Space-Time Shapelets for Action Recognition. *Workshop on Motion and Video Computing (WMVC) 2008, IEEE Winter Vision Meetings.*
104. Dhruv Batra, Bart Nabbe, and Martial Hebert. An Alternative Formulation for the Five Point Relative Pose Problem. *IEEE Workshop on Motion and Video Computing (WMVC), 2007.*

Patents

- Clint Mathialagan, Dhruv Batra.
Finding Important People in Images
U.S. Patent Application No: 62/169,634; Filing Date: June 2, 2015
- Jiebo Luo, Dhruv Batra, Andrew Gallagher.
Method for Generating Object Cutout for Topically Related Photographs.
Application number: 12/397,547; Publication number: US 2010/0226566 A1; Filing date: Mar 4, 2009.

Demos

- Deshraj Yadav, Viraj Prabhu, Prithvijit Chattopadhyay, Abhishek Das, Stefan Lee, Devi Parikh, Dhruv Batra.
Dialog agents that can see and human-AI GuessWhich games.
Demo at IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017.
- Harsh Agrawal, Aishwarya Agrawal, Jiasen Lu, Deshraj Yadav, Akrit Mohapatra, Devi Parikh, Dhruv Batra.
CloudCV Visual Question Answering (VQA).
Exhibition at the GPU Technology Conference (GTC), 2016.
Demo at IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016.
- Harsh Agrawal, Neelima Chavali, Clint Solomon, Akrit Mohapatra, Dhruv Batra.
CloudCV: Computer Vision As A Cloud Service.
Demo at IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2014.

- Adarsh Kowdle, Haochen Liu, ShaoYou Hsu, Jason Lew, Charvi Puri, Dhruv Batra, Tsuhan Chen.
iModel: Object of Interest 3D Modeling via Interactive Co-segmentation on a Mobile Device.
Demo at IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2012.
- Dhruv Batra, Adarsh Kowdle, Kevin Tang, Devi Parikh, Jiebo Luo, Tsuhan Chen.
Interactive Cosegmentation by Touch.
Demo at IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2009.

INVITED TALKS (NOT INCLUDING CONFERENCE PRESENTATIONS)

- *A-STAR: Agents that See, Talk, Act, and Reason*
 - Ⓜ CVPR Workshop Benchmarks for Deep Learning in Robotic Vision, Jun 2018.
 - Ⓜ F8 (Facebook Developer Conference), May 2018.
- *A Tale of Two Negative Results*
 - Ⓜ CVPR Workshop on Negative Results, Jul 2017.
- *Visual Dialog: Towards AI agents that can see, talk, and act*
 - Ⓜ SoCal NLP Symposium, Apr 2018
 - Ⓜ Computer Vision Workshop at the Amazon ML Conference, Apr 2018.
 - Ⓜ NVIDIA GPU Technology Conference, Mar 2018.
 - Ⓜ Xerox PARC, Feb 2018.
 - Ⓜ IIT Gandhinagar, Jan 2017.
 - Ⓜ Keynote Talk at the Dialog State Tracking Challenge Workshop, Dec 2017.
 - Ⓜ School of Computer Science Seminar, Georgia Tech, Nov 2017.
 - Ⓜ ICCV Workshop on Workshop on Closing the Loop Between Vision and Language, Oct 2017.
 - Ⓜ Facebook Faculty Summit, Oct 2017.
 - Ⓜ ICML Workshop on Deep Structured Prediction, Aug 2017.
 - Ⓜ CVPR Workshop on Vision and Language, Jul 2017.
 - Ⓜ Reinforcement Learning Summer School (RLSS), Jul 2017.
- *Visual Dialog*
 - Ⓜ eBay, Mar 2017.
 - Ⓜ Stanford University, Feb 2017.
 - Ⓜ Zoox, Jan 2017.
 - Ⓜ Mysore Park Workshop on Vision, Language and AI, Dec 2016.
 - Ⓜ Google Research, Dec 2016.
- *Towards Explainable VQA Models and AI Systems*
 - Ⓜ NIPS Workshop on Interpreting, Explaining and Visualizing Deep Learning, Dec 2017.
 - Ⓜ ICML Workshop on Visualization in Deep Learning, Aug 2017.
 - Ⓜ ECCV 2016 Workshop on Biological and Artificial Vision, Oct 2016.
 - Ⓜ Google DeepMind, Jul 2016.
 - Ⓜ International Computer Vision Summer School, Jul 2016.
 - Ⓜ Data Science Summit, Jul 2016.
 - Ⓜ CVPR 2016 Workshop on Large Scale Visual Recognition and Retrieval (BigVision), Jul 2016.
- *Diversity in Structured Prediction*
 - Ⓜ EMNLP16 Structured Prediction for Natural Language Processing, Nov 2016.
- *Towards Transparent Intelligent Systems: Diverse Predictions from Perception Modules*
 - Ⓜ Georgia Tech, Apr 2016.
 - Ⓜ Facebook AI Research, Apr 2016.
 - Ⓜ UT-Austin, Apr 2016.

- Ⓒ TTI-Chicago, Apr 2016.
- Ⓒ Microsoft Research Redmond, Mar 2016.
- Ⓒ Army Research Lab, Mar 2016.
- *Visual Question Answering (VQA)*
 - Ⓒ TTI-Chicago, Mar 2016.
 - Ⓒ Multimodal Machine Learning Workshop at NIPS, Dec 2015.
 - Ⓒ Plenary Talk, Western New York Image and Signal Processing Workshop, Dec 2015.
 - Ⓒ NYU CS Seminar, Sep 2015.
 - Ⓒ Cornell Tech, Sep 2015.
 - Ⓒ Google Research, Aug 2015.
 - Ⓒ Baidu, Aug 2015.
 - Ⓒ Data Science Summit & Dato Conference, Jul 2015.
 - Ⓒ Deep Learning Summit, May 2015.
- *Submodular meets Structured: Finding Diverse Subsets in Exponentially-Large Structured Item Sets*
 - Ⓒ NIPS Workshop on Discrete and Combinatorial Problems in Machine Learning (DISCML), Dec 2014.
- *M-Best and Diverse M-Best MAP Inference in Graphical Models*
 - Ⓒ CVPR 2014 Tutorial on Learning and Inference in Discrete Graphical Models, Jun 2014.
- *CloudCV: Large-Scale Distributed Computer Vision as a Cloud Service*
 - Ⓒ Invited Talk at GPU Technology Conference (GTC), Mar 2015.
 - Ⓒ NVIDIA GTC Express Webinar, Feb 2015.
 - Ⓒ Invited Talk at the DIMACS Workshop on Systems and Analytics of Big Data, Mar 2014.
- *Should we care about MAP Inference? MAP Inference tools for more than MAP Inference*
 - Ⓒ Workshop on Inference for Probabilistic Graphical Models at ICCV, Dec 2013.
- *Hedging Against Uncertainty via Multiple Diverse Solutions*
 - Ⓒ Machine Learning Seminar, University of Toronto, Apr 2015.
 - Ⓒ VASC Seminar, CMU, Mar 2015.
 - Ⓒ Indian Institute of Science, Bangalore, Dec 2014.
 - Ⓒ ONR Workshop on Structured Learning for Scene Understanding, Oct 2014.
 - Ⓒ Microsoft Research Cambridge, Sep 2014.
 - Ⓒ University of Oxford, Sep 2014.
 - Ⓒ IBM T. J. Watson Research Center, Aug 2014.
 - Ⓒ Workshop on Graphical models for Scene Understanding at ICCV, Dec 2013.
 - Ⓒ NICTA/ANU Machine Learning Seminar, Nov 2013.
 - Ⓒ eBay Research Labs, Aug 2013.
 - Ⓒ Amazon, Jul 2013.
 - Ⓒ Vision Reading Group, University of Washington, Jul 2013.
 - Ⓒ University of Maryland, Mar 2013.
 - Ⓒ IST Austria Symposium on Computer Vision and Machine Learning, Oct 2012.
- *Structured-Output Models for Computer Vision*
 - Ⓒ Google Research Tech Talk, Apr 2012.
 - Ⓒ University of Minnesota CSE Colloquium, Apr 2012.
 - Ⓒ Virginia Tech ECE Colloquium, Apr 2012.
 - Ⓒ Mitsubishi Electric Research Lab (MERL), Apr 2012.
 - Ⓒ Adobe Creative Technologies Lab Seattle, Feb 2012.
 - Ⓒ Microsoft Research Redmond, Feb 2012.
 - Ⓒ Michigan State University CSE Colloquium, Feb 2012.
 - Ⓒ Washington University in St. Louis CSE Colloquium, Feb 2012.
- *The M-Best Mode Problem*
 - Ⓒ Midwest Vision Workshop, UIUC, Sep 2012.

- Ⓒ Carnegie Mellon University VASC Seminar, May 2012.
 - Ⓒ Carnegie Mellon University Select Lab, May 2012.
 - Ⓒ University of California Berkeley Vision Seminar, Jan 2012.
- *Focused Inference and the M-Best Mode Problem*
 - Ⓒ University of California Santa Barbara CS/ECE Colloquium, Jan 2012.
 - Ⓒ University of California San Diego AI Seminar, Jan 2012.
 - Ⓒ University of California Irvine ICS, Jan 2012.
- *Focused Inference with Local Primal-Dual Gaps*
 - Ⓒ Cornell University CS Dept. Vision Seminar, Oct 2011.
 - Ⓒ Workshop on Inference in Graphical Models with Structured Potentials at CVPR, Jun 2011.
 - Ⓒ Brown University Computer Science Colloquium, May 2011.
 - Ⓒ Midwest Computer Vision Workshop at the University of Michigan, Ann Arbor, May 2011.
- *On Graph-Structured Discrete Labelling Problems in Computer Vision*
 - Ⓒ HP Labs Paulo Alto, Jun 2010.
 - Ⓒ UIUC CS Department, Jun 2010.
 - Ⓒ TTIC-Colloquium, May 2010.
 - Ⓒ Georgia Tech, May 2010.
 - Ⓒ UT-Austin CS Colloquium, May 2010.
 - Ⓒ MIT CSAIL, Apr 2010.
 - Ⓒ University of Georgia CS Colloquium, Apr 2010.
- *Beyond Trees: MRF Inference via Outer-Planar Decomposition (OPD)*
 - Ⓒ Carnegie Mellon University VASC Seminar, Mar 2010.
 - Ⓒ Illinois Vision Meet, TTIC, Dec 2009.
 - Ⓒ Toyota Technical Institute Chicago, Dec 2009.
 - Ⓒ Microsoft Research Redmond, Dec 2009.

PRESS COVERAGE

- *How A Virtual Scavenger Hunt Could Train Robots To Find Things In Your Home*
 - [Fast Company](#), [DigitalJournal](#)
- *Facebook built an AI system that learned to lie to get what it wants*
 - 50+ articles on our research on AI agents that are trained to negotiate via deep reinforcement learning
 - [CNN](#), [Fast Company](#), [TechCrunch](#), [Quartz](#), [The Verge](#), and [New Scientist](#).
- *Four Indian American Researchers Named Young Investigator Awardees for 2017 by Office of Naval Research*
 - IndiaWest
 - http://www.indiawest.com/news/global_indian/four-indian-american-researchers-named-young-investigator-awardees-for-by/article_dd10650a-4632-11e7-8f0c-5f1718addb51.html
- *Is Artificial Intelligence Permanently Inscrutable?*
 - Nautilus
 - <http://nautil.us/issue/40/learning/is-artificial-intelligence-permanently-inscrutable>
- *Deep learning is creating computer systems we don't fully understand*
 - The Verge
 - <http://www.theverge.com/2016/7/12/12158238/first-click-deep-learning-algorithmic-black-boxes>
- *Robot eyes and humans fix on different things to decode a scene*
 - Tech Radar

- <http://www.techradar.com/news/world-of-tech/robots-and-humans-see-the-world-differently-but-we-don-t-know-why-1324165>
- New Scientist
 - <https://www.newscientist.com/article/2095616-robot-eyes-and-humans-fix-on-different-things-to-decode-a-scene/>
- MIT Technology Review
 - <https://www.technologyreview.com/s/601819/ai-is-learning-to-see-the-world-but-not-the-way-humans-do/>
- *Teaching computers to describe images as people would*
 - Live Science
 - <http://www.livescience.com/54961-artificial-intelligence-tells-stories-from-photos.html>
 - MIT Technology Review
 - <https://www.technologyreview.com/s/601339/will-artificial-intelligence-win-the-caption-contest/>
 - Microsoft Blog
 - <http://blogs.microsoft.com/next/2016/04/14/teaching-computers-to-describe-images-as-people-would/#sm.000gnqgev13n2e4q106e2cq6togzw>
 - Venture Beat
 - <http://venturebeat.com/2016/04/14/microsoft-ai-visual-storytelling/>
- *A Giant Leap for Machine Kind; When Robots Can See*
 - Interview with WVTF Radio IQ
 - <http://wvtf.org/post/giant-leap-machine-kind-when-robots-can-see>
- *Coding jokes: Virginia Tech research team tackles the algorithm of humor*
 - Virginia Tech ECE News
 - <https://www.ece.vt.edu/news/articles/coding-jokes-virginia-tech-research-team-tackles-the-algorithm-of-humor.html>
- *Artificial Intelligence Algorithm Taught to Recognise Humor*
 - Newsweek
 - <http://www.newsweek.com/artificial-intelligence-algorithm-taught-recognise-humor-413832?rx=us>
- *AI Algorithm Identifies Humorous Pictures*
 - MIT Technology Review
 - <http://www.technologyreview.com/view/545316/ai-algorithm-identifies-humorous-pictures/>
- *What's in This Picture? AI Becomes as Smart as a Toddler*
 - Bloomberg
 - <http://www.bloomberg.com/news/articles/2015-05-22/what-s-in-this-picture-ai-becomes-as-smart-as-a-toddler>
- *How to automatically detect the most important people in a photograph*
 - The Boston Globe
 - <http://www.bostonglobe.com/ideas/2015/04/01/how-automatically-detect-most-important-people-photograph/tZND3z3epWTJu4Gvf9FSRN/story.html>
- *The Curious Case of Apple's Supposed Self-Driving Car*
 - WIRED, Feb 2015
 - <http://www.wired.com/2015/02/apple-self-driving-car/>
 - Business Insider, Feb 2015
 - <http://www.businessinsider.sg/cameras-atop-mystery-apple-van-could-be-self-driving-car-2015-2/#.VPHvGmYXMqo>
 - Realty Today, Feb 2015
 - <http://www.realtytoday.com/articles/10175/20150210/apple-news-2015-icar-under-development-yes-employee.htm>
 - Opp Trends, Feb 2015
 - <http://www.opptrends.com/2015/02/is-apple-inc-aapl-trying-to-rival-tesla-motors-inc-tsla-on-driverless-car-technology/>
 - Rt.com, Feb 2015
 - <http://rt.com/usa/230795-apple-camera-van-tesla/>
- *He aims to mimic the human brain's capabilities*
 - India Abroad
 - <http://www.indiaabroad-digital.com/indiaabroad/20140912?pg=42#pg42>

- *Dhruv Batra seeks to remove ambiguity in computer visual recognition systems*
 - Wn.com, Aug 2014
 - http://article.wn.com/view/2014/08/11/Dhruv_Batra_seeks_to_remove_ambiguity_in_computer_visual_rec/
 - HighBeam Reseach, Aug 2014
 - <http://www.highbeam.com/doc/1G1-377973708.html>
 - Veooz, Aug 2014
 - <https://www.veooz.com/news/AHPMYiC.html>
 - Phys.Org, Aug 2014
 - <http://phys.org/wire-news/169195840/dhruv-batra-seeks-to-remove-ambiguity-in-computer-visual-recogni.html>
 - Virginia Tech News, Aug 2014
 - <http://vtnews.vt.edu/articles/2014/08/081114-engineering-dhruvbatraawards.html>
 - Virginia Tech College of Engineering News, Aug 2014
 - <http://www.eng.vt.edu/news/college-engineering-researcher-seeks-remove-ambiguity-computer-visual-recognition-systems>
- *Dhruv Batra receives NSF CAREER Award for machine perception research.*
 - Discovery Analytic Center, May 2014
 - <http://dac.cs.vt.edu/wp/?p=293>
 - Virginia Tech ECE News, May 2014
 - <http://www.ece.vt.edu/news/articles/batra-CAREER.html>
- *Stocking the visual toolbox*
 - Virginia Tech ECE News, May 2013
 - http://www.ece.vt.edu/news/ar13/visual_toolbox.php

PROFESSIONAL SERVICE ACTIVITIES

Editorial Boards

- Associate Editor, Springer's Machine Vision & Applications Journal
- Guest Editor, Transactions on Pattern Analysis and Machine Intelligence (PAMI).
Special Issue on "Higher Order Graphical Models in Computer Vision: Modelling, Inference & Learning."
Together with Karteek Alahari, Srikumar Ramalingam, Nikos Paragios, and Rich Zemel.

Conference Chairs

- Area Chair, European Conference on Computer Vision (ECCV), 2018.
- Area Chair, Neural Information Processing Systems (NIPS), 2017, 2018.
- Area Chair, International Conference on Learning Representations (ICLR), 2017 – 2019.
- Area Chair, IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016 – 2018.
- Area Chair, Conference on Uncertainty in Artificial Intelligence (UAI), 2017.
- Program Chair, Indian Conference on Computer Vision, Graphics, and Image Processing (ICVGIP), 2016.
- Doctoral Consortium Chair, International Conference on Computer Vision (ICCV), 2015.
- Area Chair, Indian Conference on Computer Vision, Graphics and Image Processing (ICVGIP), 2014.

Organizer

- Workshop on Visually-Grounded Interaction and Language (ViGIL)
Neural Information Processing Systems (NIPS), 2017.
Co-organized with Florian Strub, Harm de Vries, Abhishek Das, Satwik Kottur, Stefan Lee, Mateusz Malinowski, Olivier Pietquin, Devi Parikh, Dhruv Batra, Aaron Courville, Jeremie Mary.
- 4th Workshop on Deep Vision
IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017.
Co-organized with Jose M. Alvarez, Nathan Silberman, Yann LeCun.

- 2nd Workshop/Challenge on Visual Question Answering
IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017.
Co-organized with Aishwarya Agrawal, Yash Goyal, Tejas Khot, Peng Zhang, Jiasen Lu, Larry Zitnick, and Devi Parikh.
- 1st Workshop/Challenge on Visual Question Answering
IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016.
Co-organized with Aishwarya Agrawal, Jiasen Lu, Yash Goyal, Peng Zhang, Larry Zitnick, and Devi Parikh.
- Tutorial on “Diversity meets Deep Networks”
IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016.
Co-organized with Stefan Lee, Alexander Kirillov, Bogdan Savchynskyy, and Carsten Rother.
- The 1st Mid-Atlantic Computer Vision (MACV) Workshop,
Virginia Tech, April 18, 2014.
Co-organized with Devi Parikh. (~115 attendees)
- Tutorial on “Learning and Inference in Discrete Graphical Models”.
IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2014.
Co-organized with Nikos Komodakis, Nikos Paragios and Stephen Gould.
- Tutorial on “Beyond MAP: Making Multiple Predictions from Probabilistic Models: Diversity, DPPs and more.”
IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2013.
Co-organized with Alex Kulesza, Dennis Park, and Deva Ramanan.
- Workshop on “Higher-Order Models and Global Constraints”.
European Conference on Computer Vision (ECCV), 2012.
Co-organized with Karteek Alahari, Srikumar Ramalingam, Nikos Paragios, and Rich Zemel.
- Workshop on “Supervised Large-Scale Learning of Similarity”.
Neural Information Processing Systems (NIPS), 2011.
Co-organized with Greg Shakhnarovich, Brian Kulis, and Kilian Weinberger.

Conference Program Committee / Reviewer

- IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2011-15.
- International Conference on Computer Vision (ICCV), 2009, 2011, 2013, 2015.
- European Conference on Computer Vision (ECCV), 2010, 2012, 2014, 2016.
- Neural Information Processing Systems (NIPS), 2011-16.
- International Conference on Machine Learning (ICML), 2012-15.
- Uncertainty in Artificial Intelligence (UAI), 2012-14.
- International Joint Conference on Artificial Intelligence (IJCAI), 2011.
- Energy Minimization Methods in Computer Vision and Pattern Recognition (EMMCVPR), 2011, 2013.

Workshop Program Committees

- ECCV 2014 Workshop on Graphical Models in Computer Vision.
- ICCV 2013 Workshop on Graphical Models for Scene Understanding: Challenges & Perspectives (GMSU).
- ICCV 2013 Workshop on Understanding Human Activities: Context and Interaction (HACI).
- CVPR 2013 Workshop on Structured Prediction: Tractability, Learning, and Inference.
- ICML 2013 Workshop on Inferning: Interactions between Inference & Learning.
- ECCV 2012 Workshop on Parts and Attributes.
- NIPS 2010 Workshop on Optimization in Machine Learning.

Panelist

- NSF Review Panel, CISE Directorate, 2016.
- NSF Review Panel, Information and Intelligent Systems (IIS) Division, 2015.
- VT Institute for Critical Technology and Applied Science (ICTAS) Black Swan Seminar on “Big Data Forecasting”, Apr 2014.
- NIPS 2012 Workshop Panel Discussion on Perturbation, Optimization, and Statistics.

Reviewer for Journals

- IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI).
- International Journal of Computer Vision (IJCV).

Departmental Service

- Faculty Search Committee (Machine Learning), ECE Department, Virginia Tech (2015 – 2016)
- Faculty Search Committee (Cyber-Physical Systems), ECE Department, Virginia Tech (2014 – 2015)
- Departmental Computing Committee, ECE Department, Virginia Tech (Fall 2013 – Present)
- Faculty Search Committee (Machine Learning), Computer Science Department, Virginia Tech (2013 – 2014)