

# Akshay DHARMAVARAM

[www.akshayd.com](http://www.akshayd.com) [email@akshayd.com](mailto:email@akshayd.com) [github.com/Aks-Dmv](https://github.com/Aks-Dmv) [LinkedIn](#)

## EDUCATION

---

**Carnegie Mellon University (CMU) - School of Computer Science**, Pittsburgh, USA Aug 2022  
M.S. in Robotics (Fully-Funded Research Master's) GPA: 4.22/4.00

**Birla Institute of Technology and Science, Pilani (BITS)**, Goa, India May 2020  
B.E. (Hons.) Double Majored in Computer Science and Electrical Engineering GPA: 3.73/4.00

## SKILLS AND AWARDS

---

**Expertise:** Machine Learning, Computer Vision, Reinforcement Learning, Deep Learning, Generative Models, Time-Series,  
**Languages & Tools:** Python, Java, C++, Shell, PyTorch, Tensorflow, Keras, Scikit-learn  
**Scholarships & Awards:** Graduate Tuition Waiver + Stipend (2020-22) [**Top 5%**], Senior Thesis Research Scholarship (2019) [**Top 1%**], Excellency in TAship (2017-18) [**Top 1%**], Undergraduate Tuition Waiver (2016) [**Top 2%**]

## SELECTED RESEARCH EXPERIENCE

---

**Advanced Agent Robotics Technology Lab**, Robotics Institute, CMU JAN 2020 - PRESENT  
Graduate Research Assistant | **Master's Thesis:** Self-Supervised Imitation Learning Advisor: [Prof. Katia Sycara](#)  
- Collaborated with [DARPA](#) to develop a novel Imitation Learning framework for behavior prediction and control  
- Reduced policy training time by 70% for predicting and controlling multi-modal future trajectories of agents  
- Achieved a 3x and 10x reduction in mean & variance in expert imitation by designing a novel curriculum  
- Improved zero-shot trajectory matching by 12x by creating a new graph neural network for reinforcement learning  
- Beat state-of-the-art trajectory prediction by 90% in a noisy and open-ended setting with multiple interacting agents  
**Under Review at an A\* conference** [[Pre-Print](#)]

**Stochastic Systems Laboratory**, Indian Institute of Science (IISc), Bangalore MAY 2019 - MAY 2020  
Undergraduate Research Assistant | **Bachelor's Thesis:** Average Reward Option-Critic Advisor: [Prof. Shalabh Bhatnagar](#)  
- Formulated a hierarchical reinforcement learning framework that can learn optimal strategies for repetitive tasks  
- Developed a mathematical framework to prove its convergence and demonstrated a 15% increase in rewards/cycle  
- Improved the sample-efficiency by 2.8x for data expensive settings by mathematically correcting gradient updates  
- Reduced training time by 95% by distributing the sampling and gradient calculations across 128 cores, using A3C  
**Published at [AAAI 2020](#)**

## SELECTED PROJECTS

---

**Sketch-to-Image Generation with GANs** [[Code](#)] [[Report](#)] CMU | 2021  
- Designed novel method for sketch-to-image generation that has improved gradient propagation and faster convergence  
- Developed image-to-image translation model between GAN latent-spaces and sketches that works with unpaired images  
- Beat state-of-the-art accuracy by 47% in just 50% of training steps, by avoiding overfitting and eliminating spurious features

**Object-Detection in the Absence of Ground-Truth Bounding-Boxes** [[Code](#)] CMU | 2021  
- Developed a PyTorch implementation of "Weakly Supervised Deep Detection Networks" for object localization  
- Increased accuracy by 21% against a strong baseline and generated human-interpretable heatmaps for visualization

**Visual Question Answering (VQA) using Hierarchical Question-Image Co-Attention** [[Code](#)] CMU | 2021  
- Created an alternating attention architecture for question answering over images from the [PASCAL VOC 2007](#) dataset  
- Improved relevance of answers to questions and images, thereby increasing accuracy by 26% over standard baselines

**Tackling Causal Misidentification in Robotic Manipulation** [[Report](#)] CMU | 2020  
- Demonstrated zero-shot transfer from simulation to real-world by identifying causal features by exploration in simulation  
- Interfaced to the Gazebo simulator using PyRobot and collected counterfactual data from simulation, based on curiosity  
- Deployed a low-cost object-detection pipeline on a LoCoBot and successfully pick-and-placed boxes into a bin

## PUBLICATIONS

---

- 1) **Dharmavaram, A;** Gupta, T; Li, J; Sycara, K "SS-MAIL: Self-Supervised Multi-Agent Imitation Learning" (*Under Review*) [[Arxiv](#)]
- 2) **Dharmavaram, A;** Riemer, M, Bhatnagar, S "Hierarchical Average Reward Policy Gradient Algorithms" [AAAI 2020](#) [[Arxiv](#)] [[OJS](#)]