

EDUCATION

- University College London

MSc in Computer Graphics, Vision, and Imaging

London, UK
2024 - 2025
- Birla Institute of Technology and Science, Pilani

BE in Computer Science, Minor in Data Science

Goa, India
2020 - 2024

SKILLS

- Languages

Python, C, C++, CUDA, Bash Scripting, SQL
- Libraries

PyTorch, TensorFlow, Keras, OpenCV, Diffusers, Transformers, Triton Inference Server, TensorRT, ONNX, FastAPI, NumPy, Matplotlib, Pandas, Seaborn
- Platforms

AWS (EC2, S3), GCP, Linux
- Computer Vision Tasks

Segmentation, Image Generation, Object Detection/Tracking, 3D Reconstruction, Multi-View Stereo, Camera Pose Estimation, Super-Resolution, Transformers, Representation Learning

RELEVANT WORK EXPERIENCE

- AI Research Engineer - NeuroDiscovery AI

Gurugram, India

Jan 2026 - Present

◦ Working full stack to build an MRI analysis assistant to solve complex radiologist queries through chained tool calling.

◦ Created an auto-scaling GPU backend for segmentation model inference in production- deployed on EKS leading to improved stability at lower costs.

◦ Exploring SSL methods like DINO/JEPA to train a 3D vision backbone to enable full global (between slices) attention for downstream tasks like high resolution segmentation, classification, VLMs etc.
- AI Research Intern - Jivi AI

Gurugram, India

May 2024 - Jul 2024

◦ Improved medical image representations in CLIP-like latent spaces leading to an enhancement in Vision Language Model (VLM) responses for modalities such as chest x-rays, SLO scans (eyes), and external images of skin lesions/blemishes, boosting classification accuracy by up to 30% over previous best models on some datasets

◦ Researched mixture of experts methods to enable a single VLM to provide accurate answers across various domains (ophthalmology, dermatology, etc.) while reducing catastrophic forgetting, amounting to 40% lower inference costs
- AI Research Intern - Spyne

Gurugram, India

Jul 2023 - Dec 2023

◦ Developed an end-to-end processing pipeline to improve images of food items, working on tasks such as classification, salient object segmentation, and outpainting using diffusion models

◦ Refined data collection methodologies by integrating diffusion models to synthesize realistic datasets for challenging scenarios with limited available data- growing existing datasets 10-fold

◦ Traced and deployed these models to a Triton inference server for efficient batched inference, along with coding their respective APIs using FastAPI to be used in production

◦ Took the initiative of deploying CVAT with the Segment Anything Model (SAM) backend to help the data annotation team process samples twice as fast

PROJECTS

- Master's Thesis - Smart Geometry Processing Group @ UCL

Supervised by Prof. Niloy Mitra — Code

Mar 2025 - Sep 2025

◦ Worked on an efficient, feed-forward (optimization-free) approach to multi-view stereo and camera pose estimation.

◦ Employed a modified DUST3R model to perform 3D reconstruction from sequences of key frame depth estimates and 2D point tracks, approaching state-of-the-art results with ~1% of the training data compared to other methods.
- Mangrove Forest Project

Sanctioned by the Forest Department of Goa

Sep 2022 - May 2023

◦ Created a dataset of over 3,000 images of 16 species of mangrove trees in state to assist forest department officials in mapping mangrove forests, aiding in their conservation efforts

◦ Trained an ensemble of CNNs for classification, boosting the accuracy to over 90% while reducing the model's size by 80% from the previous iteration - allowing the model to be used on edge devices