Thivyanth M V

Latest Résumé | thivyanth | thivyanthiitb | thivyanth | thivyanth.github.io | thivyanth4@gmail.com | thivyanth@iitb.ac.in |

Work Experience

ML/Perception Engineer in Team AeRoVe of UMIC Oct 2023 - present

Develop Machine Learning Models for Perception, focusing on semantic segmentation, object detection and tracking, and integrating perception subsystems with navigation modules.

RoboDrive Competition (ICRA 2024) Feb 2024 - Apr 2024

- Participated in the RoboDrive Challenge, focusing on robust multi-modal BEV 3D object detection, using corrupted camera and LiDAR data from the NuScenes dataset.
- Achieved 4th place with scores of 42.8 (NDS) and 26.4 (mAP), showcasing the ability of our model to handle camera and LiDAR corruptions.
- Got invited to present at ICRA 2024 conference, highlighting our innovative approach inspired by FUTR3D, utilizing an architecture similar to BEVFusion.

Research Projects

Training Diffusion Models using Reinforcement Learning May 2024 - present

Under the guidance of Prof. Biplab Banerjee, I am working on training diffusion models using reinforcement learning (RL). This project builds upon the work by Kevin Black, Michael Janner, Yilun Du, Ilya Kostrikov, and Sergey Levine, described in their paper "Training Diffusion Models with Reinforcement Learning." The core objective is to optimize diffusion models for downstream tasks directly, rather than merely matching a data distribution. This involves framing denoising as a multi-step decision-making problem and applying policy gradient algorithms, referred to as denoising diffusion policy optimization (DDPO). This focuses on optimizing diffusion models for specific objectives like image compressibility, aesthetic quality, and prompt-image alignment. By implementing denoising diffusion policy optimization (DDPO), we aim to improve model performance using multi-step decision-making and feedback from vision-language models, pushing the boundaries of generative model training for applications in image synthesis and beyond.

Proximal Policy Optimization for Cart-Pole Balancing Jun 2024

Implemented Proximal Policy Optimization (PPO) for CartPole Balancing. Designed and developed a PPO-based reinforcement learning agent using Python and PyTorch, integrating Generalized Advantage Estimation (GAE) and clip loss. Set up a multi-environment training pipeline with Gym, achieving a cumulative reward of X within Y episodes. Improved training stability and efficiency, leveraging TensorBoard for detailed logging and Wandb for experiment tracking. [GitHub Repository]

Education

2022 - 2026	Bachelor's Degree of Engineering Physics at **IIT Bombay**	(CPI: 7.51/10)
2022	Class 12th CBSE	(86.4%)
2020	Class 10th CBSE	(91.8%)

Skills

Machine Learning	PyTorch, TensorFlow, HuggingFace, WandB, Tensorboard, Conda
Reinforcement Learning	TorchRL, Stable Baselines 3, OpenAI Gym, TRL - Transformer Reinforcement Learning
Computer Vision	TorchVision, MMCV, MMDet, MMSegmentation, MMDet3d, PCDet
Technical/Software	Python, LaTeX, Git, Fusion 360, Linux, WSL, Docker, HTML, CSS, JS, Bash, SSH