

Saurabh Gupta

CONTACT INFORMATION

Office address: 319 Coordinated Science Laboratory,
1308 W Main Street, Urbana, IL 61801
Email address: saurabhg@illinois.edu
Webpage: <https://saurabhg.web.illinois.edu/>
Google scholar: <https://scholar.google.com/citations?user=1HO5UacAAAAJ>

AFFILIATION

- **University of Illinois at Urbana-Champaign**
Assistant Professor, Electrical and Computer Engineering (100%) 2019 –
Assistant Professor (Affiliate), Computer Science (0%) 2019 –
Assistant Professor (Affiliate), Coordinated Science Laboratory (0%) 2019 –

EDUCATION

- **University of California, Berkeley** 2011 – 2018
Ph.D. in Computer Science
Advised by Prof. Jitendra Malik
- **Indian Institute of Technology, Delhi** 2007 – 2011
B. Tech in Computer Science and Engineering
President's Gold Medalist

ACADEMIC POSITIONS

- **Facebook AI Research** 2018 – 2019
Research Scientist
- **Google Research** May 2016 – April 2017
Intern (Mentors: Rahul Sukthankar, Sergey Levine, Jitendra Malik)
- **Microsoft Research** Summer 2014
Summer Research Intern (Mentors: Piotr Dollar, Larry Zitnick, John Platt)
- **Microsoft Research** Summer 2010
Summer Research Intern (Mentors: Manik Varma, Prateek Jain)
- **Indian Institute of Technology, Delhi** Summer 2009
Summer Undergraduate Research Award (Mentor: Subhashis Banerjee)

RESEARCH INTERESTS

- Computer vision, robotics, and machine learning.

SELECTED AWARDS AND HONORS

- **Google US/Canada Fellowship in Computer Vision 2015.**
- **Berkeley Graduate Student Fellowship 2011.** Awarded to the top 4% graduate student admits every year.

- **President's Gold Medal** for securing the highest CGPA across all students who graduated from Indian Institute of Technology Delhi in 2011.
- **1st place** in AI Habitat Challenge RGB-D track, CVPR 2019.
- **1st place** in AI Habitat Challenge RGB track, CVPR 2019.
- **1st place** in MS COCO Captioning Challenge, CVPR 2015.
- **1st place** in RGB-D Indoor Scene Semantic Segmentation Challenge at Reconstruction Meets Recognition Challenge, ICCV 2013.
- **Outstanding Reviewer:** ICCV 2015.
- **Student Undergraduate Research Award (SURA) 2009** and **Kalpana Chawla Award** for the research project on Activity Analysis at IIT Delhi.

STUDENTS

- Matthew Chang (Expected CS PhD 2024)
- Rishabh Goyal (Expected CS MS 2021)
- Arjun Gupta (Expected ECE BS 2021)

INVITED TALKS

- **Learning to Act by Watching Videos** (2019)
Multi-Modal Learning from Videos Workshop at CVPR 2019.
- **Perception and Visual Navigation in 3D Scenes** (2018)
University of Illinois at Urbana-Champaign, Carnegie Mellon University, Cornell University, University of Washington, ETH Zürich, University of Pennsylvania, University of Texas Austin, Princeton University, Georgia Institute of Technology, Brown University, New York University, University of Wisconsin at Madison, Toyota Institute of Technology, Simon Fraser University, University of British Columbia, Berkeley AI Research Seminar, Facebook AI Research, NVidia, SILO Seminar (University of Wisconsin at Madison).
- **Representations for Visually Guided Actions** (2017)
GRASP Seminar at University of Pennsylvania.
- **Cognitive Mapping and Planning** (2017)
Redwood Center for Theoretical Neuroscience, UC Berkeley, Indian Institute of Technology (IIT) Delhi, Berkeley AI Research Seminar, Google Brain, OpenAI, Lyft
- **Detailed Scene Understanding from RGB-D Images** (2016)
CMU VASC seminar, Toyota Technological Institute at Chicago (TTIC) Young Researcher Seminar Series, Allen Institute for Artificial Intelligence (AI2), Workshop on Understanding 3D and Visuo-Motor Learning at 3D Vision (3DV), Conference, Indian Institute of Technology (IIT) Delhi.
- **Image Captioning** (2015)
Large-scale Scene Understanding (LSUN) workshop at CVPR 2015.
- **RGB-D Scene Understanding** (2014)
Reconstruction Meets Recognition Challenge workshop at ICCV 2013 and ECCV 2014, Microsoft Research, Redmond.
- **Perceptual Organization & Recognition of Indoor Scenes from RGB-D Images** (2013)
Conference Talk at CVPR 2013, Visual Computing Lab, UC Berkeley.

TEACHING EXPERIENCE

- **University of Illinois Urbana-Champaign:**
 - ECE 598SG: Special Topics in Learning-based Robotics Fall 2020
 - CS 543 / ECE 549: Computer Vision Spring 2020
 - ECE 598SG: Special Topics in Learning-based Robotics Fall 2019
- **Guest Lectures:**
 - Computer Vision Class at IIT Delhi on Robot Navigation Fall 2017
 - Computer Vision Class at UC Berkeley on Robot Navigation Spring 2017
- **Teaching Assistant:**
 - **CS 280: Computer Vision**, Prof. Jitendra Malik, Prof. Alexei Efros Fall 2013
 - **CS 188: Artificial Intelligence**, Prof. Dan Klein, Prof. Pieter Abbeel Fall 2012

SERVICE

- Workshop on *Benchmarking in Robotics* in August 2019.
- Tutorial at CVPR 2019 on *Bringing Robots to the Vision Community*.
- Workshop at CVPR 2019 on *Computer Vision in 5 Years*.
- Tutorial at ICVGIP 2018 on *Learning-based techniques for visually-guided robotic manipulation and navigation*.
- Workshop at ECCV 2018 on *Visual Learning and Embodied Agents in Simulation Environments*.
- Area Chair for CVPR 2019, CVPR 2021.
- Reviewer for PAMI, IJCV, CVPR, ECCV, ICCV, ICML, NIPS, SIGGRAPH, ICLR, ICRA, CoRL, BMVC, IROS, IJCAI, ACCV, 3DV, ICVGIP, IKDD CODS, SIGGRAPH Asia.
- Student Reviewer for Graduate Admissions Applications at UC Berkeley 2012-2014.

CONFERENCE PUBLICATIONS

1. Devendra Singh Chaplot, Helen Jiang, **Saurabh Gupta**, and Abhinav Gupta. Semantic curiosity for active visual learning. In *European Conference on Computer Vision (ECCV)*, 2020
2. Senthil Purushwalkam, Tian Ye, **Saurabh Gupta**, and Abhinav Gupta. Aligning videos in space and time. In *European Conference on Computer Vision (ECCV)*, 2020
3. Devendra Singh Chaplot, Ruslan Salakhutdinov, Abhinav Gupta, and **Saurabh Gupta**. Neural topological SLAM for visual navigation. In *Computer Vision and Pattern Recognition (CVPR)*, 2020
4. Kiana Ehsani, Shubham Tulsiani, **Saurabh Gupta**, Ali Farhadi, and Abhinav Gupta. Use the force, luke! learning to predict physical forces by simulating effects. In *Computer Vision and Pattern Recognition (CVPR)*, 2020
5. Junfeng Guan, Sohrab Madani, Suraj Jog, **Saurabh Gupta**, and Haitham Hassanieh. Through fog high resolution imaging using millimeter wave radar. In *Computer Vision and Pattern Recognition (CVPR)*, 2020

6. Rohan Chitnis, Shubham Tulsiani, **Saurabh Gupta**, and Abhinav Gupta. Efficient bimanual manipulation using learned task schemas. In *International Conference on Robotics and Automation (ICRA)*, 2020
7. Rohan Chitnis, Shubham Tulsiani, **Saurabh Gupta**, and Abhinav Gupta. Intrinsic motivation for encouraging synergistic behavior. In *International Conference on Learning Representations (ICLR)*, 2020
8. Devendra Singh Chaplot, **Saurabh Gupta**, Dhiraj Gandhi, Abhinav Gupta, and Ruslan Salakhutdinov. Learning to explore using active neural mapping. In *International Conference on Learning Representations (ICLR)*, 2020
9. William Qi, Ravi Teja Mullapudi, **Saurabh Gupta**, and Deva Ramanan. Learning to move with affordance maps. In *International Conference on Learning Representations (ICLR)*, 2020
10. Ashish Kumar, **Saurabh Gupta**, and Jitendra Malik. Learning navigation subroutines by watching videos. In *Conference on Robot Learning (CoRL)*, 2019
11. Somil Bansal, Varun Tolani, **Saurabh Gupta**, Jitendra Malik, and Claire Tomlin. Combining optimal control and learning for visual navigation in novel environments. In *Conference on Robot Learning (CoRL)*, 2019
12. Tao Chen, **Saurabh Gupta**, and Abhinav Gupta. Learning exploration policies for navigation. In *International Conference on Learning Representations (ICLR)*, 2019
13. Michael Danielczuk, Matthew Matl, **Saurabh Gupta**, Andrew Li, Andrew Lee, Jeffrey Mahler, and Ken Goldberg. Segmenting unknown 3D objects from real depth images using mask R-CNN trained on synthetic point clouds. In *International Conference on Robotics and Automation (ICRA)*, 2019
14. Ashish Kumar*, **Saurabh Gupta***, David Fouhey, Sergey Levine, and Jitendra Malik. Visual memory for robust path following. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018. *denotes equal contribution.
15. Shubham Tulsiani, **Saurabh Gupta**, David Fouhey, Alexei A Efros, and Jitendra Malik. Factoring shape, pose, and layout from the 2D image of a 3D scene. In *Computer Vision and Pattern Recognition (CVPR)*, 2018
16. **Saurabh Gupta**, James Davidson, Sergey Levine, Rahul Sukthankar, and Jitendra Malik. Cognitive mapping and planning for visual navigation. In *Computer Vision and Pattern Recognition (CVPR)*, 2017
17. **Saurabh Gupta**, Judy Hoffman, and Jitendra Malik. Cross modal distillation for supervision transfer. In *Computer Vision and Pattern Recognition (CVPR)*, 2016
18. Judy Hoffman, **Saurabh Gupta**, and Trevor Darrell. Learning with side information through modality hallucination. In *Computer Vision and Pattern Recognition (CVPR)*, 2016
19. Judy Hoffman, **Saurabh Gupta**, Jian Leong, Sergio Guadarrama, and Trevor Darrell. Cross-modal adaptation for RGB-D detection. In *International Conference on Robotics and Automation (ICRA)*, 2016
20. **Saurabh Gupta**, Pablo Arbeláez, Ross Girshick, and Jitendra Malik. Aligning 3d models to RGB-D images of cluttered scenes. In *Computer Vision and Pattern Recognition (CVPR)*, 2015
21. Hao Fang*, **Saurabh Gupta***, Forrest Iandola*, Rupesh K Srivastava*, Li Deng, Piotr Dollár, Jianfeng Gao, Xiaodong He, Margaret Mitchell, John C Platt, C Lawrence Zitnick, and Geoffrey Zweig. From captions to visual concepts and back. In *Computer Vision and Pattern Recognition (CVPR)*, 2015. *denotes equal contribution.

22. Jacob Devlin, Hao Cheng, Hao Fang, **Saurabh Gupta**, Li Deng, Xiaodong He, Geoffrey Zweig, and Margaret Mitchell. Language models for image captioning: The quirks and what works. In *Association for Computational Linguistics (ACL)*, 2015
23. **Saurabh Gupta**, Ross Girshick, Pablo Arbeláez, and Jitendra Malik. Learning rich features from RGB-D images for object detection and segmentation. In *European Conference on Computer Vision (ECCV)*, 2014
24. **Saurabh Gupta**, Pablo Arbelaez, and Jitendra Malik. Perceptual organization and recognition of indoor scenes from RGB-D images. In *Computer Vision and Pattern Recognition (CVPR)*, 2013
25. Rahul Sharma, **Saurabh Gupta**, Bharath Hariharan, Alex Aiken, Percy Liang, and Aditya V Nori. A data driven approach for algebraic loop invariants. In *European Symposium on Programming (ESOP)*, 2013
26. Rahul Sharma, **Saurabh Gupta**, Bharath Hariharan, Alex Aiken, and Aditya V Nori. Verification as learning geometric concepts. In *International Static Analysis Symposium*, 2013
27. Pablo Arbeláez, Bharath Hariharan, Chunhui Gu, **Saurabh Gupta**, Lubomir Bourdev, and Jitendra Malik. Semantic segmentation using regions and parts. In *Computer Vision and Pattern Recognition (CVPR)*, 2012

JOURNAL PUBLICATIONS

1. **Saurabh Gupta**, Varun Tolani, James Davidson, Sergey Levine, Rahul Sukthankar, and Jitendra Malik. Cognitive mapping and planning for visual navigation. *International Journal of Computer Vision (IJCV)*, 2019
2. Jitendra Malik, Pablo Arbeláez, Joao Carreira, Katerina Fragkiadaki, Ross Girshick, Georgia Gkioxari, **Saurabh Gupta**, Bharath Hariharan, Abhishek Kar, and Shubham Tulsiani. The three R's of computer vision: Recognition, reconstruction and reorganization. *Pattern Recognition Letters*, 72, 2016
3. **Saurabh Gupta**, Pablo Arbeláez, Ross Girshick, and Jitendra Malik. Indoor scene understanding with RGB-D images: Bottom-up segmentation, object detection and semantic segmentation. *International Journal of Computer Vision (IJCV)*, 2015

TECHNICAL REPORTS

1. Matthew Chang, Arjun Gupta, and **Saurabh Gupta**. Semantic visual navigation by watching youtube videos. *arXiv preprint arXiv:2006.10034*, 2020
2. Adithyavairavan Murali*, Tao Chen*, Kalyan Vasudev Alwala*, Dhiraj Gandhi*, Lerrel Pinto, **Saurabh Gupta**, and Abhinav Gupta. Pyrobot: An open-source robotics framework for research and benchmarking. *arXiv preprint arXiv:1906.08236*, 2019
3. Peter Anderson, Angel Chang, Devendra Singh Chaplot, Alexey Dosovitskiy, **Saurabh Gupta**, Vladlen Koltun, Jana Kosecka, Jitendra Malik, Roozbeh Mottaghi, Manolis Savva, and Amir Zamir. On evaluation of embodied navigation agents. *arXiv preprint arXiv:1807.06757*, 2018
4. **Saurabh Gupta**, David Fouhey, Sergey Levine, and Jitendra Malik. Unifying map and landmark based representations for visual navigation. *arXiv preprint arXiv:1712.08125*, 2017
5. Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, **Saurabh Gupta**, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015

6. **Saurabh Gupta** and Jitendra Malik. Visual semantic role labeling. *arXiv preprint arXiv:1505.04474*, 2015
7. **Saurabh Gupta**, Bharath Hariharan, and Jitendra Malik. Exploring person context and local scene context for object detection. *arXiv preprint arXiv:1511.08177*, 2015
8. Jacob Devlin, **Saurabh Gupta**, Ross Girshick, Margaret Mitchell, and C Lawrence Zitnick. Exploring nearest neighbor approaches for image captioning. *arXiv preprint arXiv:1505.04467*, 2015