

DHRUV BATRA

CURRENT POSITION

Senior Director Aug 2021 – Present
Fundamental AI Research (FAIR), Meta
Lead: Embodied AI team

Associate Professor Fall 2019 – Present
Lead: Machine Learning & Perception Group
School of Interactive Computing
Georgia Institute of Technology

Research Interests: Artificial Intelligence, Machine Learning, Computer Vision

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 Paper Award** 2023
Emergence of Maps in the Memories of Blind Navigation Agents.
International Conference on Learning Representations (ICLR), 2023.

One of 4 outstanding papers out of 4900 submissions (or top 0.08%) to ICLR 2023, one of the premier publication venues for research in machine learning.

• **Best Paper Award Finalist** 2022
Ego4D: Around the World in 3,000 Hours of Egocentric Video
IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2022.

One of 33 papers nominated for best paper awards out of 8161 submissions (or top 0.4%) to CVPR 2022, one of the premier publication venues for research in computer vision.

• **Presidential Early Career Award for Scientists and Engineers (PECASE)** 2019

PECASE is the highest honor bestowed by the U.S. government to outstanding scientists and engineers who show exceptional promise for leadership in science and technology. PECASE is awarded by the President of the United States and administered by the White House Office of Science and Technology Policy following nominations from participating agencies like the National Science Foundation and Department of Defense.

• **Best Paper Nomination** 2019
Habitat: A Platform for Embodied AI Research

International Conference on Computer Vision (ICCV), 2019.

One of 11 papers nominated for best paper awards out of 4303 submissions (or top 0.25%) to ICCV 2019, one of the premier publication venues for research in computer vision.

- **Early Career Award for Scientists and Engineers (ECASE-Army)**, Army Research Office 2018
The ECASE-Army is modeled after the Presidential Early Career Award for Scientists and Engineers (PECASE). ECASE-Army is awarded by the Army Research Office (ARO) following a rigorous selection process among all ARO Young Investigator award recipients across all areas of science and engineering. ECASE-Army award is supported by \$1M in new funding.
- **Outstanding Junior Faculty Research award** 2018
College of Computing, Georgia Tech
Awarded the by Dean and Awards Committee, 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** 2017
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
"[T]he 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 whose scientific pursuits show outstanding promise for supporting the Department of Defense, while also promoting their professional development."
- **Amazon Academic Research Award** 2016
- **Best Student Paper Award** 2016
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** 2016
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."
- **Best Poster Paper Award** 2015
Visual Question Answering, Antol et al., Workshop on Object Understanding for Interaction.
International Conference on Computer Vision (ICCV), 2015.
- **Outstanding Reviewer Award** 2015
Conference on Computer Vision and Pattern Recognition (CVPR), 2015

- **Outstanding New Assistant Professor award** 2015
 College of Engineering, Virginia Tech

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 2014
 Scholarship Recognition
- Virginia Tech Center for Instructional Development and Education Research 2013
 Teaching Recognition

PAST APPOINTMENTS

- Research Director** Aug 2021 – Jan 2024
 Fundamental AI Research (FAIR), Meta
Lead: Embodied AI team
- Research Scientist** Aug 2017 – Aug 2021
 Facebook AI Research (FAIR)
- Assistant Professor** Fall 2016 – Spring 2019
 School of Interactive Computing
 Georgia Institute of Technology
- Visiting Researcher** Fall 2016 – Summer 2017
 Facebook AI Research (FAIR)
- Assistant Professor** Jan 2013 – Aug 2016
 Bradley Department of Electrical and Computer Engineering, Virginia Tech
Lead: Machine Learning & Perception Group @ VT
- Visiting Researcher** May – Aug 2015
 Microsoft Research, Redmond, WA
 Hosts: Larry Zitnick, Meg Mitchell
- Research Assistant Professor** Sep 2010 – Dec 2012
 Toyota Technological Institute at Chicago
- Visiting Faculty**, Carnegie Mellon University May – Aug 2012
 Host: Carlos Guestrin

Visiting Scientist, CSAIL Massachusetts Institute of Technology
Host: Bill Freeman

Mar – May 2011

TEACHING

Fall 2018, 2019, 2020, 2021	GT CS 4803/7643 Deep Learning
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

RESEARCH GROUP AND STUDENT ADVISING

Students

- Arjun Majumdar PhD student (2019 – Present)
- Joanne Truong PhD student (2019 – Present), Co-advised with Sonia Chernova
Apple PhD Fellowship
Adobe Fellowship
NSF Graduate Research Fellowship
Google Women Techmakers Scholarship
National Defense Science and Engineering Graduate Fellowship (declined)
- Andrew Szot PhD student (2020 – Present); Co-advised with Zsolt Kira
- Naoki Yokoyama PhD student (2020 – Present); Co-advised with Sehoon Ha
ARCS Scholar Award Fellowship '22
- Ram Ramrakhya PhD student (2023 – Present); Co-advised with Zsolt Kira
- Karmesh Yadav PhD student (2023 – Present); Co-advised with Zsolt Kira

- Mukul Khanna MS student (2022 – Present)
- Gunjan Chhablani MS student (2022 – Present)

Alumni (Graduated Students & Former Interns)

- Dr. Peter Anderson Research Scientist (2018 – 2019)
(next position: Research Scientist in Language and Vision at Google)

- Dr. Zhile Ren Postdoc (2018 – 2019)
(next position: Apple)

- Dr. Stefan Lee Postdoc (2016 – 2017), Research Scientist (2017 – 2019)
(next position: Assistant Professor, Oregon State University)
Outstanding Research Scientist Award, College of Computing, Georgia Tech
DARPA Riser

- Harsh Agrawal PhD, May 2023 (next position: Apple)
Snap Inc. Research Fellowship

- Nirbhay Modhe PhD, Dec 2022
- Erik Wijmans PhD, Aug 2022 (next position: Apple); Co-advised with Irfan Essa
ARCS Scholar Award Fellowship ('19, '20, '21)
Georgia Tech College of Computing Dissertation Award, 2023
Sigma Xi Best Ph.D. Thesis Award, 2023
- Ashwin Kalyan PhD, Jul 2020 (next position: Allen Institute for Artificial Intelligence)
JP Morgan PhD Fellowship (declined)
- Abhishek Das PhD, Mar 2020 (next position: Facebook AI Research)
Facebook Fellowship
Microsoft Research Ph.D. Fellowship (declined)
NVIDIA Graduate Fellowship (declined)
Adobe Research Fellowship
Snap Inc. Research Fellowship
Georgia Tech College of Computing Dissertation Award, 2021
Sigma Xi Best Ph.D. Thesis Award, 2021
AAAI/ACM-SIGAI Doctoral Dissertation Award Honorable Mention, 2020
- Michael Cogswell PhD, Mar 2020 (next position: SRI)
Bradley Fellowship
- Yash Goyal PhD, Mar 2020 (next position: Samsung AI Lab)
- Aishwarya Agrawal PhD, Aug 2019 (next position: Google DeepMind)
Google Fellowship (declined)
Facebook Fellowship (declined)
NVIDIA Graduate Fellowship
Sigma Xi Best Ph.D. Thesis Award, 2020
Georgia Tech College of Computing Dissertation Award, 2020
AAAI/ACM-SIGAI Doctoral Dissertation Award Honorable Mention, 2019
- Qing Sun PhD, Nov 2017 (next position: Facebook AI)
- Gordon Christie PhD, Nov 2016 (next position: Army Research Lab)
- Rishabh Jain MS 2021 (next position: eBay)
Marshall D. Williamson Fellowship award, 2021
- Joel Ye MS 2021 (next position: PhD student at CMU)
Donald V. Jackson Fellowship award, 2021
- Sameer Dharur MS 2021 (next position: Apple)
- Deshraj Yadav MS, Dec 2018 (next position: Tesla Autopilot)
Snap Inc. Research Scholarship, 2017
- Akrit Mohapatra MS, Apr 2018 (next position: Ebay)
- Aroma Mahendru MS, Apr 2017 (next position: Bloomberg)
- Latha Pemula MS, Sep 2016 (next position: Amazon)
- Harsh Agrawal MS, May 2016 (next position: Snapchat Research)
- Prakriti Banik MS, Jul 2015 (next position: Bloomberg)
- Neelima Chavali MS, Jul 2015 (next position: Many Trees Inc)
- Clint Solomon MS, May 2015 (next position: PhotoKharma)
- Yash Kant Intern, Fall 2019 – Spring 2020 (next position: PhD student at Univ of Toronto)
- Rishabh Jain Intern, Fall 2018 – Spring 2019 (next position: MS student at Georgia Tech)
- Sanyam Agrawal Intern, Fall 2018 – Spring 2019 (next position: Google)
- Deshraj Yadav Intern, Fall 2016 – Spring 2017 (next position: MS student at Georgia Tech)

-
- Tejas Khot Intern, Fall 2016 – Spring 2017 (next position: MS student at RI, CMU)
 - Viraj Prabhu Intern, Fall 2016 – Spring 2017 (next position: MS student at Georgia Tech)
 - Avi Singh Intern, Summer 2016 (next position; PhD student at UC Berkeley)
 - Khushi Gupta Intern, Spring 2016 (next position: MS student at RI, CMU)
 - Ashwin Kalyan Intern, Fall 2015 – Spring 2016 (next position: PhD student at VT)
 - Abhishek Das Intern, Fall 2015 – Spring 2016 (next position: PhD student at VT)
 - Ahmed Osman Intern, Summer 2015 (next position: MS student at UCL)
 - Faruk Ahmed Intern, Fall 2013 – Spring 2015 (next position: PhD student at Univ of Montreal)
 - Senthil PS Intern, Summer13, Fall 2014 – Spring 2014 (next position: MS student at CMU)
 - Aroma Mahendru Intern, Summer13, Fall 2014 – Spring 2014 (next position: MS student at VT)
 - Ratnesh Kumar Intern, Fall 2015 (next position: MERL)
 - Qi Lou Intern, Fall 2013 – Spring 2014 (next position: PhD student at UC Irvine)
 - Ankit Laddha Intern, Fall 2013 – Spring 2014 (next position: MS student at RI, CMU)
 - Adarsh Prasad Intern, Fall 2013 – Fall 2014 (next position: PhD student at CMU)
 - Harsh Agrawal Intern, Summer 2013 (next position: MS student at VT)
 - Abhimanyu Dubey Intern, Summer 2013 (next position: PhD student at MIT)

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. Theophile Gervet, Soumith Chintala, Dhruv Batra, Jitendra Malik, Devendra Singh Chaplot. Navigating to objects in the real world. *Science Robotics*, Volume 8, Vol 8, Issue 79, 2023.

5. Brennan Shacklett, Luc Guy Rosenzweig, Zhiqiang Xie, Bidipta Sarkar, Andrew Szot, Erik Wijmans, Vladlen Koltun, Dhruv Batra, Kayvon Fatahalian. An Extensible, Data-Oriented Architecture for High-Performance, Many-World Simulation. *SIGGRAPH / ACM Transactions on Graphics (TOG)*, Volume 42, Issue 4, 2023.
6. William Ferguson, Dhruv Batra, Raymond Mooney, Devi Parikh, Antonio Torralba, David Bau, David Diller, Josh Fasching, Jaden Fiotto-Kaufman, Yash Goyal, Jeff Miller, Kerry Moffitt, Alex Montes de Oca, Ramprasaath R. Selvaraju, Ayush Shrivastava, Jialin Wu, Stefan Lee. *Applied AI Letters*, Volume 2, Issue 4, 2021.
7. Joanne Truong, Sonia Chernova, Dhruv Batra. Bi-directional Domain Adaptation for Sim2Real Transfer of Embodied Navigation Agents. *IEEE Robotics and Automation Letters (RA-L)*, 2021.
8. Abhishek Kadian*, Joanne Truong*, Aaron Gokaslan, Alexander Clegg, Erik Wijmans, Stefan Lee, Manolis Savva, Sonia Chernova, Dhruv Batra. Sim2Real Predictivity: Does Evaluation in Simulation Predict Real-World Performance. *IEEE Robotics and Automation Letters (RA-L)*, 2020.
9. Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, Dhruv Batra. Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization. *International Journal on Computer Vision (IJCV)*, 2019.
10. Yash Goyal, Aishwarya Agrawal, 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. *International Journal of Computer Vision (IJCV)*, 2019.
11. 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.
12. 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.
13. 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.
14. 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.
15. 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.
16. Vittal Premachandran, Daniel Tarlow, Alan Yuille, Dhruv Batra. Empirical Minimum Bayes Risk Prediction. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2016.
17. 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.
18. 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.
19. 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.

20. Kun Duan, Dhruv Batra, David Crandall. Human Pose Estimation via Multi-layer Composite Models. *Signal Processing*, 2014.
21. 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.
22. 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%)

23. Arjun Majumdar, Anurag Ajay, Xiaohan Zhang, Sriram Yenamandra, Mikael Henaff, Alexander Sax, Sneha Silwal, Paul Mccvay, Oleksandr Maksymets, Sergio Arnaud, Pranav Putta, Karmesh Yadav, Qiyang Li, Ben Newman, Mohit Sharma, Vincent-Pierre Berges, Shiqi Zhang, Pulkit Agrawal, Dhruv Batra, Yonatan Bisk, Mrinal Kalakrishnan, Franziska Meier, Chris Paxton, Aravind Rajeswaran. OpenEQA: Embodied Question Answering in the Era of Foundation Models. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2024.
24. Mukul Khanna, Yongsun Mao, Hanxiao Jiang, Sanjay Haresh, Brennan Shacklett, Dhruv Batra, Alexander Clegg, Eric Undersander, Angel X Chang, Manolis Savva. Habitat Synthetic Scenes Dataset (HSSD-200): An Analysis of 3D Scene Scale and Realism Tradeoffs for ObjectGoal Navigation. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2024.
25. Mukul Khanna, Ram Ramrakhya, Gunjan Chhablani, Sriram Yenamandra, Theophile Gervet, Matthew Chang, Zsolt Kira, Devendra Singh Chaplot, Dhruv Batra, Roozbeh Mottaghi. GOAT-Bench: A Benchmark for Multi-modal Lifelong Navigation. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2024.
26. Ram Ramrakhya, Aniruddha Kembhavi, Dhruv Batra, Zsolt Kira, Kuo-Hao Zeng, Luca Weihs. Seeing the Unseen: Visual Common Sense for Semantic Placement. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2024.
27. Xavier Puig, Eric Undersander, Andrew Szot, Mikael Dallaire Cote, Tsung-Yen Yang, Ruslan Partsey, Ruta Desai, Alexander William Clegg, Michal Hlavac, So Yeon Min, Vladimír Vondruš, Theophile Gervet, Vincent-Pierre Berges, John M. Turner, Oleksandr Maksymets, Zsolt Kira, Mrinal Kalakrishnan, Jitendra Malik, Devendra Singh Chaplot, Unnat Jain, Dhruv Batra, Akshara Rai, Roozbeh Mottaghi. Habitat 3.0: A Co-Habitat for Humans, Avatars and Robots. International Conference on Learning Representations (ICLR), 2024.
28. Sneha Silwal, Karmesh Yadav, Tingfan Wu, Jay Vakil, Arjun Majumdar, Sergio Arnaud, Claire Chen, Vincent-Pierre Berges, Dhruv Batra, Aravind Rajeswaran, Mrinal Kalakrishnan, Franziska Meier, Oleksandr Maksymets. What Do We Learn from a Large-Scale Study of Pre-Trained Visual Representations in Sim and Real Environments? IEEE International Conference on Robotics and Automation (ICRA), 2024.
29. Naoki Yokoyama, Sehoon Ha, Dhruv Batra, Jiuguang Wang, Bernadette Bucher. VLFM: Vision-Language Frontier Maps for Zero-Shot Semantic Navigation. IEEE International Conference on Robotics and Automation (ICRA), 2024.
30. Arjun Majumdar, Karmesh Yadav, Sergio Arnaud, Yecheng Jason Ma, Claire Chen, Sneha Silwal, Aryan Jain, Vincent-Pierre Berges, Pieter Abbeel, Jitendra Malik, Dhruv Batra, Yixin Lin, Oleksandr Maksymets, Aravind Rajeswaran, Franziska Meier. Where are we in the search for an Artificial Visual Cortex for Embodied Intelligence? Neural Information Processing Systems (NeurIPS), 2023.

31. Sriram Yenamandra, Arun Ramachandran, Karmesh Yadav, Austin Wang, Mukul Khanna, Theophile Gervet, Tsung-Yen Yang, Vidhi Jain, Alexander William Clegg, John Turner, Zsolt Kira, Manolis Savva, Angel Chang, Devendra Singh Chaplot, Dhruv Batra, Roozbeh Mottaghi, Yonatan Bisk, Chris Paxton. HomeRobot: Open-Vocabulary Mobile Manipulation. Conference on Robot Learning (CoRL), 2023.
32. Arjun Majumdar, Fei Xia, Brian Ichter, Dhruv Batra, Leonidas Guibas. FindThis: Language-Driven Object Disambiguation in Indoor Environments. Conference on Robot Learning (CoRL), 2023.
33. Xiaoyu Huang, Dhruv Batra, Akshara Rai, Andrew Szot. Skill Transformer: A Monolithic Policy for Mobile Manipulation. International Conference on Computer Vision (ICCV), 2023.
34. Jacob Krantz, Theophile Gervet, Karmesh Yadav, Austin Wang, Chris Paxton, Roozbeh Mottaghi, Dhruv Batra, Jitendra Malik, Stefan Lee, Devendra Singh Chaplot. Navigating to Objects Specified by Images. International Conference on Computer Vision (ICCV), 2023.
35. Andrew Szot, Unnat Jain, Dhruv Batra, Zsolt Kira, Ruta Desai, Akshara Rai. Adaptive Coordination in Social Embodied Rearrangement. International Conference on Machine Learning (ICML), 2023.
36. Karmesh Yadav, Ram Ramrakhya, Santhosh Kumar Ramakrishnan, Theo Gervet, John Turner, Aaron Gokaslan, Noah Maestre, Angel Xuan Chang, Dhruv Batra, Manolis Savva, Alexander William Clegg, Devendra Singh Chaplot. Habitat-Matterport 3D Semantics Dataset. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2023.
Highlight Paper (Selection rate: 235/9155 submissions = 2.5%)
37. Ram Ramrakhya, Dhruv Batra, Erik Wijmans, Abhishek Das. PIRLNav: Pretraining with Imitation and RL Finetuning for ObjectNav. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2023.
38. Vincent-Pierre Berges, Andrew Szot, Devendra Singh Chaplot, Aaron Gokaslan, Roozbeh Mottaghi, Dhruv Batra, Eric Undersander. Galactic: Scaling End-to-End Reinforcement Learning for Rearrangement at 100k Steps-Per-Second. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2023.
39. Erik Wijmans, Manolis Savva, Irfan Essa, Stefan Lee, Ari S. Morcos, Dhruv Batra. Emergence of Maps in the Memories of Blind Navigation Agents. International Conference on Learning Representations (ICLR), 2023.
Outstanding Paper Award (Selection rate: 4/4900 submissions = 0.08%)
40. Andrew Szot, Amy Zhang, Dhruv Batra, Zsolt Kira, Franziska Meier. BC-IRL: Learning Generalizable Reward Functions from Demonstrations. International Conference on Learning Representations (ICLR), 2023.
Notable top-25% paper (top 8% of 4900 submissions)
41. Simar Kareer*, Naoki Yokoyama*, Dhruv Batra, Sehoon Ha, and Joanne Truong. ViNL: Visual Navigation and Locomotion Over Obstacles. IEEE International Conference on Robotics and Automation (ICRA), 2023.
Best Paper Award at Learning for Agile Robotics Workshop at CoRL, 2022.
42. Jing Yu Koh, Harsh Agrawal, Dhruv Batra, Richard Tucker, Austin Waters, Honglak Lee, Yinfei Yang, Jason Baldridge, Peter Anderson. Simple and Effective Synthesis of Indoor 3D Scenes. AAAI Conference on Artificial Intelligence (AAAI), 2023.
43. Arjun Majumdar, Gunjan Aggarwal, Bhavika Devnani, Judy Hoffman, Dhruv Batra. ZSON: Zero-Shot Object-Goal Navigation using Multimodal Goal Embeddings. Neural Information Processing Systems (NeurIPS), 2022.
44. Erik Wijmans, Irfan Essa, Dhruv Batra. Variable Experience Rollout: Training Robust Skill Policies for Rearrangement. Neural Information Processing Systems (NeurIPS), 2022.

-
45. Changan Chen, Carl Schissler, Sanchit Garg, Philip Kobernik, Alexander Clegg, Paul Calamia, Dhruv Batra, Philip W Robinson, Kristen Grauman. SoundSpaces 2.0: A Simulation Platform for Visual-Acoustic Learning. Neural Information Processing Systems (NeurIPS) Track on Datasets and Benchmarks, 2022.
 46. Joanne Truong, Max Rudolph, Naoki Yokoyama, Sonia Chernova, Dhruv Batra, Akshara Rai. Sim2Real: Lower Fidelity Simulation Leads to Higher Sim2Real Transfer in Navigation. Conference on Robot Learning (CoRL), 2022.
 47. Karl Pertsch, Ruta Desai, Vikash Kumar, Franziska Meier, Joseph J Lim, Dhruv Batra, Akshara Rai. Cross-Domain Transfer via Semantic Skill Imitation. Conference on Robot Learning (CoRL), 2022.
 48. Yash Kant, Arun Ramachandran, Sriram Yenamandra, Igor Gilitschenski, Dhruv Batra, Andrew Szot, Harsh Agrawal. Housekeep: Tidying Virtual Households using Commonsense Reasoning. European Conference on Computer Vision (ECCV), 2022.
 49. Lina Mezghani, Sainbayar Sukhbaatar, Thibaut Lavril, Oleksandr Maksymets, Dhruv Batra, Piotr Bojanowski, Karteek Alahari. Memory-Augmented Reinforcement Learning for Image-Goal Navigation. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2022.
 50. Naoki Yokoyama, Qian Luo, Dhruv Batra, Sehoon Ha. Benchmarking Augmentation Methods for Learning Robust Navigation Agents: the Winning Entry of the 2021 iGibson Challenge. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2022.
 51. Samyak Datta, Sameer Dharur, Vincent Cartillier, Ruta Desai, Mukul Khanna, Dhruv Batra, Devi Parikh. Episodic Memory Question Answering. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2022.
Oral (Selection Rate: 344/8161 = 4.2%)

Kristen Grauman et al. Ego4D: Around the World in 3,000 Hours of Egocentric Video. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2022.
Best Paper Award Finalist (Selection Rate: 33/8161 = 0.4%)
 52. Ram Ramrakhya, Eric Undersander, Dhruv Batra, Abhishek Das. Habitat-Web: Learning Embodied Object-Search Strategies from Human Demonstrations at Scale. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2022.
 53. Ruslan Partsey, Erik Wijmans, Naoki Yokoyama, Oles Doboševych, Dhruv Batra, Oleksandr Maksymets. Is Mapping Necessary for Realistic PointGoal Navigation? IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2022.
 54. Erik Wijmans, Irfan Essa, Dhruv Batra. How to Train PointGoal Navigation Agents on a (Sample and Compute) Budget. International Conference on Autonomous Agents and Multiagent Systems (AAMAS), 2022.
 55. Andrew Szot, Alex Clegg, Eric Undersander, Erik Wijmans, Yili Zhao, John Turner, Noah Maestre, Mustafa Mukadam, Devendra Chaplot, Oleksandr Maksymets, Aaron Gokaslan, Vladimir Vondrus, Sameer Dharur, Franziska Meier, Wojciech Galuba, Angel Chang, Zsolt Kira, Vladlen Koltun, Jitendra Malik, Manolis Savva, Dhruv Batra. Habitat 2.0: Training Home Assistants to Rearrange their Habitat. Neural Information Processing Systems (NeurIPS), 2021.
Spotlight paper (Selection Rate: 273/9122 = 3%)
 56. Santhosh K. Ramakrishnan, Aaron Gokaslan, Erik Wijmans, Oleksandr Maksymets, Alex Clegg, John Turner, Eric Undersander, Wojciech Galuba, Andrew Westbury, Angel X. Chang, Manolis Savva, Yili Zhao, Dhruv Batra. Habitat-Matterport 3D Dataset (HM3D): 1000 Large-scale 3D Environments for Embodied AI. Neural Information Processing Systems (NeurIPS) Track on Datasets and Benchmarks, 2021.

-
57. Abhinav Moudgil, Arjun Majumdar, Harsh Agrawal, Stefan Lee, Dhruv Batra. SOAT: A Scene- and Object-Aware Transformer for Vision-and-Language Navigation. Neural Information Processing Systems (NeurIPS), 2021.
 58. Oleksandr Maksymets, Vincent Cartillier, Aaron Gokaslan, Erik Wijmans, Stefan Lee, Wojciech Galuba, Dhruv Batra. THDA: Treasure Hunt Data Augmentation for Semantic Navigation. International Conference on Computer Vision (ICCV), 2021.
 59. Yash Kant, Abhinav Moudgil, Dhruv Batra, Devi Parikh, Harsh Agrawal. Contrast and Classify: Training Robust VQA Models. International Conference on Computer Vision (ICCV), 2021.
 60. Joel Ye, Dhruv Batra, Abhishek Das, Erik Wijmans. Auxiliary Tasks and Exploration Enable ObjectNav. International Conference on Computer Vision (ICCV), 2021.
 61. Xiaoming Zhao, Harsh Agrawal, Dhruv Batra, Alexander Schwing. The Surprising Effectiveness of Visual Odometry Techniques for Embodied PointGoal Navigation. International Conference on Computer Vision (ICCV), 2021.
 62. Jacob Krantz, Aaron Gokaslan, Dhruv Batra, Stefan Lee, Oleksandr Maksymets. Waypoint Models for Instruction-guided Navigation in Continuous Environments. International Conference on Computer Vision (ICCV), 2021.
Oral (Selection Rate: 210/6236 = 3.3%)
 63. Brennan Shacklett, Erik Wijmans, Aleksei Petrenko, Manolis Savva, Dhruv Batra, Vladlen Koltun, Kayvon Fatahalian. Large Batch Simulation for Deep Reinforcement Learning. International Conference on Learning Representations (ICLR), 2021.
 64. Naoki Yokoyama, Sehoon Ha, Dhruv Batra. Success Weighted by Completion Time: A Dynamics-Aware Evaluation Criteria for Embodied Navigation. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2021.
 65. Joanne Truong, Denis Yarats, Tianyu Li, Franziska Meier, Sonia Chernova, Dhruv Batra, Akshara Rai. Learning Navigation Skills for Legged Robots with Learned Robot Embeddings. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2021.
 66. Joanne Truong, Sonia Chernova, Dhruv Batra. Bi-directional Domain Adaptation for Sim2Real Transfer of Embodied Navigation Agents. IEEE International Conference on Robotics and Automation (ICRA), 2021.
 67. Sameer Dharur, Purva Tendulkar, Dhruv Batra, Devi Parikh, Ramprasaath R. Selvaraju. SoT-ing VQA Models: Contrastive Gradient Learning for Improved Consistency. Conference of the North American Chapter of the Association for Computational Linguistics (NAACL), 2021.
 68. Vincent Cartillier, Zhile Ren, Neha Jain, Stefan Lee, Irfan Essa, Dhruv Batra. Semantic MapNet: Building Allocentric Semantic Maps and Representations from Egocentric Views. AAAI Conference on Artificial Intelligence (AAAI), 2021.
 69. Peter Anderson, Ayush Shrivastava, Joanne Truong, Arjun Majumdar, Devi Parikh, Dhruv Batra, Stefan Lee. Sim-to-Real Transfer for Vision-and-Language Navigation. Conference on Robot Learning (CoRL), 2020.
 70. Joel Ye, Dhruv Batra, Erik Wijmans, Abhishek Das. Auxiliary Tasks Speed Up Learning PointGoal Navigation. Conference on Robot Learning (CoRL), 2020.
 71. Samyak Datta, Oleksandr Maksymets, Judy Hoffman, Stefan Lee, Dhruv Batra, Devi Parikh. Integrating Egocentric Localization for More Realistic Point-Goal Navigation Agents. Conference on Robot Learning (CoRL), 2020.

72. Michael Cogswell, Jiasen Lu, Rishabh Jain, Stefan Lee, Devi Parikh, Dhruv Batra. Dialog without Dialog Data: Learning Visual Dialog Agents from VQA Data. Neural Information Processing Systems (NeurIPS), 2020.
73. Meera Hahn, Jacob Krantz, Dhruv Batra, Devi Parikh, James Rehg, Stefan Lee, Peter Anderson. Where Are You? Localization from Embodied Dialog. Conference on Empirical Methods in Natural Language Processing (EMNLP), 2020.
74. Abhishek Kadian*, Joanne Truong*, Aaron Gokaslan, Alexander Clegg, Erik Wijmans, Stefan Lee, Manolis Savva, Sonia Chernova, Dhruv Batra. Sim2Real Predictivity: Does Evaluation in Simulation Predict Real-World Performance. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2020.
75. Arjun Majumdar, Ayush Shrivastava, Stefan Lee, Peter Anderson, Devi Parikh, Dhruv Batra. Improving Vision-and-Language Navigation with Image-Text Pairs from the Web. European Conference on Computer Vision (ECCV), 2020.
Spotlight (Selection Rate: 265/5025 = 5.3%)
76. Jacob Krantz, Erik Wijmans, Arjun Majumdar, Dhruv Batra, Stefan Lee. Beyond the Nav-Graph: Vision-and-Language Navigation in Continuous Environments. European Conference on Computer Vision (ECCV), 2020.
77. Medhini Narasimhan, Erik Wijmans, Xinlei Chen, Trevor Darrell, Dhruv Batra, Devi Parikh, Amanpreet Singh. Seeing the Un-Scene: Learning Amodal Semantic Maps for Room Navigation. European Conference on Computer Vision (ECCV), 2020.
78. Vishvak Murahari, Dhruv Batra, Devi Parikh, Abhishek Das. Large-scale Pretraining for Visual Dialog: A Simple State-of-the-Art Baseline. European Conference on Computer Vision (ECCV), 2020.
79. Yash Kant, Dhruv Batra, Peter Anderson, Alex Schwing, Devi Parikh, Jiasen Lu, Harsh Agrawal. Spatially Aware Multimodal Transformers for TextVQA. European Conference on Computer Vision (ECCV), 2020.
80. Erik Wijmans, Abhishek Kadian, Ari Morcos, Stefan Lee, Irfan Essa, Devi Parikh, Manolis Savva, Dhruv Batra. DD-PPO: Learning Near-Perfect PointGoal Navigators from 2.5 Billion Frames. International Conference on Learning Representations (ICLR), 2020.
81. Nirbhay Modhe, Prithvijit Chattopadhyay, Mohit Sharma, Abhishek Das, Devi Parikh, Dhruv Batra, Ramakrishna Vedantam. IR-VIC: Unsupervised Discovery of Sub-goals for Transfer in RL. International Joint Conference on Artificial Intelligence (IJCAI), 2020.
82. Peter Anderson*, Ayush Shrivastava*, Devi Parikh, Dhruv Batra, Stefan Lee. Chasing Ghosts: Instruction Following as Bayesian State Tracking. Neural Information Processing Systems (NeurIPS), 2019.
83. ViLBERT: Pretraining Task-Agnostic Visiolinguistic Representations for Vision-and-Language Tasks. Jiasen Lu, Dhruv Batra, Devi Parikh, Stefan Lee. Neural Information Processing Systems (NeurIPS), 2019.
84. Manolis Savva*, Abhishek Kadian*, Oleksandr Maksymets*, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, Devi Parikh, Dhruv Batra. Habitat: A Platform for Embodied AI Research. International Conference on Computer Vision (ICCV), 2019.
Best Paper Award Nominee (Selection Rate: 11/4303 = 0.25%)
85. Harsh Agrawal*, Karan Desai*, Xinlei Chen, Rishabh Jain, Dhruv Batra, Devi Parikh, Stefan Lee, Peter Anderson. nocaps: novel object captioning at scale. International Conference on Computer Vision (ICCV), 2019.
86. Daniel Gordon, Abhishek Kadian, Devi Parikh, Judy Hoffman, Dhruv Batra. SplitNet: Sim2Sim and Task2Task Transfer for Embodied Visual Navigation. International Conference on Computer Vision (ICCV),

2019.

87. Jianwei Yang*, Zhile Ren*, Mingze Xu, Xinlei Chen, David Crandall, Devi Parikh, Dhruv Batra. Embodied Visual Recognition. International Conference on Computer Vision (ICCV), 2019.
88. Ramprasaath R. Selvaraju, Stefan Lee, Yilin Shen, Hongxia Jin, Dhruv Batra, Devi Parikh. Taking a HINT: Leveraging Explanations to Make Vision and Language Models More Grounded. International Conference on Computer Vision (ICCV), 2019.
89. Jyoti Aneja, Harsh Agrawal, Dhruv Batra, Alexander Schwing. Sequential Latent Spaces for Modeling the Intention During Diverse Image Captioning. International Conference on Computer Vision (ICCV), 2019.
90. Vishvak Murahari, Prithvijit Chattopadhyay, Dhruv Batra, Devi Parikh, Abhishek Das. Improving Generative Visual Dialog by Answering Diverse Questions. Conference on Empirical Methods in Natural Language Processing (EMNLP), 2019.
91. Ashwin Kalyan, Peter Anderson, Stefan Lee, Dhruv Batra. Differentiable Decoding of Sets of Sequences for Neural Sequence Models. International Conference on Machine Learning (ICML), 2019.
92. Ramakrishna Vedantam, Karan Desai, Stefan Lee, Marcus Rohrbach, Dhruv Batra, Devi Parikh. Probabilistic Neural-symbolic Models for Interpretable Visual Question Answering. International Conference on Machine Learning (ICML), 2019.
Oral (Long Talk)
93. Abhishek Das, Théophile Gervet, Joshua Romoff, Dhruv Batra, Devi Parikh, Michael Rabbat, Joelle Pineau. TarMAC: Targeted Multi-Agent Communication. International Conference on Machine Learning (ICML), 2019.
94. Yash Goyal, Ziyang Wu, Jan Ernst, Dhruv Batra, Devi Parikh, Stefan Lee. Counterfactual Visual Explanations. International Conference on Machine Learning (ICML), 2019.
95. Jin-Hwa Kim*, Nikita Kitaev*, Xinlei Chen, Marcus Rohrbach, Yuandong Tian, Dhruv Batra, Devi Parikh. CoDraw: Collaborative Drawing as a Testbed for Grounded Goal-driven Communication. Annual Meeting of the Association for Computational Linguistics (ACL), 2019.
96. Erik Wijmans, Samyak Datta, Oleksandr Maksymets, Abhishek Das, Georgia Gkioxari, Stefan Lee, Irfan Essa, Devi Parikh, Dhruv Batra. Embodied Question Answering in Photorealistic Environments with Point Cloud Perception. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2019.
Oral (Selection Rate: 288/5160 = 5.5%)
97. Licheng Yu, Xinlei Chen, Georgia Gkioxari, Mohit Bansal, Tamara L. Berg, Dhruv Batra. Multi-target Embodied Question Answering. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2019.
98. Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, Marcus Rohrbach. Towards VQA Models That Can Read. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2019.
99. Huda Alamri, Vincent Cartillier, Abhishek Das, Jue Wang, Stefan Lee, Peter Anderson, Irfan Essa, Devi Parikh, Dhruv Batra, Anoop Cherian, Tim K. Marks, Chiori Hori. Audio-Visual Scene-Aware Dialog. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2019.
100. Satwik Kottur, José M. F. Moura, Devi Parikh, Dhruv Batra, Marcus Rohrbach. CLEVR-Dialog: A Diagnostic Dataset for Multi-Round Reasoning in Visual Dialog. Conference of the North American Chapter of the Association for Computational Linguistics (NAACL), 2019.

-
101. Nan Rosemary Ke, Amanpreet Singh, Ahmed Touati, Anirudh Goyal, Yoshua Bengio, Devi Parikh, Dhruv Batra. Modeling the Long-Term Future in Model-Based Reinforcement Learning. *International Conference on Learning Representations (ICLR)*, 2019.
 102. Abhishek Das, Georgia Gkioxari, Stefan Lee, Devi Parikh, Dhruv Batra. Neural Modular Control for Embodied Question Answering. *Conference on Robot Learning (CoRL)*, 2018.
 103. Jianwei Yang*, Jiasen Lu*, Stefan Lee, Dhruv Batra, Devi Parikh. Learning to Ask Questions to Learn Visual Recognition. *Conference on Robot Learning (CoRL)*, 2018.
 104. 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.
 105. 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.
 106. Jianwei Yang*, Jiasen Lu*, Stefan Lee, Dhruv Batra, Devi Parikh. Graph R-CNN for Scene Graph Generation. *European Conference on Computer Vision (ECCV)*, 2018.
 107. 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)
 108. 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%)
 109. 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.
 110. Jiasen Lu, Jianwei Yang, Dhruv Batra, Devi Parikh. Neural Baby Talk. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
Spotlight (Selection Rate: 294/3309 = 8.9%)
 111. 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.
 112. 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.
 113. 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 (NeurIPS)*, 2017.
 114. 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%)
 115. 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.

-
116. 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%)
 117. 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%)
 118. 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.
 119. 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.
 120. 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.
 121. 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)
 122. 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.
 123. 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)
 124. 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.
 125. 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.
 126. 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 (NeurIPS)*, 2016.
 127. Jiasen Lu, Jianwei Yang, Dhruv Batra, Devi Parikh. Hierarchical Question-Image Co-Attention for Visual Question Answering. *Neural Information Processing Systems (NeurIPS)*, 2016.
 128. 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.
 129. Aishwarya Agrawal, Dhruv Batra, Devi Parikh. Analyzing the Behavior of Visual Question Answering Models. *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2016.

-
130. 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.
131. 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.
132. 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.
133. 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.
134. 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 (Selection Rate: 206/2123 = 9.7%)
135. 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 (Selection Rate: 206/2123 = 9.7%)
136. 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.
137. Jianwei Yang, Devi Parikh, Dhruv Batra. Simultaneously Discovering Image Clusters and Deep Representations. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
138. 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.
139. 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
140. 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.
141. Ratnesh Kumar, Dhruv Batra. Pose Tracking by Efficiently Exploiting Global Features. *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2016.
142. Dustin Bales, Pablo Tarazaga, Mary Kasarda, Dhruv Batra. Gender Classification using Vibration Measurements. *IMAC Conference on Human Induced Vibrations*, 2016. Society for Experimental Mechanics.
143. Qing Sun, Dhruv Batra. SubmodBoxes: Near-Optimal Search for a Set of Diverse Object Proposals. *Neural Information Processing Systems (NeurIPS)*, 2015.

-
144. Stanislav 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.
 145. Faruk Ahmed, Daniel Tarlow, Dhruv Batra. Optimizing expected Intersection-over-Union with Candidate-Constrained CRFs. *International Conference on Computer Vision (ICCV)*, 2015.
 146. 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 (Selection Rate: 71/2151 = 3.3%)
 147. Clint Solomon Mathialagan, Andrew Gallagher, Dhruv Batra. VIP: Finding Important People in Group Photographs. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
 148. Adarsh Prasad, Stefanie Jegelka, Dhruv Batra. Submodular meets Structured: Finding Diverse Subsets in Exponentially-Large Structured Item Sets. *Neural Information Processing Systems (NeurIPS)*, 2014.
Spotlight (Selection Rate: 82/1678 = 4.9%)
 149. 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.
 150. Kun Duan, David Crandall, Dhruv Batra. Multi-modal Learning in Loosely-Organized Web Images. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.
 151. 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 (Selection Rate 7.7%)
 152. Daozheng Chen, Dhruv Batra, William T. Freeman. Group Norm for Learning Structured SVMs with Unstructured Latent Variables. *International Conference on Computer Vision (ICCV)*, 2013.
 153. 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.
 154. Payman Yadollahpour, Dhruv Batra, Greg Shakhnarovich. Discriminative Re-ranking of Diverse Segmentations. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013.
 155. 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 (Selection Rate 4%)
 156. 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 (Selection Rate 11%)
 157. Abner Guzman-Rivera, Dhruv Batra, Pushmeet Kohli. Multiple Choice Learning: Learning to Produce Multiple Structured Outputs. *Neural Information Processing Systems (NeurIPS)*, 2012.
 158. 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 (Selection Rate 2.6%)

-
159. Dhruv Batra. An Efficient Message-Passing Algorithm for the M-Best MAP Problem. *The Conference on Uncertainty in Artificial Intelligence (UAI)*, 2012.
Oral (Selection Rate 7.8%)
160. Tanmay Verma, Dhruv Batra. MaxFlow Revisited: An Empirical Comparison of Maxflow Algorithms for Dense Vision Problems. *British Machine Vision Conference (BMVC)*, 2012.
161. Kun Duan, Dhruv Batra, David J. Crandall. A Multi-layer Composite Model for Human Pose Estimation. *British Machine Vision Conference (BMVC)*, 2012.
162. 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.
163. Daniel Tarlow, Dhruv Batra, Pushmeet Kohli, Vladimir Kolmogorov. Dynamic Tree Block Coordinate Descent. *International Conference on Machine Learning (ICML)*, 2011.
Oral
164. 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.
165. Andrew Gallagher, Dhruv Batra, Devi Parikh. Inference for Order Reduction in Markov Random Fields. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011.
166. 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.
167. 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 (Selection Rate 12.8%)
168. 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.
169. 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.
170. Dhruv Batra, Devi Parikh, Adarsh Kowdle, Tsuhan Chen. Seed Image Selection in Interactive Cosegmentation. *IEEE International Conference on Image Processing (ICIP)*, 2009.
171. Dhruv Batra, Rahul Sukthankar and Tsuhan Chen. Semi-Supervised Clustering via Learnt Codeword Distances. *British Machine Vision Conference (BMVC)*, 2008.
172. Dhruv Batra, Rahul Sukthankar and Tsuhan Chen. Learning Class-Specific Affinities for Image Labelling. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2008.
173. Dhruv Batra, Girish Singhal, and Santanu Chaudhury. Gabor Filter based Fingerprint Classification using Support Vector Machines. *IEEE India Annual Conference (INDICON)*, 2004.

Workshop Papers

174. 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)

175. 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)

176. 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.

177. 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.

178. 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.*

179. 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.*

180. 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.*

181. Michael Cogswell, Dhruv Batra. Semantic Segmentation with Deep Learning. *SUNw: Scene Understanding Workshop, IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2014.*

182. 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 (NeurIPS), 2013.*

183. 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 (NeurIPS), 2012.*

184. Varun Ramakrishna, Dhruv Batra. Mode-Marginals: Expressing Uncertainty via Diverse M-Best Solutions. *Workshop on Perturbations, Optimization, and Statistics, Neural Information Processing Systems (NeurIPS), 2012.*

185. 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 (NeurIPS), 2011.

186. 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 (NeurIPS)*, 2011.
187. Dhruv Batra, Greg Shakhnarovich. Similarity Sensitive Nonlinear Embeddings. *Workshop on Kernels and Distances for Computer Vision, International Conference on Computer Vision (ICCV)*, 2011.
188. 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.
189. 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 (NeurIPS)*, 2009.
190. 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.
191. 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*.
192. 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)

- *Visual Cortex and Embodied Question Answering*
 - ☞ NeurIPS workshop on Robot Learning, Dec 2023
 - ☞ CoRL workshop on Pre-Training for Robot Learning, Nov 2023
- *Robot Skill Acquisition: Where do things stand in simulation & sim2real land?*
 - ☞ CoRL workshop Towards Generalist Robots, Nov 2023
 - ☞ University of California San Diego CRI Seminar, Apr 2023
- *The Scaling Hypothesis and Embodied AI*
 - ☞ CVPR Workshop on Embodied AI, Jun 2022
- *The Simulation Argument for Embodied AI*
 - ☞ Berkeley AI Seminar, Oct 2021
 - ☞ Georgia Tech IRIM Fall Symposium, Aug 2021
- *Habitat 2.0: Training Home Assistants to Rearrange their Habitat*
 - ☞ ICCV Workshop on Human Interaction for Robotic Navigation, Oct 2021
 - ☞ CVPR Workshop on 3D Scene Understanding, Jun 2021
- *Emergence of Maps in the Memories of Blind Navigation Agents*
 - ☞ CVPR workshop on Fair, Data-efficient, and Trusted Computer Vision (TCV), Jun 2023
 - ☞ Computer Vision Meetup, Apr 2023
 - ☞ CVPR VizWiz workshop, Jun 2021
 - ☞ USC ICI Seminar, Apr 2021
 - ☞ AI2 Embodied AI Seminar, Apr 2021
- *From Disembodied to Embodied Multimodal Learning and AI*
 - ☞ Summer School at IIIT Hyderabad, Aug 2021
 - ☞ Keynote talk at the Conference of the North American Chapter of the Association for Computational Linguistics (NAACL), Jun 2021
 - ☞ CVPR Language for 3D workshop, Jun 2021
 - ☞ Alexa AI, Feb 2021
 - ☞ West Coast NLP (WeNLP) Summit, Oct 2020.
 - ☞ ECCV Commands For Autonomous Vehicles Workshop, Aug 2020.
 - ☞ ECCV Workshop on Assistive Computer Vision and Robotics, Aug 2020.
 - ☞ CVPR Workshop on Multimodal Learning, June 2020.
- *How far can we scale model-free RL for embodied AI?*
 - ☞ ICRA Workshop on Perception, Action, and Learning, June 2020.
- *Habitat: A Platform for Embodied AI Research*
 - ☞ IIT Gandhinagar, Dec 2019.
 - ☞ Georgia Tech IRIM Fall Symposium, Aug 2019.
 - ☞ NAACL Combined Workshop on Spatial Language Understanding (SpLU) & Grounded Communication for Robotics (RoboNLP), Jun 2019.
 - ☞ CVPR Workshop on Conceptual Captions, Jun 2019.
 - ☞ CVPR Workshop on Semantic Information, Jun 2019.
- *A-STAR: Agents that See, Talk, Act, and Reason*
 - ☞ Samsung AI Summit, Jan 2019.
 - ☞ Keynote talk at IEEE Winter Conference on Applications of Computer Vision (WACV), Jan 2019.

-
- Ⓒ ECCV Workshop on Visual Learning and Embodied Agents in Simulation Environments, Sep 2018.
 - Ⓒ RSS Workshop on Models and Representations for Natural Human-Robot Communication, Jun 2018.
 - Ⓒ 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*
 - Ⓒ NeurIPS 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 NeurIPS, 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 & Data Conference, Jul 2015.
- Ⓢ Deep Learning Summit, May 2015.
- *Submodular meets Structured: Finding Diverse Subsets in Exponentially-Large Structured Item Sets*
 - Ⓢ NeurIPS 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

- *Facebook open-sources AI Habitat to help robots navigate realistic environments*
 - [MIT Technology Review](#), [Venture Beat](#)
- *Indian-Americans shine in White House honors for rising stars in the world science and tech*
 - [Hindustan Times](#), [Georgia Tech News](#)
- *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 Research, 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

FUNDING & GIFTS

Research Grants & Gifts

- Early Career Awards for Scientist and Engineers (ECASE-Army), Army Research Office 2019 – 2024
“Towards Transparent Machine Perception Systems”
- Samsung Global Research Outreach (GRO) Grant 2018
“Leveraging Explanations to Improve VQA Models Through Focused Feedback”
Co-PI: Prof. Devi Parikh, Dr. Stefan Lee
- Siemens Research Grant 2018
“Learning Visual Curiosity”
Co-PI: Prof. Devi Parikh
- Gift funding from Mitsubishi Electric Research Lab (MERL) 2018
- AWS Machine Learning Research Award 2018
“Towards AI Agents that can See, Talk, and Act”
- DARPA Explainable AI (XAI) 2017 – 2021
“EQUAS: Explainable Question Answering System”
Lead Academic PI: Dhruv Batra. Co-PIs: Profs. Ray Mooney, Antonio Torralba, Devi Parikh. Prime: Raytheon.
- Office of Naval Research (ONR) Young Investigator Program (YIP) award 2017 – 2020
“Explainable and Trustworthy Intelligent Systems”
- Amazon Academic Research Award 2016
“Visual Dialog”
- Army Research Lab (ARL) Grant W911NF-15-2-0080 2015 – 2018
“Answering Binary Questions about Images”
Co-PI: Prof. Devi Parikh

-
- Defense Advanced Research Projects Agency (DARPA) Grant HR0011-16-1-0002 2015 – 2016
“Deep Cognitive Radio System Study”
Co-PI: Tim O’Shea (Hume Center, Virginia Tech)
 - Google Faculty Research Award 2015
“Visual Question Answering (VQA)”
Co-PI: Prof. Devi Parikh
 - Office of Naval Research (ONR) Grant N00014-14-1-0679 2014 – 2017
“Feedback-Enabled Joint Reasoning over Uncertain Sub-components of Perception”
 - Army Research Office (ARO) Young Investigator Program (YIP) Award W911NF-14-1-0180 2014 – 2017
“Building Reflective, Transparent, and Integrated Intelligent Systems”
 - NSF CAREER Award IIS-1350553 2014 – 2019
“Holistic Scene Understanding with Multiple Hypotheses from Vision Modules”
 - Virginia Tech Institute for Critical Technology and Applied Science (ICTAS) 2014 – 2016
Junior Faculty Award
“CloudCV: Large-Scale Distributed Computer Vision As A Cloud Service”
 - L-3 Communications (sub-award from VT Hume Center) 2014 – 2015
“Detection of Naval Objects in Images”
 - NSF Award IIS-1353694 2013 – 2015
“EAGER: Diverse M-Best Predictions from Probabilistic Models”

Equipment, Educational, Organizational Grants & Gifts

- Google Cloud Computing Support for CS 7643 Deep Learning 2017
- NRT-DESE: UrbComp: Data Science for Modeling, 2015 – 2020
“Understanding, and Advancing Urban Populations”
PI: Naren Ramakrishnan (CS, VT)
- State Council of Higher Education for Virginia (SCHEV) Equipment Trust Grant 2014
- Oak Ridge Associated University (ORAU) Event Organization Grant 2014
1st Mid Atlantic Computer Vision workshop at Virginia Tech; Co-PI: Devi Parikh.
- 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

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 (NeurIPS), 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 (NeurIPS), 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 (NeurIPS), 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 (NeurIPS), 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.
- NeurIPS 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.
- NeurIPS 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)