

Pranav Vinodh

Residence/domicile: Al Jimi Avenue, Khalidiyah, Abu Dhabi, UAE

E-mail: prnv2005@gmail.com * *Telephone number:* +91 9108589097

LinkedIn: Pranav Vinodh * *GitHub:* Pranav-Vinodh

Education

Bachelor of Technology in Information Technology *National Institute of Technology Karnataka
Btech-IT* *August 2023 - present*
provisional grade: 8.55 CGPA

High School Degree *Gems United Indian School*
High school's diploma program *April 2021 - April 2023*
Final grade: 95 percentage

Work experience

Carnegie Mellon University *May 2026 - Present*
Researcher at Xu Lab *Pittsburgh, USA*

- Developed a Hierarchical Domain Adaptation framework for few-shot Cryo-ET subtomogram classification, integrating jointly optimized input-level transformations (Spatial Transformer Networks, Difference-of-Gaussian intensity filters, and global color transforms) with feature-level alignment losses (MMD and CORAL) on a Kinetics-400 pretrained Swin3D backbone to bridge the synthetic-to-real domain gap.
- Engineered extensive parameter sweeps and ablation studies across 10 random seeds and multiple low-shot target datasets (Noble and Qiang), demonstrating that combining input-level domain adjustment with feature alignment consistently outperformed target-only baselines and improved 3-shot classification accuracy by up to +6.05%.
- Optimized few-shot generalization in challenging 5-shot regimes, achieving up to 63.50% accuracy on Noble and 58.83% on Qiang datasets (+7.16% over target-only models), while diagnosing optimization collapse, mitigating overfitting, and evaluating Adam versus AdamW convergence under strict low-data constraints.

Nanyang Technological University (NTU) *July 2025 - Present*
Research Intern (VIDA LAB) *Singapore*

- Built and released a large-scale infographic generation dataset by programmatically scraping open-source visualization repositories, performing quality filtering and structural curation, and constructing aligned image-structured instruction pairs via LLM-based semantic parsing; used the dataset for diffusion model pretraining and published it on Kaggle, Which has been cited in an independent research publication (InfoAffect, 2025). ([Infographics Dataset](#)).
- Developed and optimized an LLM-driven instruction synthesis pipeline, leveraging advanced prompt engineering and controlled generation strategies to produce hierarchical, step-wise infographic construction instructions; iteratively improved structural fidelity and semantic alignment for text-to-infographic supervision.
- Surveyed and synthesized 50 papers (IEEE VIS, IEEE TVCG, ACM CHI) to frame the methodological landscape of automated infographic generation; evaluated dataset scale, annotation density, and GPU memory constraints to design a Stable Diffusion fine-tuning strategy for domain-adaptive text-to-infographic generation.

NITK Surathkal *May 2026 - July 2026*
Computer Vision Research Intern (HALE Lab) *Mangalore, India*

- Fine-tuned and aligned Ornit 9B and 35B parameter LLMs, improving downstream task accuracy by 24% and reducing preference alignment compute costs by 50% by developing multi-stage SFT and Direct Preference Optimization (DPO) pipelines with closed-form policy rewards.
- Scaled test-time reasoning capacity across long-horizon outputs (>20k tokens), increasing complex problem-solving accuracy by 18% while preventing policy entropy collapse by implementing Group Relative Policy Optimization (GRPO) and dynamic entropy-regularized output chunk masking.
- Optimized high-throughput model inference across multi-GPU setups, reducing peak VRAM consumption by 40% and eliminating context boundary overflow across 100+ evaluation pipelines by engineering custom 2D tensor auto-scaling, dynamic context pruning, and automated execution verification harnesses.

MBZUAI

Summer Research Intern

May 2025 - July 2025

Abu Dhabi, UAE

- Designed and deployed a semi-supervised deep clustering framework for 3D cryo-electron tomograms, integrating a YOPO-based feature extractor with a vectorized GMM initialized via k-means and enhanced with label constraints to enable structure discovery from sparse annotations.
- Engineered a robust PyTorch pipeline with modular support for feature extraction, latent space modeling, and dynamic label refinement, improving clustering fidelity on low-SNR volumetric datasets and outperforming baseline unsupervised methods.
- Conducted extensive evaluation of learned representations across multiple cryo-ET benchmarks, demonstrating improved compactness and class separability in feature space, and laying groundwork for downstream subtomogram classification and anomaly detection.

Carnegie Mellon University

Research Assistant at Xulab

October 2024 - February 2025

Pittsburgh, USA

- Under the guidance of Dr. Min Xu at Xulab(<https://xulabs.github.io/>), I am currently immersed in a research role focused on advancing computer vision and machine learning techniques for Cryo ET and Cryo EM images.
- Built complete processing pipelines for Cryo-EM and Cryo-ET volume reconstruction, utilizing tools like EMAN2, CryoSPARC, and Warp for subtomogram averaging, 3D classification, and multi-scale alignment of protein complexes.
- Automated high-throughput analysis across hundreds of tomograms, improving lab efficiency in structure discovery while ensuring reproducibility in ultrastructural reconstruction of macromolecular assemblies.

IRIS,NITK

Machine Learning Developer

September 2024 - Present

Mangalore, India

- Worked with backend and UI teams to deploy ML modules into IRIS—a production-grade ERP platform serving over 7,000+ students and staff.
- Developed a robust profanity detection and sentiment analysis system to filter harmful feedback submitted by users across the university ERP platform, leveraging pretrained transformer-based models and custom rule-based augmentations.
- Used Ollama's LLaMA 7B model to build a summarization tool for over 10+ years of placement data, predicting key skills and qualifications linked to successful company-specific outcomes.

Projects

Movie-Recommender-Systems

Deep Learning

May 2024

Mangalore, India

Content-based recommender systems tailor recommendations by analyzing item attributes and aligning them with user preferences. Collaborative filtering suggests items based on user-item interactions, using either similar users (user-based) or similar items (item-based) as reference points.

Vision Kinect

May 2024

Computer Vision

Mangalore,India

Vision Kinect is a state of the art Machine Learning project, which applies the concepts of Computer Vision, Neural Networks and YOLOv5,(an object detection model) to recreate a hands free version of Tetris, but making it detect different hand-gestures representing different in-game actions. All of this is wrapped up with a web application.

Irish Poem Generator

June 2024

Natural Language Processing

Mangalore,India

The Irish Poem Generator is an NLP project that generates poetic text by using previous words from a seed input to predict new words. Using an LSTM layer to capture long-term dependencies, along with dense and embedding layers for word representation and prediction. Tokenizers are used to preprocess the text, converting words into numerical tokens.

Lunar Lander

June 2024

Reinforcement Learning

Mangalore,India

In this project, we train an agent to successfully land a lunar module on a designated landing pad using reinforcement learning. The goal is to achieve a soft landing, minimizing impact velocity and landing within the designated area. The agent receives rewards based on its actions, such as maintaining proper orientation, landing with gentleness, and avoiding crashing.

Certifications

- Supervised Machine Learning – Stanford University (Coursera)**
- Advanced Learning Algorithms – Stanford University (Coursera)**
- Unsupervised Learning, Recommenders and RL – Stanford University (Coursera)**

Technical skills

Programming Languages	Python, C++, C, Java, JavaScript
Frameworks	PyTorch, TensorFlow, Keras, OpenCV, YOLOv5, HuggingFace Transformers
Tools and Libraries	Git, GitHub, FastAPI, LaTeX, MLflow, Docker
Domains of Research	Computer Vision, Natural Language Processing, Reinforcement Learning, Active Learning, Gen AI, Computational Biology

Language proficiencies

English	Full Proficiency
Hindi	Elementary Proficiency
Japanese	Elementary Proficiency

Extracurricular Activities

WebClub NITK(WEC)	Conducted seminars on CNNs and RNNs as part of Intel SIG. Participated in monthly research paper discussions and hands-on workshops.
DebSoc NITK	Represented NITK in national-level British and Asian Parliamentary debates and Model United Nations conferences.