

Education

- **University of Michigan**, Ann Arbor (ongoing) Aug. 2025 – May. 2027
MS (Computer Science) - **GPA: 3.98**
Research Assistant on 3D Generative AI ([Dr. Jeong Joon Park](#))
- **SRM Institute of Science and Technology**, Chennai, Tamil Nadu Aug. 2021 – Aug. 2025
B.Tech (Computer Science)
CGPA: 9.47/10

Publications

- *Pagoda: Roofline Characterization of Energy and Time for DNN Inference and Training on Edge Accelerators*, ICPP (accepted). ([link](#))
- *Fulcrum: Power and Latency-aware Optimization of Concurrent DNN Training and Inference on Edge*, ACM TECS (under review). ([link](#))
- *ViDAS - Vision-based Danger Assessment and Scoring*, ICVGIP 2024. ([link](#))
- *ECHO - Environmental Sound Classification with Hierarchical Ontology-Guided Semi-Supervised Learning*, IEEE CONECCT 2024. ([link](#))
- *Managing Congregations of People by Predicting Likelihood of a Person Being Infected by a Contagious Disease like the COVID Virus*, IEEE CCEM 2020 - **Best Short Research Paper Award**. ([link](#))

Work Experience

1. **Research Assistant, University of Michigan, PI - [Dr. Jeong Joon Park](#)** Sept. 2025 – Present
 - Working on a **flow** model that generates voxel-based **surface fields** inspired by **TRELLIS** to extract high-quality and **topologically robust** meshes.
 - Working on **hyperspectral** foundation models to classify images of chemical solutions and concentrations in spillages.
2. **Research Intern, Stanford University, PI - [Dr. Teddy Akiki](#)** Jun. 2024 – Jun. 2025
 - Developed autoregressive and diffusion models (**Chronos**, **TimeLLM**) for time-series forecasting to uncover latent “language-like” patterns in **BOLD fMRI** signals.
 - Designed and trained a **VQ-VAE + GPT2/LLaMA2** pipeline to discretize continuous fMRI signals into latent tokens and predict future sequences.
3. **SDE Intern, Kiwi** Apr. 2025 – May. 2025
 - Constructed a **LangGraph multi-agent framework** handling **100+** customer service **interactions per day** across varied query types like Transactions, Rewards, and Card applications.
 - Framework included **15+** **tools** and Freshdesk and Mixpanel integrations delivering accurate results in under a minute.
4. **Research Intern, Indian Institute of Science, PI - [Dr. Yogesh Simmhan](#)** Aug. 2024 – Feb. 2025
 - Proposed and improved baseline neural networks through a **penalizable MAPE loss function** in optimizing concurrent training and inference on edge accelerators from MAPE scores of 10% to 2%.
 - Built a **novel Active Learning technique** that samples the best power modes on Jetson devices for DNN workloads with losses close to optimal (0%).
5. **Research Intern, Samsung R&D, PI - [Mr. Shwetank Choudhary](#)** Sept. 2023 – Jul. 2024
 - Devised a **novel** two-stage pipeline that trains a CNN with semi-supervised learning(ResNet50, EfficientNet) on the label ontology of environmental sound classes as a **pretext** task.
 - Produced an accuracy **increase** of 1-8% compared to baseline models, achieving 97.5% **state-of-the-art** accuracy in the ESC-10 benchmark dataset.
6. **SDE Intern, Upthrust** Jun. 2023 – Aug. 2023
 - Integrated **PageSpeed API** and **Screaming Frog API** to produce and fix website issues with **Langchain** and **LLMs**.
 - Linked the **Meta API** with **GPT** to create a chatbot providing **recommendations to optimize ads**.
7. **SDE Intern, AarogyaAI** Jun. 2020 – Sept. 2021
 - Deployed a dashboard to display datasets and model results stored in AWS S3 bucket, showcasing what drugs are resistant to the given variant of Tuberculosis bacteria.

Positions of Responsibility

Odyssey Lab ([link](#))

Jan. 2024 – Jan. 2025

- Started a lab to help students across the best universities in India to work on research.
 - Built a network of mentors from IISc, IIIT-H, Amazon, IBM, and Samsung and created 3 teams to work on Multimodal and LLMs reasoning projects.
 - Published a paper to ICVGIP 2024, India's top Computer vision Conference.
-

Skills

- **Languages:** Python, C++, SQL
 - **ML / Deep Learning:** PyTorch, TensorFlow, Hugging Face, Scikit-learn, LangGraph, Semi-supervised learning, Self-Supervised learning, NumPy, Pandas, Weights and Biases, Active Learning
 - **Computer Vision / Generative AI:** Transformers, Diffusion Models, Flow Matching, Vision-Language Models, Large Language Models, Agentic AI, 3D Generative Models, Hyperspectral Imaging, Research
 - **3D Vision / Graphics:** 3D Mesh Generation, Voxel Representations, Surface Reconstruction, Geometric Deep Learning
 - **HPC / Systems:** SLURM, GPU Computing, Distributed Training, Linux, HPC Clusters, Edge Accelerators, Jetson
 - **Software:** Data Structures, Object-Oriented Programming, REST APIs, API Integration, Debugging, Unit Testing, Backend Development
 - **Tools / Platforms:** Git, GitHub, AWS, Jupyter, Conda
-

Relevant Coursework

- | | |
|---|--|
| • EECS 542 - Advanced Computer Vision | • CSE 598 - Action and Perception |
| • CSE 595 - Natural Language Processing | • CSE 598 - Graphics and GenModels |
| • CSE 585 - Advanced Scalable Systems in Agentic AI | • CSE 598 - Building Small Language Models |