Prahlad Anand

+1 408 359 7678 | prahladanand01@gmail.com | LinkedIn | GitHub

EDUCATION

Johns Hopkins University, Baltimore, MD **2024 – 2026 (Expected)**
Master of Science in Engineering, Computer Science GPA: 3.93/4.0
Vellore Institute of Technology, Vellore, India **2020 – 2024**
Bachelor of Technology, Computer Science and Engineering GPA: 8.86/10.0

EXPERIENCE

Image Analysis and Communications Lab, Johns Hopkins University
Graduate Research Assistant, *Advised by Dr. Jerry Prince and Dr. Blake Dewey* **Aug 2024 – Present**
- Used **self-supervised multimodal learning** to **disentangle** anatomy and contrast representations from clinical brain MRIs and their corresponding acquisition parameters. Improved on SOTA DICOM tag prediction accuracy by 15%.
- Adapted **Grad-CAM** to a CVAE-transformer network to analyze failure cases for brain MRI harmonization.
- Designed a **gradient-based adversarial attack** strategy using feature perturbation to improve test-time performance.

Artificial Intelligence and Robotics Lab, Indian Institute of Science, Bangalore, India
Research Assistant, *Advised by Dr. Suresh Sundaram* **Jan 2024 – Jun 2024**
- Trained a **semi-supervised RGB-to-IR image translation** network, improving autonomous vehicle perception by 8.6% using synthetic data as augmentation over baseline translation networks with the same strategy.
- Achieved a 4.2% performance improvement for semantic segmentation (DeepLabv3+) and object detection (YOLO, RCNN) models with **zero-shot** synthesis for simulation environments **(real-to-sim domain adaptation)**.

Research Intern **May 2023 – Jul 2023**
- Enhanced the accuracy of **object detection and semantic segmentation** networks by 7% on RGB-IR datasets through multimodal fusion and **data augmentation using generative image translation** models (GANs, diffusion).
- Improved model accuracy by 10.2% over SOTA using supervised IR image synthesis and super-resolution.

Research Intern **May 2022 – Jul 2022**
- Experimented with complex-valued neural networks, loss functions, architectures, augmentation and hyperparameter tuning to boost multi-modal aerial image classification by 7.5% under severe class imbalance.

PUBLICATIONS

- Sikdar, A., Saadiyean, Q., **Anand, P.**, Sundaram, S. (2024). *SSL-RGB2IR: Semi-supervised RGB-to-IR Image-to-Image Translation for Enhancing Vision Task Training in Semantic Segmentation and Object Detection*. **IROS**.
- **Anand, P.**, Saadiyean, Q., Sikdar, A., N., Nalini, Sundaram, S. (2024). *Supervised Image Translation from Visible to Infrared Domain for Object Detection.* **arXiv.**
- **Anand, P.,** Carass, A., Mowry, E. M., Newsome, S. D., Prince, J. L., Dewey, B. E. (2026). *DICOM-CLIP: Zero-Shot Acquisition Parameter Retrieval from Unprocessed DICOM Files.* **ISBI**.

SKILLS

Programming - Python, Java, C/C++, JavaScript, SQL, Bash, PHP, HTML/CSS
Technologies - PyTorch, TensorFlow, Jupyter Notebooks, NumPy, pandas, scikit-learn, HuggingFace, Unix/Linux, Git
Coursework - Medical Image Analysis, Medical Imaging Systems, Computer Vision, NLP, RL, Parallel Computing

PROJECTS

Neuroimage Registration and Synthesis
- Skull-stripped brain MRIs using a U-Net. Performed deformable registration using the ANTs framework, and extended a CVAE-transformer harmonization network to modality fusion and cross-modal synthesis.

Product Recommendation System - Link
- Developed a recommendation system pipeline to web crawl Amazon and create datasets in real-time using Selenium. Suggested related products based on similarity in learned representations of product images and text metadata.

Video Sign Language Recognition
- Achieved +6.7% accuracy and 2.5× faster inference over VideoMAE, a SOTA network using fine-tuned lightweight image classification networks (ResNet50/ViT) on video frames with time-averaged predictions.

Question Answering Model for Passage Based Reading Comprehension
- Modeled a BERT LLM using traditional NLP techniques including NER, POS tagging, TF-IDF and Jaccard index on benchmark datasets SQuAD, QAsper and HotpotQA.

Dots: Geospatial Data Visualization for Businesses - Link
- Built Dots, a geospatial business intelligence tool integrating LLM-powered RAG with DuckDB/Postgres to let users visualize and query data as map-based knowledge graphs. Placed top-10 from 310 participants at HopHacks 2025.