Gokul M K

IDDD in Robotics, AI

ed21b026@smail.iitm.ac.in Github Website LinkedIn

Education

Indian Institute of Technology, Madras (CGPA 8.94) **Nov. 2021 – Present**
Bachelor of Technology in Engineering Design, IDDD in Robotics (Interdisciplinary Dual Degree) *Chennai, Tamil Nadu*

GRD Public School (Grade 12 CBSE (97.2 %), Grade 10 CBSE (93.8 %)) **April 2007 – April 2021**
Primary, Seconday and Higher Secondary *Coimbatore, Tamil Nadu*

Relevant Coursework

- Optimization: Theory and Algorithms
- Reinforcement Learning (CS6700)
- Recent Advances in RL (DA7400)
- Introduction to Large Language Models
- Deep Learning for Imaging (EE5179)
- Introduction to Motion Planning (ED5215)
- Programming Data Structures and Algorithms
- Probability, Statistics and Stochastic Processes
- Optimal Control (EE6415)
- Non-Linear System Analysis (EE6412)

Research Experience, Projects and Competitions

Nanyang Technological University, Singapore **July 2025 – August 2025**
Global Connect Fellow, Assoc Prof. Arvind Easwaran

- Developed generalizable and trustworthy Reinforcement Learning planning agents for the F1Tenth Racing.
- Developing an hierarchical framework for decision making which considers both high-level maneuver selection and low-level motion control for racing cases, will submit a conference paper on the same.

Qneuro India Pvt Limited, Chennai **Jan 2025 – May 2025**
AI Research Intern, CTO Dr. Rahul Bhardwaj *Chennai, Tamil Nadu*

- Developed Deep Learning models that capture spatio-temporal voltage patterns from EEG Signals to infer patient's stress levels, cognitive load, achieving a 5% higher test accuracy while including spatial features.
- Built a PCG prototype for real-time heart sound capture, generating datasets to train an Hidden Markov Model for murmur classification.

Inter-IIT TechMeet 13.0, IIT Bombay **Nov 2024 – Dec 2024**
Team Lead - BharatForge Problem Statement *Chennai, Tamil Nadu*

- Developed End-End Reinforcement Learning solution for multi-agent autonomous navigation of warehouse robots in dynamically evolving environments.
- Benchmarked various multi-agent RL algorithms, with MA-TD3 proving most sample-efficient and transferable to various scenarios.

Robert Bosch Centre for Cyber-Physical Systems (RBCCPS) - IISc Bangalore **May 2024 – July 2024**
Robotic Summer Intern, HiRo Lab, Dr.Ravi Prakash *Bangalore, Karnataka*

- Developed a Learning-Based Approach for **Bimanual Robotic Manipulation** to Toss Objects Efficiently.
- Worked on implementing one of the prominent research from Google AI, TossingBot, which is a deep learning based grabbing and throwing in a cluttered scenario.

Team Anveshak, Mars Rover Team, IIT Madras **June 2022 – July 2024**
Embedded System Lead, Electronics and Software Engineer *Chennai, Tamil Nadu*

- Worked on embedded systems of the rover's manipulator, gaining experience in control systems and CAN integration. Worked with various **Visual SLAM** algorithms to address visual odometry and mapping.
- Coded the rover's manipulator in simulation to evaluate its functionality using the **Moveit** Framework for tasks such as obstacle avoidance with **OctoMap** and grasp planning.

Eyantra Summer Intern, IIT Bombay **May 2023 – July 2023**

Exploring Various Algorithms for Grasping Unknown Objects using a two finger gripper *Mumbai, Maharashtra*

- Conducted research on the **Point Cloud Library**, harnessing its robust features to devise a lightweight grasp detection method. This approach estimates grasp poses based on Euclidean Clustering and K-D Trees.
- Integrated various grasping algorithms, including Graspnet, HAF, and GPD, with the Industrial UR5 arm and 2-finger gripper using **ROS**. Implemented these algorithms in both simulation and in hardware setups.

Eyantra Robotics Competition 2022-23, IIT Bombay **Sep. 2022 – Mar. 2023**

AIR-4 Krishibot Theme *Mumbai, Maharashtra*

- Programmed an Autonomous Ground Vehicle (AGV) for greenhouse navigation and fruit harvesting, using LiDAR-based PD control for lane keeping and OpenCV-based centroid detection for target localization.

PROJECTS

Reasoning Enhancement of Base LLM using GRPO and TPO (Personal Project) **Oct. 2025 - Present**

- Fine-tuned a base LLaMA model using **GRPO** to produce structured chain-of-thought from Reasoning Gym dataset, enforcing to reason in a $< think >< /think >$ and $< final > .. < /final >$ output format.
- Further optimizing reasoning quality with **TPO(DPO)** preference alignment, comparing good vs. bad thought sequences to refine the model's ability to reason correctly and consistently.

Model-Based Control, Koopman Learning for Quadruped Locomotion | *Prof. Anuj Tiwari* **Mar. 2025 - Present**

- Working on implementing a Convex Model Predictive Control framework for a quadruped robot based on a state-of-the-art formulation, enabling real-time gait optimization across varying terrains.
- Designing a modular control pipeline with plans to integrate Koopman operator-based learning for capturing nonlinear dynamics, enabling smooth transitions from model-based to data-driven control.

Denoising and Deblurring MVTEC AD Dataset | *EE5179 Course Project* **July. 2024 - Nov. 2024**

- This project was done in collaboration with KLA to develop deep learning models for denoising and deblurring MVTEC images. The goal was to achieve good PSNR and SSIM, ensuring high-quality image restoration.
- RIDNet, a deep learning-based model, was employed for this task and achieved a PSNR of 32.8 on the training set and 34.78 on the test set, as evaluated by KLA.

Implicit Reinforcement without Interaction at Scale (IRIS) | *Prof. Balaraman Ravindran* **July. 2024 - Nov. 2024**

- The project deals with addressing sub-optimality and diversity challenges in large datasets for long-horizon manipulation tasks using **Offline Reinforcement Learning**.
- The project involves the basic implementation of IRIS in a few toy environments with dataset, and fine-tuning the architecture to process image-based datasets (like DREAMER) for running IRIS in latent space.

OPEN SOURCE PROJECTS

Project Groot: Designing and Training Transformer LM from First Principles **Dec. 2025 - Present**

- Currently building a collection of Transformer-based language models (GrootTiny, GrootSmall, and larger variants) from scratch, focusing on architecture design and stable training.
- Pretraining models on narrative and encyclopedic datasets and applying instruction fine-tuning to improve fluency, coherence, and factual grounding.

SCHOLASTIC ACHIEVEMENTS

- Achieved an **All India Rank of 6253** in the **JEE Advanced 2021** examination, surpassing a competitive pool of 1.5 lakh students nationwide.
- Attained the highest ranking within my school in the Grade 12 CBSE examination.

TECHNICAL SKILLS

Languages: Python, C, C++, MySQL, Rust, MATLAB, AVR Assembly Language

Software Tools: ROS/ROS2, Git, Docker, Kubernetes, Fusion360, Abaqus, Ansys, Altium

Framework Libraries: Pytorch, Jax, Tensorflow, Scikit-learn, Mujoco, Bullet, Langchain, HuggingFace