

# Aravind Venugopal

GitHub || LinkedIn || Google Scholar || Website

User ID: [avenugo2@andrew.cmu.edu](mailto:avenugo2@andrew.cmu.edu)

Mobile: +1-412-273-5354

## EDUCATION

---

### Carnegie Mellon University

*M.S. in Machine Learning*

Pittsburgh, PA, USA

*Grad. Dec. 2023, CGPA: 4.11/4*

### BITS Pilani K. K. Birla Goa Campus

*B.E. (Hons) in C.S. and M.Sc. (Hons) in Economics*

Goa, India

*Grad. Aug. 2020, CGPA : 8.53/10 (85.3%)*

### Loyola School

*Indian School Certificate (ISC), 12th Grade*

Trivandrum, India

*Grad. May 2015, 97.2%*

## EXPERIENCE

---

### Graduate Researcher

*Carnegie Mellon University | Advisors: Prof. Fei Fang, Prof. Jeff Schneider*

Aug. 2022 - Present

*Pittsburgh, PA*

- Developed a hierarchical latent variable model, Multi-Agent Bi-Level Model (MABL) for multi-agent RL.
- Extended MABL to learn temporally abstracted action and state sequence representations for RL.
- Working on Denoising Probabilistic Diffusion Models to learn stochastic dynamics from real-world, offline RL data.

### Machine Learning Intern

*Blue River Technology, L.L.C, Manager: Ben Cline*

May. 2023 - Aug. 2023

*Santa Clara, CA*

- Designed a real-time 3D Transformer framework for object detection on autonomous tractors in the field.
- Constructed large-scale pseudo-label datasets and pipelines applying classical computer vision.

### Post Baccalaureate Research Fellow

*RBCDSAI, Indian Institute of Technology | Advisor: Prof. Balaraman Ravindran*

Aug. 2020 - June 2022

*Madras, India*

- Developed, CombSGPO, which uses game theory and RL to combat wildlife poaching.
- Developed an RL approach for empirical comparison with evolutionary algorithms in Green Security Games (GSGs), collaborating with Prof. Jacek Mańdziuk.
- Developed an optimization framework to achieve robust equilibrium performance in Markov games.

### Research Intern

*RBCDSAI, Indian Institute of Technology | Advisor: Prof. Balaraman Ravindran*

Feb. 2020 - Aug. 2020

*Madras, India*

- Develop an RL algorithm for resource allocation in GSGs using contextual, combinatorial multi-armed bandits.
- Worked on optimizing defender patrolling strategies in GSGs, using Neural Fictitious Self Play.

### Research Intern (Thesis)

*Robotics Institute, Carnegie Mellon University | Advisor: Prof. Cameron Riviere*

Sept. 2019 - Feb. 2020

*Pittsburgh, PA*

- Developed a real-time virtual fixture strategy for Micron, a handheld surgical robotic tool, to avoid blood vessels during neurosurgery and conducted real-world trials.

### Research Intern

*NearLab, Politecnico di Milano | Advisor: Prof. Elena De Momi*

July 2019 - Sept. 2019

*Milan, Italy*

- Worked towards vessel avoidance using Micron by implementing accurate and fast deep-learning algorithms for real-time, intra-operative blood vessel segmentation.

## RESEARCH PROJECTS

---

### Diffusion Models for RL-based Tokamak Control

*Advisor: Prof. Jeff Schneider*

- Developing approaches based on diffusion models accurately learn noisy and stochastic plasma dynamics from real offline Tokamak data.
- Exploring the use of diffusion models as generalizable planners to create robust policies capable of out-of-distribution generalization.

### Latent Variable Models for Multi-Agent Reinforcement Learning

*Advisors: Prof. Fei Fang, Prof. Balaraman Ravindran*

- Developed a latent variable world model for MARL, LaVa, in collaboration with Prof. Elizabeth Bondi, with potential applications to real-world scenarios ([paper](#)).
- Developed a scalable hierarchical latent-variable model, MABL, that disentangles agent-centric and global information, respectively, from multi-agent trajectories ([paper](#)).
- Exploring temporally abstract hierarchical models to use as representation learners for hierarchical RL.

## Optimization for N-player Games

Advisor: Prof. Balaraman Ravindran

- Developed and implemented an actor-critic algorithm that provably converges to SSE, collaborating with Prof. Elizabeth Bondi and Prof. Fei Fang.
- Evaluated the algorithm empirically on Stackelberg games and showed improved sample efficiency and rate of convergence.
- Derived and implemented a framework that uses pair-wise bilinear approximations of objectives to solve for NE in fully cooperative Markov games.

## Reinforcement Learning for Unified Allocation and Patrolling in Green Security Games

Advisor: Prof. Balaraman Ravindran

- Developed Combined Security Game Policy Optimization (CombSGPO), combining defender strategies for resource allocation and patrolling in GSGs, collaborating with Elizabeth Bondi and Prof. Milind Tambe ([paper](#)).
- Achieved **3x** better utility and **2.5x** better scalability than state-of-the-art RL and MILP-based approaches.
- CombSGPO uses action representation learning, competitive optimization and MARL. It has applications to combat wildlife poaching, illegal logging, illegal fishing etc. ([code](#))

## Real-time 3D Object Detection with Visual Transformers

Advisor: Ben Cline

- Developed a transformer-based 3D object detection model to detect stopclass objects for autonomous tractors.
- Developed all stages of the 3D model pipeline from 3D dataset creation, camera calibration and 3D labelling workflows to model training, modification and evaluation.
- Achieved 3D bounding box prediction with a median error of **40 cm** at **25 fps**.

## Real-time Segmentation for Vessel Avoidance with Surgical Robot

Advisors: Prof. Cameron Riviere and Prof. Elena De Momi

- Developed an algorithm integrating real-time segmentation and real-time 3D reconstruction of a blood vessel with a virtual-fixture-based control strategy. ([paper](#))
- Built segmentation datasets of intra-operative neurosurgery images for training and evaluation.
- Used the control strategy for real-time vessel avoidance by interfacing with Micron, a surgical robot, and conducted real-world trials.
- Achieved vessel avoidance at **33.3 fps** enabling safe operation at up to **400  $\mu\text{m}$**  near a vessel. ([code](#) | [data](#))

## Undergraduate Projects

- Visual Question Answering on CLEVR dataset using Stacked Attention Networks (SANs). ([code](#))
- Locomotion tasks with Turtlebot3 robots using ROS. ([code](#))
- Predicting Party Behavior in 2019 Indian Parliamentary Elections from electoral data using game theory and machine learning. ([report](#))

## PUBLICATIONS

---

Aravind Venugopal, Stephanie Milani, Fei Fang and Balaraman Ravindran. **Bi-level Latent Variable Model for Sample-Efficient Multi-Agent Reinforcement Learning**. (*under review, AAMAS 2024*). ([paper](#))

Aravind Venugopal, Elizabeth Bondi, Fei Fang and Balaraman Ravindran. **LaVa: Latent Variable Models for Sample Efficient Multi-Agent Reinforcement Learning**. In *RLDM*, June 2022. ([paper](#))

Adam Zychowski, Jacek Mańdziuk, Elizabeth Bondi, Aravind Venugopal, Milind Tambe and Balaraman Ravindran. **Evolutionary Approach to Security Games with Signaling**. In *IJCAI-ECAI*, July 2022. ([paper](#))

Aravind Venugopal, Sara Moccia, Arpita Routray, Simone Foti, Elena De Momi and Cameron N. Riviere. **Real-time vessel segmentation and reconstruction for virtual fixtures for an active handheld micro-neurosurgical instrument**. In *International Journal of Computer Assisted Radiology and Surgery*, 2022. ([paper](#))

Aravind Venugopal, Elizabeth Bondi, Harshavardhan Kamarthi, Keval Dholakia, Balaraman Ravindran and Milind Tambe. **Reinforcement Learning for Unified Allocation and Patrolling in Signaling Games with Uncertainty**. In *AAMAS*, May 2021. ([paper](#), [media](#))

Aravind Venugopal, Sara Moccia, Arpita Routray, Simone Foti, Elena De Momi and Cameron N. Riviere. **Real-time Compensation of Handheld Surgical Robotic Tool for Safer Neurosurgery**. In *Workshop on “Integrating Sensor Fusion and Perception for Human-robot Interaction” at IEEE RO-MAN*, September 2020. ([paper](#))

---

## TECHNICAL SKILLS

**Languages:** Python, C++

**Tools and Frameworks:** PyTorch, Tensorflow, ROS, Visual Studio, Git

**Relevant Coursework (Masters)** Reinforcement Learning (10703), Probabilistic Graphical Models (10708), Computer Vision (16720), Convex Optimization (10725), Probability and Mathematical Statistics (36700), Advanced Introduction to Machine Learning (10715), Machine Learning in Practice (10718)

**Relevant Coursework (Undergrad):** Probability and Statistics, Econometrics, Machine Learning, Neural Networks and Fuzzy Logic, Logic in Computer Science, Artificial Intelligence, Reinforcement Learning, Data Structures and Algorithms, Object Oriented Programming

---

## TEACHING AND EXTRA-CURRICULARS

**Teaching Assistant** Aug. 2022 - Dec 2022  
*Carnegie Mellon University* Pittsburgh, PA  
• Teaching Assistant for Exec. Education, 17759: Advanced Topics in ML and Game Theory, taught by Prof. Fei Fang

**Teaching Assistant** Jan. 2022 - May 2022  
*Indian Institute of Technology* Madras, India  
• Teaching Assistant for CS 6700: Reinforcement Learning, taught by Prof. Balaraman Ravindran

**Operations Head** Oct. 2017  
*Waves* Goa, India  
• Led a 3-tier team of 50+ volunteers to organize 25+ events at Waves, a national level inter-collegiate competition organized by BITS Pilani, Goa.

**Registered Player, Division I League** Sept.2016 - Sept. 2020  
*Goa Football Association* Goa, India  
• Registered player contract in Division I Goa Football League under the purview of GFA - Goa Football Association.

---

## ACADEMIC SERVICES

**Subreviewer:** AAMAS 2024

---

## MERITS

**2011: National Talent Search Scholar** (0.5% acceptance out of over 1,000,000)

**2015: Innovation in Science Pursuit for Inspired Research Scholar**(top 1% nationally in ISC, 12th grade)

---

## TEST SCORES

GRE: 337 (Q: 169, V:168, W: 5.0)

IELTS: 8.5/9