

Education

- **SRM Institute of Science and Technology**, Chennai, Tamil Nadu Aug. 2021 – Aug. 2025
Bachelor of Technology in Computer Science and Engineering
CGPA: 9.45/10

Publications

- *Fulcrum: Power and Latency-aware Optimization of Concurrent DNN Training and Inference on Edge*: Prashanthi S. K., Saisamarth Taluri, **Pranav Gupta**, Amartya Saikia, Lakshya Karwa, Kedar Dhule, Yogesh Simmhan. Submitted to ACM SenSys 2025 (CORE A*).
- *VARS: Vision-based Assessment of Risk in Security Systems*: **Pranav Gupta**, Pratham Gohil, Sridhar S. Submitted to IEEE Access Journal.
- *ViDAS - Vision-based Danger Assessment and Scoring*: **Pranav Gupta**, Advith Krishnan, Naman Nanda, Ananth Eswar, Deeksha Agarwal, Pratham Gohil, Pratyush Goel. Accepted at ICVGIP 2024 ([link to preprint](#)).
- *ECHO - Environmental Sound Classification with Hierarchical Ontology-Guided Semi-Supervised Learning*: **Pranav Gupta**, Raunak Sharma, Rashmi Kumari, Sri Krishna Aditya, Shwetank Choudhary, Sumit Kumar, Kanchana M, R Thilagavathy. IEEE CONECCT 2024 ([link to paper](#)).
- *ISAAApp - Image Based Attendance Application*: Aritra Dutta, G Suseela, G, Niranjana, Pushpita Boral, **Pranav Gupta**, Subha Bal Pal. International Conference on Advances in Artificial Intelligence and Machine Learning in Big Data Processing, 2023 ([link to paper](#)).
- *Managing Congregations of People by Predicting Likelihood of a Person Being Infected by a Contagious Disease like the COVID Virus*: **Pranav Gupta**, Manish Gupta. IEEE Cloud Computing in Emerging Markets, 2020 - **Best Short Research Paper Award** ([link to paper](#)).

Work Experience

- Research Intern, Stanford University, PI - Dr. Teddy Akiki** Jun. 2024 – Present
- Implemented transformer-based **time-series forecasting** models like **Chronos** and **TimeLLM** on data from fMRI scans to find “language” patterns.
 - Proposed a decoder-only transformer architecture like TimesFM to forecast future brain states and the preliminary finetuning results with R^2 scores of 0.4+ are promising.
 - Proposed multivariate forecasting as an inpainting problem using GANs and diffusion models.

- Research Intern, Indian Institute of Science, PI - Dr. Yogesh Simmhan** Aug. 2024 – Present
- 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%).

- Research Intern, Samsung R&D, PI - Mr. Shwetank Choudhary** Sept. 2023 – Jul. 2024
- **Lead** researcher in a semi-supervised learning project focusing on environmental sound classification.
 - Devised a **novel** two-phase pipeline that trains a CNN (ResNet50, EfficientNet) on the label ontology of 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.

- AI 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**.

- Data Science Intern, AarogyaAI** Jun. 2020 – Sept. 2021
- Developed 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.
 - Automated results analysis and plotting of ML models in drug prediction.
 - Automated and detected 100+ anomalies in data of tuberculosis mutations.

Positions of Responsibility

- Odyssey Lab** Jan. 2024 – Present
- 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.

- Next Tech Lab** Nov. 2021 – Present
- Appointed as the Head AI Researcher and led 100+ undergraduates through mentorship, and research collaborations.

Open-Source and Community Contributions

Hackthons and Team Collaborations

- Quantathon 1.0 - 1st Runner's Up - Created a sentiment analysis model using qiskit's quantum computing library
- MesoHacks 2022 - Special Mention Prize - Created a Surveillance Anomaly detector
- Smart India Hackathon 2023 - Selected for National - Developed a de-hazing algorithm and smoke remover model for CCTV footages and drone surveillance.
- Participated in MLH Hack-the-mountains 3.0 and a Datathon hosted by IIT Kanpur on document classification

Technical Articles on Medium.com

- Wrote an article on Face Recognition and embeddings with a novel pipeline for creating an attendance app that received **6.3K views** and **3.5K reads**. ([link to article](#))
- Written articles on the AI Wordle Solver and Genetic Handwritten Digits project.

Github Contributions

- Developed a repository for a template code for the usage of the OpenAI Assistants API which received **86 stars** and **19 forks**. ([link to project](#))
 - Completed HacktoberFest 2022 by getting pull requests merged in 4+ open-source repositories.
-

Personal Projects

- **AI Wordle Solver:** Created a Wordle Solver by producing the next best word to play using AI and algorithms given a screenshot of a partially filled Wordle. Computed with OpenCV and TensorFlow to produce two models to recognize alphabets and classify 3 different colors - green, yellow, and grey ([link to project](#)).
 - **Splitwise-GPT-Vision:** Combined GPT-Vision and PyTesseract to perform OCR on a bill and convert the bill into a Pandas dataframe using function calling. Integrated the Splitwise API on a Streamlit dashboard to select people paying for specific items and adds personalized splits directly to the Splitwise App ([link to project](#)).
 - **SimCLR-UrbanSound8K:** Implemented the SimCLR contrastive learning model from scratch on the Urban-Sound8K dataset for audio classification using PyTorch. Applied random augmentations like random cropping and time shifting as instructed in the paper. Produced an accuracy of 81% on mel-spectrograms of the audios in the dataset ([link to project](#)).
 - **MusicLM Generation:** Trained a text2music AudioLDM model on MusicCaps Dataset using the HuggingFace diffusers library. Finetuned the Stable Diffusion model on the melspectrogram images of the dataset using the librosa library ([link to project](#)).
 - **Genetic Handwritten Digits:** Optimized the training of a handwritten digit's model using Evolutionary genetic algorithms from scratch on the MNIST dataset. Showed proof of concept as the model evolved from an accuracy of 92% to 96% ([link to project](#)).
-

Relevant Coursework

- | | | |
|---|--|---------------------------------------|
| • Linear Algebra | • Analog and Digital Electronics* | • Artificial Intelligence* |
| • Advanced Calculus | • Design and Analysis of Algorithms* | • Digital Image Processing |
| • Probability and Queueing Theory | • Operating Systems* | • Neuro-Fuzzy and Genetic Programming |
| • Discrete Mathematics | • Software Engineering and Project Management* | • Computer Networks* |
| • Algorithms* | • Advanced Programming* | • Data Mining |
| • Object-Oriented Programming* | • Database Systems* | • Information Storage and Management |
| • Computer Organization and Architecture* | • Compilers* | |

(* marked courses include laboratory component)

Technical Skills

- **Programming Languages:** Python, C, C++, Java, JavaScript
- **Web Technologies:** HTML, CSS, Node.js
- **Databases:** MySQL
- **Machine Learning:** TensorFlow, PyTorch, scikit-learn
- **Deep Learning:** Hugging Face, Keras, OpenCV
- **Cloud Platforms:** AWS, Azure, Google Cloud
- **Frameworks:** LangChain, Flask, Django, Streamlit, Gradio
- **Data Visualization:** Matplotlib, Seaborn