Thivyanth M V

thivyanth | thivyanthiitb | thivyanth | thivyanth.github.io | thivyanth4@gmail.com | thivyanth@iitb.ac.in | tvenkat3@ucsc.edu

WARNING: DO NOT SHARE - Conference rules prohibit sharing publication details before final results

PUBLICATIONS

Can Inference-Time Scaling Incentivize the Discriminative —Capability of Diffusion Models?

Oct 2024 - Feb 2025

Research Collaboration with Xuehai He (PhD Candidate, UCSC) **[Under Review at ACL-ARR]** **[First Author]**

- Investigated inference-time scaling for diffusion discriminators, revealing the inefficiency of naive step scaling for discriminative tasks and motivating the design of the "Sweet Diffusion Discriminator".
- Developed a novel **Heuristic Relative Sampling Policy (HRSP)** integrated with **Phased Consistency Models (PCMs)** to dynamically allocate computation, achieving ∼6.5x faster inference (vs. DiffusionITM).
- Demonstrated state-of-the-art zero-shot performance on diverse image/video-text retrieval benchmarks (Winoground, ImageCode, MSRVTT, ActivityNet), significantly outperforming prior diffusion-based methods.

Grounded Document QA with GRPO (Work in Progress) — Reinforcing Evidence-Grounded Reasoning

May 2025 - Present

Research Collaboration with Haozhe Wang (PhD Candidate, HKUST) **[Targeting ICLR]** **[First Author]**

- Benchmarked on MMLongBenchDoc, DUDE, and an in-house synthetic QA set derived from documents like research papers (custom document QA benchmark)
- Fine-tuned Qwen2.5-VL via SFT then GRPO to incentivize grounding on cited document evidence during reasoning
- Used ANLS (Average Normalized Levenshtein Similarity) to score generated answers against ground truth answers
- Tooling: VERL (Volcano Engine Reinforcement Learning for LLMs) and OpenRLHF for SFT/RL training and evaluation
- Research in progress; intermediate results suggest the method is effective

GrAD-BEn: Multimodal Prompting using Gradient Aligned Distillation —& Beta Ensembling for Few-shot Domain Incremental Learning

Aug 2024 - Mar 2025

Research Collaboration with Samrat Mukherjee (PhD Candidate, IIT Bombay) — Guide: Prof. Biplab Banerjee, Centre of Studies in Resources Engineering (CSRE) (IITB) **[Under Review at NeurIPS 2025]** [Second Author]

- Contributed to the development and experimentation of GrAD-BEn, a novel method for few-shot domain incremental learning using gradient aligned distillation and beta ensembling.
- Developed and implemented multiple baseline methods including Elastic Weight Consolidation (EWC), Learning Without Forgetting (LWF), and a novel gradient projection approach for continual prompt learning in CLIP models for image classification in datasets with different domains like DomainNet, Office-Home, and PACS.
- Implemented Dual Gradient Projection Memory (DualGPM) techniques to maintain orthogonal gradient spaces between current and previous tasks, effectively minimizing interference while preserving past knowledge.
- Designed and conducted comprehensive experiments across DomainNet, CoRE50, and CDDB-Hard benchmarks, achieving **state-of-the-art performance** with higher average accuracy and better forgetting alleviation compared to existing methods.

Morpho4D: An Interactive, Controllable, and Editable Text-to-4D World —Simulator

Jan 2025 - Mar 2025

Research Collaboration with Xuehai He (PhD Candidate, UCSC) **[Targeting ICCV/CVPR]** [Second Author]

- Contributed to the development of a comprehensive **text-to-4D scene generation and editing system** with a multi-stage pipeline: (1) Command Parameterization using LLM agents to interpret user instructions, (2) Scene Generation with text-to-video diffusion models and 3D reconstruction, and (3) Scene Editing with dynamic control and static edit submodules for motion and appearance modifications.
- Implemented and optimized the **scene reconstruction module** that transforms 2D video representations into dynamic 3D Gaussian-based 4D scenes, integrating latent feature embeddings for effective scene editing and developing feature distillation techniques to maintain consistency across views and frames.
- Contributed to key components of the **Dynamic Control Submodule**, enabling precise motion direction control of objects through trajectory-aligned attention weighting and bounding box guidance to control object motion direction and speed.
- Designed and conducted comprehensive **evaluation methodologies** comparing generated 4D scenes against real-world videos from the DAVIS dataset, demonstrating comparable or superior quality across multiple metrics including BRISQUE, NIQE, CLIP Similarity, and QAlign.

PROJECTS

Improving Image-Text Alignment in Diffusion Models with Reinforcement Learning

May 2024 - Jul 2024

Guide: Prof. Biplab Banerjee, Centre of Studies in Resources Engineering (CSRE)

- Conducted a detailed literature review on the use of reinforcement learning (RL) for optimizing diffusion models, specifically implementing the **Denoising Diffusion Policy Optimization (DDPO)** from *Training Diffusion Models with Reinforcement Learning*, enhancing image-text alignment using CLIP scores from *Fine-grained Image Captioning with CLIP Reward*.
- Developed and integrated the DDPO algorithm in PyTorch, treating the sequence of denoising steps as a Markov Decision Process (MDP). Implemented CLIP reward, leveraging image-text similarity scores from a pretrained CLIP model to optimize image generation towards specific textual descriptions. Enabled **LoRA finetuning** for lesser memory usage.
- Employed Hugging Face's `transformers` for efficient model training; tailored the **CLIP reward function** to dynamically adjust training objectives, ensuring the generation of images closely aligned with textual prompts.

The RoboDrive Challenge (ICRA 2024) Feb 2024 - Apr 2024

Competition Project — International Conference on Robotics and Automation (ICRA)

- Participated in the RoboDrive Challenge, focusing on **advanced multi-modal BEV 3D object detection**, addressing challenges posed by **corrupted sensor data** from cameras and LiDAR on the NuScenes dataset.
- Developed a **robust detection model** using a Modality-Agnostic Feature Sampler (MAFS) for unified multi-scale feature processing from cameras and LiDAR & a transformer-based decoder for precise iterative refinement of 3D detection outputs.
- **Secured 4th place** with scores of **42.8 (NDS)** and **26.4 (mAP)**, demonstrating the model's robustness to corruptions.
- **Got invited to present the team's findings and methodologies** at ICRA 2024, underscoring our innovative approach that integrates concepts from FUTR3D and BEVFusion, leading to advancements in sensor fusion and 3D detection.

Generalist Robotics Policy (GRP) Implementation Oct 2024

Self-Project — ML Implementation Project

- Implemented a **Generalist Robotics Policy (GRP) model** using **Vision Transformers** for multi-modal robotic control, integrating image, text, and goal image inputs to generate continuous action outputs. Programmed key components including **patch embedding**, **positional encoding**, and **multi-head attention mechanisms**, adapting Transformer architectures for robotics applications.
- Implemented a comprehensive training pipeline with **AdamW optimizer** and **MSE loss**, along with simulation-based testing. GitHub Repository: `thivyanth/grp`

Proximal Policy Optimization for Cart-Pole Balancing Jun 2024

Self-Project — ML Implementation Project

- Developed a robust reinforcement learning model using **Proximal Policy Optimization** (PPO) tailored for the **CartPole balancing problem**, focusing on advanced policy learning strategies.
- Engineered an agent using Python and PyTorch, integrating techniques such as **Generalized Advantage Estimation** (GAE) and **clip loss** to optimize performance and stability.
- Configured a vectorized training environment using **Gym** to facilitate **efficient and parallel simulations**, significantly improving the model's performance metrics.

OTHER PROJECTS

Smart Appliance IR Controller Mar 2024 - Apr 2024

Course Project: Digital Electronics and Microprocessor — Guide: Prof. Maniraj Mahalingam

- Developed a **Smart Appliance Infrared-Remote Controller** using Arduino and Arduino-based components.
- **Integrated 2 Arduino units** via Serial Communication protocol, enabling the storage of signals for over 5 different devices.
- Utilized the **HC-05 Bluetooth module** to enhance functionality by enabling **remote control** via a Bluetooth device.

Obstacle Removing Line Follower Robot Nov 2022 - Feb 2023

MS 101 Course Project — Guide: Prof(s). D.K. Sharma and Joseph John (Department of Electrical Engineering)

- **Designed and developed** a line follower robot using **Arduino**, **Fusion 360** for 3D modeling, **LaserCAD** for laser-cutting templates, and **Fractory** for 3D printing, integrating these with circuit design to enhance reliability.
- Achieved recognition as a top-performing team, placing among the **top 24 out of 120** and presenting our project.

POSITIONS OF RESPONSIBILITY

Senior Engineer (Machine Learning & Localization Subsystem) Apr 2024 - Jul 2024

Unmesh Mashruwala Innovation Cell, IIT Bombay

- **Team Participation:** Part of AeRoVe, which competes in UAV competitions worldwide.
- **Technical Leadership:** Responsible for technical challenges in the ML subsystem, participating in ML challenges like RoboDrive (ICRA), Autonomous Grand Challenge (CVPR), Stranger Sections, Field Area Segmentation and more projects.
- **Recruitment Panel:** Served on a panel interviewing 20 students, selecting 2 for team membership in the subsystem.
- **Mentorship:** Guided 2 students through 5-week beginner projects, covering technical and non-technical aspects.

EDUCATION

2022 - 2026 Bachelor's Degree of Engineering Physics at **IIT Bombay**

SKILLS

Machine Learning	PyTorch, NumPy, TensorFlow, HuggingFace, WandB, Tensorboard, Transformers
Reinforcement Learning	TorchRL, Stable Baselines 3, OpenAI Gym, TRL - Transformer Reinforcement Learning
Computer Vision	TorchVision, MMCV, MMDet, MMSegmentation, MMDet3d, PCDet
Programming Languages	Python, HTML, CSS, JS, Bash
Tools and Platforms	Conda, Poetry, LaTeX, Git, Fusion 360, Linux, WSL, Docker, SSH

KEY COURSES UNDERTAKEN

CS & Robotics: Computer Programming and Utilisation, Makerspace, Artificial Intelligence and Data Science

Physics and Electronics: Quantum Physics I & II, Classical Mechanics, Analog Electronics, Digital Electronics & Microprocessors, Statistical Physics, Electromagnetic Theory, Non-linear Dynamics, Magnetism and Superconductivity, Molecular Spectroscopy and Optical Physics, Optics and Spectroscopy Lab

Math: Calculus, Linear Algebra, Differential Equations, Complex Analysis and Integral Transform, Estimation on Lie Groups, Numerical Analysis, Analytical and Geometric Dynamics

Miscellaneous: Physical, Organic & Inorganic Chemistry, Biology, Economics, Environmental Studies, Physics of Biological Systems, Biofuels: Technology and Policy Perspective

EXTRACURRICULARS

- Participated in the National Sports Organization for **Weightlifting** for a year, showcasing my grit and fitness passion.

Last updated: August 15, 2025