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

- **Indian Institute of Technology, Delhi**   2007 – 2011
  B. Tech in Computer Science and Engineering
  President's Gold Medalist

- **University of California, Berkeley**   2011 – 2018
  Ph.D. in Computer Science
  Advisor: Prof. Jitendra Malik
  Thesis title: Representations for Visually Guided Actions

ACADEMIC POSITIONS

- **Indian Institute of Technology, Delhi**   Summer 2009
  Summer Undergraduate Research Award (Mentor: Subhashis Banerjee)

- **Microsoft Research**   Summer 2010
  Summer Research Intern (Mentors: Manik Varma, Prateek Jain)

- **Microsoft Research**   Summer 2014
  Summer Research Intern (Mentors: Piotr Dollar, Larry Zitnick, John Platt)

- **Google Research**   May 2016 – April 2017
  Intern (Mentors: Rahul Sukthankar, Sergey Levine, Jitendra Malik)

- **Facebook AI Research**   2018 – 2019
  Research Scientist (Research Manager: Abhinav Gupta)

RESEARCH INTERESTS

- Computer vision, robotics, and machine learning.

SELECTED AWARDS AND HONORS

- **Student Undergraduate Research Award (SURA) 2009** and **Kalpana Chawla Award** for the research project on Activity Analysis at IIT Delhi.

- **President's Gold Medal** for securing the highest CGPA across all students who graduated from Indian Institute of Technology Delhi in 2011.

- **Berkeley Graduate Student Fellowship 2011.** Awarded to the top 4% graduate student admits every year.
- **1st place** in RGB-D Indoor Scene Semantic Segmentation Challenge at Reconstruction Meets Recognition Challenge, ICCV 2013.
- **Outstanding Reviewer:** ICCV 2015.
- **Google US/Canada Fellowship in Computer Vision 2015.**
- **1st place** in MS COCO Captioning Challenge, CVPR 2015.
- **1st place** in AI Habitat Challenge RGB track, CVPR 2019.
- **1st place** in AI Habitat Challenge RGB-D track, CVPR 2019.
- **Amazon Research Award 2020.**
- **NSF CAREER Award 2022.**
- **Andrew T. Yang Research and Entrepreneurship Award 2022.**

STUDENTS

Ph.D.
- Matthew Chang (CS, Fall 2019 –), passed prelim exam in CS in Spring 2023.
- Arjun Gupta (ECE, Spring 2021 –), passed qualifying exam in ECE in Fall 2022.
- Aditya Prakash (CS, Fall 2021 –), passed qualifying exam in CS in Spring 2023.
- Shaowei Liu (CS, Fall 2021 –), co-advised with Prof. Shenlong Wang, passed qualifying exam in CS in Fall 2022.
- Xiaoyu Zhang (CS, Fall 2022 –)

Master's
- Rishabh Goyal (CS, 2019 – 2021), now Computer Vision Engineer at
- Sahil Modi (CS BS+MS, 2021 – 2022), next Software Engineer at NVidia.
- Mohit Goyal (ECE, 2020 – 2022), next Software Engineer at Google.
- Michelle Zhang (ECE, 2023 –)

Undergraduates
- Peiyi Chen (ECE BS 2022), next Meta.
- Rohini Amy Sharma (CS BS 2022).
- Sanjeev Venkatesan (CS BS 2021), next MS student at USC.
- Arjun Gupta (ECE BS 2020), next PhD student at UIUC.
- Sahil Modi (CS BS 2021), next BS+MS student at UIUC.
- Max E.Shepherd (ECE BS 2023).
- Matthew Jin (ECE BS 2023).
- Ruisen Tu (ECE BS 2024).

- **Perceptual Organization & Recognition of Indoor Scenes from RGB-D Images** (2013)
  Conference Talk at CVPR 2013.
  Visual Computing Lab, UC Berkeley.

- **RGB-D Scene Understanding** (2014)
  Reconstruction Meets Recognition Challenge workshop at ECCV 2014.
  Reconstruction Meets Recognition Challenge workshop at ICCV 2013.
  Microsoft Research, Redmond.

- **Image Captioning** (2015)
  Large-scale Scene Understanding (LSUN) workshop at CVPR 2015.

- **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.

- **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.

- **Representations for Visually Guided Actions** (2017)
  GRASP Seminar at University of Pennsylvania.

- **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).

- **Learning to Act by Watching Videos** (2019)
  Multi-Modal Learning from Videos Workshop at CVPR 2019.

- **Learning to Speculate about People and Places for Visual Navigation** (2021)
  Perception and Control for Autonomous Navigation in Crowded, Dynamic Environments Workshop at RSS 2021.

- **Robot Learning by Understanding Egocentric Videos** (2021)
  Computational Sensorimotor Learning (CSL) Seminar at CSAIL@MIT.

- **Learning to Move and Moving to Learn** (2021)
  3D Scene Understanding for Vision, Graphics, and Robotics Workshop at CVPR 2021.
  Visual Learning and Reasoning for Robotics Workshop at RSS 2021.

- **Representations for Visual Navigation and How to Train Them** (2021)
  Virtual Seminar Series on Geometry Processing and 3D Computer Vision (3DGV).

- **Use of Machine Learning in Robotics** (2022)
  Faculty Development Program, Swami Keshvanand Institute of Technology, Management & Gramothan, Jaipur, India.

- **Imbuing Robots with Common Sense** (2022)
  GFX Research Scholars seminar, UIUC

- **Robot Learning by Understanding Egocentric Videos** (2023)
  VASC Seminar, CMU, 2023

- **Learning to Articulate Doors and Manipulate Vines** (2023)
  Illinois Scholars Undergraduate Research (ISUR) Seminar, UIUC 2023

- **Robot Learning by Understanding Egocentric Videos** (2023)
  Embodied AI Workshop, CVPR 2023

- **Understanding and Articulating Articulated Objects** (2023)
  Workshop on 3D Vision and Robotics, CVPR 2023

- **Robot Learning by Understanding Egocentric Videos** (2023)
  Learning by Asking for Intelligent Robots and Agents Workshop at RO-MAN 2023.

TEACHING EXPERIENCE

- **University of Illinois Urbana-Champaign:**

  - ECE 598SG: Special Topics in Learning-based Robotics   Fall 2019, Enrollment: 38
  - CS 543 / ECE 549: Computer Vision   Spring 2020, Enrollment: 179
  - ECE 598SG: Special Topics in Learning-based Robotics   Fall 2020, Enrollment: 29
  - CS 543 / ECE 549: Computer Vision   Spring 2021, Enrollment: 149
  - ECE 598SG: Special Topics in Learning-based Robotics   Fall 2021, Enrollment: 56
  - ECE 598SG: Special Topics in Learning-based Robotics   Fall 2022, Enrollment: 30
  - CS 543 / ECE 549: Computer Vision   Spring 2023, Enrollment: 125
  - CS 444: Deep Learning for Computer Vision   Fall 2023, Enrollment: 100

- **Guest Lectures:**

  - Computer Vision Class at UC Berkeley on "Robot Navigation"   Spring 2017
  - Computer Vision Class at IIT Delhi on "Robot Navigation"   Fall 2017

– Machine Learning for Signal Processing at at UIUC on "Computer Vision"   Fall 2019

- **Teaching Assistant:**

    – **CS 188: Artificial Intelligence**, Prof. Dan Klein, Prof. Pieter Abbeel   Fall 2012
    – **CS 280: Computer Vision**, Prof. Jitendra Malik, Prof. Alexei Efros   Fall 2013

EXTERNAL SERVICE

- Reviewer for PAMI, IJCV, CVPR, ECCV, ICCV, ICML, NIPS, SIGGRAPH, ICLR, ICRA, CoRL, BMVC, IROS, IJCAI, ACCV, 3DV, ICVGIP, IKDD CODS, SIGGRAPH Asia.

- Organized workshop at ECCV 2018 on *Visual Learning and Embodied Agents in Simulation Environments*.

- Organized tutorial at ICVGIP 2018 on *Learning-based techniques for visually-guided robotic manipulation and navigation*.

- Organized workshop at CVPR 2019 on *Computer Vision in 5 Years*.

- Organized tutorial at CVPR 2019 on *Bringing Robots to the Vision Community*.

- Organized workshop on *Benchmarking in Robotics* in August 2019.

- Area Chair for CVPR 2019, CVPR 2021, NeurIPS 2022, CVPR 2023, ICCV 2023, CVPR 2024.

- NSF Panelist for NRI 2.0, 2020.

- Associate Editor for ICRA 2021.

- Reviewer for RSS Pioneers 2023.

DIVERSITY, EQUITY, INCLUSION, OR ACCESS ACTIVITIES

- Presented at the Faculty Development Program at Swami Keshvanand Institute of Technology, Management and Gramothan, Jaipur, India in 2022.

- Hosted students from the Engineering Pathways' GearUP program for rising juniors who are transitioning from Illinois community colleges to the University of Illinois at Urbana-Champaign in 2022 and 2023.

- Mentoring a high-school student on a machine learning and computer vision project (2022 – 2023).

- Introduced Computer Vision to high-school students at Illinois Summer Academy Summer Camp at ECE@UIUC, 2023.

- Spoke about Grad School 101 to undergraduate students at Lunch and Learn at CS@UIUC, 2023.

- Participated in Speed Mentoring Event for Graduate Students at CVPR 2023.

- Introduced undergraduates to research at UIUC by giving talks at the GFX Research Scholars seminar (2022) and the Illinois Scholars Undergraduate Research Seminar (2023).

Conference Publications

1. 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. (Acceptance rate: 24.0%.)

2. Rahul Sharma, **Saurabh Gupta**, Bharath Hariharan, Alex Aiken, and Aditya V Nori. Verification as learning geometric concepts. In *International Static Analysis Symposium*, 2013. (Acceptance rate: 41.1%.)

3. 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. (Acceptance rate: 25.8%.)

4. **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. (Acceptance rate: 25.2%.)

5. **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. (Acceptance rate: 27.9%.)

6. 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. (Acceptance rate: 25.0%.)

7. 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. (Acceptance rate: 28.3%.). *denotes equal contribution.

8. **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. (Acceptance rate: 28.3%.)

9. 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. (Acceptance rate: 35.0%.)

10. Judy Hoffman, **Saurabh Gupta**, and Trevor Darrell. Learning with side information through modality hallucination. In *Computer Vision and Pattern Recognition (CVPR)*, 2016. (Acceptance rate: 29.9%.)

11. **Saurabh Gupta**, Judy Hoffman, and Jitendra Malik. Cross modal distillation for supervision transfer. In *Computer Vision and Pattern Recognition (CVPR)*, 2016. (Acceptance rate: 29.9%.)

12. **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. (Acceptance rate: 29.9%.)

13. 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. (Acceptance rate: 29.6%.)

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. (Acceptance rate: 20.8%.). *denotes equal contribution.

15. 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. (Acceptance rate: 45.0%.)

16. Tao Chen, **Saurabh Gupta**, and Abhinav Gupta. Learning exploration policies for navigation. In *International Conference on Learning Representations (ICLR)*, 2019. (Acceptance rate: 31.4%.)

17. 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. (Acceptance rate: 27.6%.)

18. Ashish Kumar, **Saurabh Gupta**, and Jitendra Malik. Learning navigation subroutines by watching videos. In *Conference on Robot Learning (CoRL)*, 2019. (Acceptance rate: 27.6%.)

19. William Qi, Ravi Teja Mullapudi, **Saurabh Gupta**, and Deva Ramanan. Learning to move with affordance maps. In *International Conference on Learning Representations (ICLR)*, 2020. (Acceptance rate: 26.5%.)

20. Devendra Singh Chaplot, Dhiraj Gandhi, **Saurabh Gupta***, Abhinav Gupta*, and Ruslan Salakhutdinov*. Learning to explore using active neural mapping. In *International Conference on Learning Representations (ICLR)*, 2020. (Acceptance rate: 26.5%.). *denotes equal contribution.

21. Rohan Chitnis, Shubham Tulsiani, **Saurabh Gupta**, and Abhinav Gupta. Intrinsic motivation for encouraging synergistic behavior. In *International Conference on Learning Representations (ICLR)*, 2020. (Acceptance rate: 26.5%.)

22. 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. (Acceptance rate: 44.0%.)

23. 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. (Acceptance rate: 22.1%.)

24. 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. (Acceptance rate: 22.1%.)

25. Devendra Singh Chaplot, Ruslan Salakhutdinov*, Abhinav Gupta*, and **Saurabh Gupta***. Neural topological SLAM for visual navigation. In *Computer Vision and Pattern Recognition (CVPR)*, 2020. (Acceptance rate: 22.1%.). *denotes equal contribution.

26. Senthil Purushwalkam, Tian Ye, **Saurabh Gupta**, and Abhinav Gupta. Aligning videos in space and time. In *European Conference on Computer Vision (ECCV)*, 2020. (Acceptance rate: 27.1%.)

27. Devendra Singh Chaplot, Helen Jiang, **Saurabh Gupta**, and Abhinav Gupta. Semantic curiosity for active visual learning. In *European Conference on Computer Vision (ECCV)*, 2020. (Acceptance rate: 27.1%.)

28. Matthew Chang, Arjun Gupta, and **Saurabh Gupta**. Semantic visual navigation by watching youtube videos. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. (Acceptance rate: 20.1%.)

29. Silvery Fu, **Saurabh Gupta**, Radhika Mittal, and Sylvia Ratnasamy. On the use of ML for blackbox system performance prediction. In *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2021. (Acceptance rate: 16%.)

30. Arun Narenthiran Sivakumar, Sahil Modi, Mateus Valverde Gasparino, Che Ellis, Andres Baquero Velasquez, Girish Chowdhary*, and **Saurabh Gupta***. Learned visual navigation for under-canopy agriculture robots. In *Robotics: Science and Systems (RSS)*, 2021. (Acceptance rate: 27%.). *denotes equal contribution.

31. Devendra Chaplot, Murtaza Dalal, **Saurabh Gupta**, Jitendra Malik, and Ruslan Salakhutdinov. SEAL: Self-supervised embodied active learning using exploration and 3D consistency. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. (Acceptance rate: 26%.)

32. Sudeep Dasari, Jianren Wang, Joyce Hong, Shikhar Bahl, Yixin Lin, Austin Wang, Abitha Thankaraj, Karanbir Chahal, Berk Calli, **Saurabh Gupta**, David Held, Lerrel Pinto, Deepak Pathak, Vikash Kumar, and Abhinav Gupta. RB2: Robotic manipulation benchmarking with a twist. In *Advances in Neural Information Processing Systems (Track on Datasets and Benchmarks) (NeurIPS)*, 2021. (Acceptance rate: 26%.)

33. Matthew Chang*, Arjun Gupta*, and **Saurabh Gupta**. Learning value functions from undirected state-only experience. In *International Conference on Learning Representations (ICLR)*, 2022. (Acceptance rate: 33%.). *denotes equal contribution.

34. Mohit Goyal, Sahil Modi, Rishabh Goyal, and **Saurabh Gupta**. Human hands as probes for interactive object understanding. In *Computer Vision and Pattern Recognition (CVPR)*, 2022. (Acceptance rate: 25%.).

35. Sohrab Madani, Junfeng Guan, Waleed Ahmed, **Saurabh Gupta**, and Haitham Hassanieh. Radatron: Accurate detection using multi-resolution cascaded MIMO radar. In *European Conference on Computer Vision (ECCV)*, 2022. (Acceptance rate: 28%.).

36. Gabriel Sarch, Zhaoyuan Fang, Adam Harley, Paul Schydlo, Michael Tarr, **Saurabh Gupta**, and Katerina Fragkiadaki. TIDEE: Room reorganization using visuo-symbolic common sense priors. In *European Conference on Computer Vision (ECCV)*, 2022. (Acceptance rate: 28%.).

37. Aditi Partap, Samuel Grayson, Muhammad Huzaifa, Sarita Adve, Brighten Godfrey, **Saurabh Gupta**, Kris Hauser, and Radhika Mittal. On-device cpu scheduling for robot systems. In *International Conference on Intelligent Robots and Systems (IROS)*, 2022. (Acceptance rate: 48%.).

38. Arjun Gupta, Max Shepherd, and **Saurabh Gupta**. Predicting motion plans for articulating everyday objects. In *International Conference on Robotics and Automation (ICRA)*, 2023. (Acceptance rate: 43%.).

39. Matthew Chang and **Saurabh Gupta**. One-shot visual imitation via attributed waypoints and demonstration augmentation. In *International Conference on Robotics and Automation (ICRA)*, 2023. (Acceptance rate: 43%.).

40. Shaowei Liu, **Saurabh Gupta***, and Shenlong Wang*. Building rearticulable models for arbitrary 3d objects from 4d point clouds. In *Computer Vision and Pattern Recognition (CVPR)*, 2023. (Acceptance rate: 26%.). *denotes equal contribution.

41. Hailan Shanbhag, Sohrab Madani, Akhil Isanaka, Deepak Nair, **Saurabh Gupta**, and Haitham Hassanieh. Contactless material identification with millimeter wave vibrometry.

In *International Conference on Mobile Systems, Applications and Services (MobiSys)*, 2023.

42. Xiaoyu Zhang and **Saurabh Gupta**. Push past green: Learning to look behind plant foliage by moving it. In *Conference on Robot Learning (CoRL)*, 2023. (Acceptance Rate: 39.9%.).

43. Shaowei Liu, Yang Zhou, Jimei Yang, **Saurabh Gupta***, and Shenlong Wang*. Object-centric contact field for grasp generation. In *International Conference on Computer Vision (ICCV)*, 2023. *denotes equal contribution.

44. Matthew Chang, Aditya Prakash, and **Saurabh Gupta**. Look ma, no hands! agent-environment factorization of egocentric videos. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. (Acceptance Rate: 26.1%.).

JOURNAL PUBLICATIONS

1. **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

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**, Varun Tolani, James Davidson, Sergey Levine, Rahul Sukthankar, and Jitendra Malik. Cognitive mapping and planning for visual navigation. *International Journal of Computer Vision (IJCV)*, 2020

TECHNICAL REPORTS AND PREPRINTS

1. 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

2. **Saurabh Gupta**, Bharath Hariharan, and Jitendra Malik. Exploring person context and local scene context for object detection. *arXiv preprint arXiv:1511.08177*, 2015

3. **Saurabh Gupta** and Jitendra Malik. Visual semantic role labeling. *arXiv preprint arXiv:1505.04474*, 2015

4. 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

5. **Saurabh Gupta**, David Fouhey, Sergey Levine, and Jitendra Malik. Unifying map and landmark based representations for visual navigation. *arXiv preprint arXiv:1712.08125*, 2017

6. 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

7. 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

8. Aditya Prakash, Matthew Chang, Matthew Jin, and **Saurabh Gupta**. Learning hand-held object reconstruction from in-the-wild videos. *arXiv*, 2305.03036, 2023.

PATENTS

1. Jianfeng Gao, Xiaodong He, **Saurabh Gupta**, Geoffrey G Zweig, Forrest Iandola, Li Deng, Hao Fang, Margaret A Mitchell, John C Platt, Rupesh Kumar Srivastava, and others. Discovery of semantic similarities between images and text, 2017. US Patent 9,836,671

2. Rahul Sukthankar, **Saurabh Gupta**, James Christopher Davidson, Sergey Vladimir Levine, and Jitendra Malik. Agent navigation using visual inputs, 2021. US Patent App. US11010948B2

3. Junfeng Guan, Seyedsohrab Madani, Suraj S Jog, Haitham Al Hassanieh, and **Saurabh Gupta**. Neural network-based millimeter-wave imaging system, 2021. US Patent App. 17/124,637