

Arjun Majumdar

<https://arjunmajum.github.io>

EDUCATION

Georgia Institute of Technology
PhD, Computer Science, 2019 – present
Advisor: Dhruv Batra

University of Illinois at Urbana/Champaign
MS, Electrical and Computer Engineering, 2007
Concentration: Computer Vision

University of Illinois at Urbana/Champaign
BS, Electrical and Computer Engineering, 2005
Concentration: Digital Signal Processing

EXPERIENCE

Meta AI (FAIR) Summer/Fall 2023
Research Intern with Franziska Meier and Aravind Rajeswaran

- Developed a modern formulation of the Embodied Question Answering (EQA) task and curated a benchmark dataset (OpenEQA) to evaluate state-of-the-art vision-language models (VLMs).

Google Robotics Summer/Fall 2022
Student Researcher with Leonidas Guibas and Fei Xia

- Designed a robotic system using pretrained foundation models that can find objects in real-world environments through a conversation with a user about the object’s attributes and appearance.

Amazon Alexa AI Summer/Fall 2021
Research Intern with Jesse Thomason and Gaurav Sukhatme

- Explored new techniques for learning visual navigation policies with reinforcement learning using only sparse rewards.

Facebook AI Research (FAIR) Summer/Fall 2020
Research Intern with Ross Girshick

- Designed object-centric, self-supervised visual representation learning algorithms.

MIT Lincoln Laboratory 2014 – 2019
Technical Staff – Intelligence and Decision Technologies Group

- Secured funding for a multi-year computer vision program. As a principal investigator (PI), I directed research on visual question answering, active learning, image-to-image translation, and semantic segmentation resulting in publications at CVPR, ICASSP, and EMBC.

Northrop Grumman 2007 – 2014
Fellow Engineer – Automated Sensor Exploitation Technology Center

- Led computer vision and machine learning projects focused on various applications including object detection, multi-object tracking, and various 3D point cloud processing tasks such as helicopter landing zone estimation.
- Designed an efficient planning algorithm for a distributed set of video and radar systems tasked with recognizing activities (e.g., two people meeting). The main idea was to approximate the value function in a POMDP framework with Monte-Carlo sampling and online learning.

- Developed a novel person re-identification algorithm in which feature representations and a similarity measure were learned in an offline setting. The method matched the performance of existing online learning methods without depending on complex bootstrapping schemes.

U.S. Army Corps of Engineers

2005 – 2007

Research Assistant in the Construction Engineering Research Laboratory

AWARDS

Adobe Research Fellowship Finalist 2022
 Best Paper Honorable Mention Award – Language in Reinforcement Learning (LaReL) 2020
 MIT Lincoln Laboratory Director’s Office Technical Presentation Award
 Northrop Grumman Copernicus Innovation Prize (Included 6 months of research funding)
 Northrop Grumman Timely Awards Program (15 Awards)

PUBLICATIONS

- [20] **Arjun Majumdar***, Anurag Ajay*, Xiaohan Zhang*, Pranav Putta, Sriram Yenamandra, Mikael Henaff, Sneha Silwal, Paul Mccvay, Oleksandr Maksymets, Sergio Arnaud, Karmesh Yadav, Qiyang Li, Ben Newman, Mohit Sharma, Vincent Berges, Shiqi Zhang, Pulkit Agrawal, Yonatan Bisk, Dhruv Batra, Mrinal Kalakrishnan, Franziska Meier, Chris Paxton, Sasha Sax, and Aravind Rajeswaran, “OpenEQA: Embodied Question Answering in the Era of Foundation Models,” 2023
- [19] 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[†], and Oleksandr Maksymets[†], “What do we learn from a large-scale study of pre-trained visual representations in sim and real environments?” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2024
- [18] **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[†], and Franziska Meier[†], “Where are we in the search for an Artificial Visual Cortex for Embodied Intelligence?” in *Conference on Neural Information Processing Systems (NeurIPS)*, 2023
- [17] **Arjun Majumdar**, Fei Xia, Brian Ichter, Dhruv Batra, and Leonidas Guibas, “FindThis: Language-Driven Object Disambiguation in Indoor Environments,” in *Conference on Robot Learning (CoRL)*, 2023
- [16] Philipp Wu, **Arjun Majumdar**[†], Kevin Stone[†], Yixin Lin[†], Igor Mordatch, Pieter Abbeel, and Aravind Rajeswaran, “Masked Trajectory Models for Prediction, Representation, and Control,” in *International Conference on Machine Learning (ICML)*, 2023
- [15] Zijiao Yang, **Arjun Majumdar**, and Stefan Lee, “Behavioral Analysis of Vision-and-Language Navigation Agents,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023
- [14] Karmesh Yadav*, **Arjun Majumdar***, Ram Ramrakhya, Naoki Yokoyama, Alexei Baevski, Zolt Kira, Oleksandr Maksymets, and Dhruv Batra, “OVRL-V2: A simple state-of-art baseline for ImageNav and ObjectNav,” *arXiv preprint arXiv:2303.07798*, 2023

- [13] Karmesh Yadav, Ram Ramrakhya, **Arjun Majumdar**, Vincent-Pierre Berges, Sachit Kuhar, Dhruv Batra, Alexei Baevski, and Oleksandr Maksymets, “Offline Visual Representation Learning for Embodied Navigation,” in *ICLR Workshop on Reincarnating Reinforcement Learning*, 2023
- [12] **Arjun Majumdar***, Gunjan Aggarwal*, Bhavika Devnani, Judy Hoffman, and Dhruv Batra, “ZSON: Zero-Shot Object-Goal Navigation using Multimodal Goal Embeddings,” in *Conference on Neural Information Processing Systems (NeurIPS)*, 2022
- [11] Abhinav Moudgil, **Arjun Majumdar**, Harsh Agrawal, Stefan Lee, and Dhruv Batra, “SOAT: A Scene- and Object-Aware Transformer for Vision-and-Language Navigation,” in *Conference on Neural Information Processing Systems (NeurIPS)*, 2021
- [10] Peter Anderson, Ayush Shrivastava, Joanne Truong, **Arjun Majumdar**, Devi Parikh, Dhruv Batra, and Stefan Lee, “Sim-to-Real Transfer for Vision-and-Language Navigation,” in *Conference on Robot Learning (CoRL)*, 2020
- [9] **Arjun Majumdar**, Ayush Shrivastava, Stefan Lee, Peter Anderson, Devi Parikh, and Dhruv Batra, “Improving Vision-and-Language Navigation with Image-Text Pairs from the Web,” in *European Conference on Computer Vision (ECCV)*, 2020
(Spotlight; Selection Rate: $265/5025 = 5.3\%$)
- [8] Jacob Krantz, Erik Wijmans, **Arjun Majumdar**, Dhruv Batra, and Stefan Lee, “Beyond the Nav-Graph: Vision-and-Language Navigation in Continuous Environments,” in *European Conference on Computer Vision (ECCV)*, 2020
- [7] David Mascharka*, Philip Tran, Ryan Soklaski, and **Arjun Majumdar***, “Transparency by Design: Closing the Gap Between Performance and Interpretability in Visual Reasoning,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018
- [6] Miriam Cha, **Arjun Majumdar**, H. T. Kung, and Jarred Barber, “Improving SAR Automatic Target Recognition Using Simulated Images Under Deep Residual Refinements,” in *IEEE International Conference on Acoustics Speech and Signal Processing (ICASSP)*, 2018
- [5] **Arjun Majumdar**, Laura J. Brattain, Brian A. Telfer, Chad Farris, and Jonathan Scalera, “Detecting Intracranial Hemorrhage with Deep Learning,” in *International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2018
- [4] Mark Hernandez, Adam Brewster, Larry Thul, Brian A. Telfer, **Arjun Majumdar**, Heejin Choi, Taeyun Ku, Kwanghun Chung, and Laura J. Brattain, “Learning-Based Long-Range Axon Tracing in Dense Scenes,” in *IEEE International Symposium on Biomedical Imaging (ISBI)*, 2018
- [3] Christian Anderson, Paul Breimyer, Stephanie Foster, Kelly Geyer, J Daniel Griffith, Andrew Heier, **Arjun Majumdar***, Danelle C Shah, Olga Simek, Nicholas Stanisha, and Frederick R. Waugh, “A Network Science Approach to Open Source Data Fusion and Analytics for Disaster Response,” in *IEEE International Conference on Information Fusion (Fusion)*, 2015,
*corresponding author
- [2] Michael K. McInerney, Jonathan C. Trovillion, Robert Lozar, Tarek Abdallah, and **Arjun Majumdar**, “Classifying Infrastructure Using Thermal IR Signatures,” in *American Society for Photogrammetry and Remote Sensing (ASPRS) Conference*, 2006

- [1] John Carlyle, Michael K. McInerney, Sean Morefield, Vincent Hock, and **Arjun Majumdar**, “Acoustic Measurement of Tension in Steel Dam Tainter Gate Tendon Rods and Lock Drag Links,” in *31st PIANC International Navigation Congress*, 2006

SELECTED
INVITED
TALKS

- “Keynote Talk: Visual Cortex and OpenEQA”
6th Robot Learning Workshop at NeurIPS 2023 (jointly with Dhruv Batra)
- “Attention and Transformers”
Georgia Tech Deep Learning Course, FA 2020, SP 2021, FA 2021, SP 2022, SP 2023
- “Computer on Watch – Applying Deep Learning to Imagery Analysis”
MIT Advanced Research and Technology Symposium (ARTS), 2018
- “Artificial Intelligence for National Security”
MIT Lincoln Laboratory ISR Workshop, 2018
MIT Lincoln Laboratory Staff Seminar, 2017
- “A Network Science Approach to Open Source Data Fusion and Analytics for Disaster Response”
International Conference on Information Fusion (FUSION), 2015

WORKSHOP
ORGANIZING

- “2nd Workshop on Pre-Training for Robot Learning (PRL)”
Conference on Robot Learning (CoRL), 2023

TEACHING

- Georgia Institute of Technology
Head Teaching Assistant – Computer Vision, Fall 2020
Teaching Assistant – Computer Vision, Fall 2019
- MIT Beaver Works Summer Institute
Guest Lecture – “Medical Image Segmentation.” Medlytics, 2018
*Guest Lecture – “Object Detection.” Cog*Works, 2017*
- MIT Lincoln Laboratory Technical Education
Co-Instructor – Text Analysis, Summer 2016

MENTORSHIP

- Pranav Putta – Georgia Tech BS *Fall 2023* (publication #20)
Gunjan Aggarwal – Georgia Tech MS *Spring 2022* (publication #12, next Adobe)
Bhavika Devnani – Georgia Tech MS *Spring 2022* (publication #12, next Apple)
Sachit Kuhar – Georgia Tech MS *Fall 2021* (publication #13,)
Akhil Goel – Georgia Tech BS *Fall 2021*
Abhinav Moudgil – Georgia Tech Intern *Spring 2021* (publication #11, next PhD Mila)
Ceslee Montgomery – Georgia Tech MS *Fall 2020* (next Google AI Residency)
David Mascharka – MIT LL Intern *Summer 2016 and 2017* (publication #7, next Covariant)

REVIEWING

- CVPR 2023, NeurIPS 2023, PT4R 2023, ECCV 2022, ICCV 2021, ICML 2021/2022, ViGIL 2019/2021, RA-L 2021