Visweswar Sirish Parupudi

+1 475-348-4238 | sirish.p@nyu.edu | LinkedIn | GitHub | Website

Education

New York University, Tandon School of Engineering — Brooklyn, NY
Computer Science, M.S. — *May 2026*

Birla Institute of Technology and Science Pilani — Hyd, India
Electronics and Communication, B.E. | *Minor in Computing and Intelligence* — *May 2024*

Technical Skills

Programming Languages: Python, C, C++, C#, SQL, JavaScript, MATLAB, Assembly
ML / AI: PyTorch, TensorFlow, scikit-learn, OpenCV, pandas, NumPy, Matplotlib; NLP: NLTK, spaCy; RL: Gymnasium
Frameworks / Web / Tools: Django, React, HTML, CSS, Unity, PySpark, Hadoop, Docker, Kubernetes, Arduino

Experience

Researcher | *EEG–fMRI Cross Modal Translation, Foundation Models* — June. 2025 – Present
Neural Translation and Foundation Models — *Neuroinformatics Lab, NYU, NY*

- Building multimodal neural foundation models (CBraMod, BrainLM) over 1.1M+ EEG windows (TUH) and large-scale fMRI datasets (UK Biobank, HCP, NATVIEW) to **learn shared latent spaces** across electrophysiology and hemodynamics.
- Implemented cross-modal translators with asymmetric positional encodings, criss-cross attention, and staged Huber-loss training; achieved ROI-level correlations up to **0.83** on NATVIEW held-out subjects for visual decoding tasks.

Research Assistant (Undergraduate Thesis) | *LangChain, LLMs, PyTorch* — June. 2023 – Dec. 2023
Text Augmentation for NER and Language Modeling — *Nanyang Technological University, Singapore*

- Improved Named Entity Recognition (NER) performance on WNUT17 dataset of tweets by mitigating challenges from noisy data. Utilized entity placeholder replacements and PEGASUS-based paraphrasing to clean noisy text, resulting in more accurate NER predictions.
- Employed PEGASUS to produce multiple paraphrases, selected based on named entity retention, fluency and BERTscore. Boosted performance up to F1 score 0.558, with a 43% increase from the baseline, succeeding in handling noisy input more effectively.
- Leveraged **LLaMA 2** for advanced paraphrase generation through zero-shot and few-shot prompt engineering techniques, with a focus on preserving entities and improving fluency. Achieved a peak F1 score of 0.581 and an **85% increase** from the baseline with 6% entity loss.

Software Engineer Intern | *Python, SQL, Web3* — May. 2022 – July. 2022
Ethereum Gas Price Prediction — *DLT Labs (KNNX), Hyderabad, IN*

- Devised a prediction model to forecast Ethereum gas prices using time series data forecasting techniques, with a focus on ARIMA methodologies. Successfully extracted and processed gas price data directly from Ethereum Blockchain APIs.
- Delivered accurate daily and minute-by-minute gas price forecasts through rigorous hyperparameter tuning and model diagnostics, including ACF/PACF analysis, to **minimize MAE by 27%**, contributing to enhanced decision-making for participants.

Publications

A Smartphone-Enabled Deep Learning Approach for Myocardial Infarction Detection Using ECG Traces for IoT-Based Healthcare Applications, in **IEEE Sensors Letters**, vol. 7, no. 11, pp. 1-4 *Python, Tensorflow, TFlite, Android Studio* | Nov. 2023

- The model uses a majority voting ensemble approach with five branches, each containing a quantized transfer learning architecture (FP16 and INT8 formats) to achieve high performance while maintaining efficiency on mobile devices for medical health imaging.
- The overall accuracy of the developed ensemble DL-based model has been obtained as **99.34%** using the fivefold CV scheme. The proposed smartphone-app based healthcare system is portable and can be utilized for **real-time detection of MI** and other heart abnormalities.

Selected Projects

PinToBeans – Pinterest-Like Social Platform | *Django, PostgreSQL, HTML/CSS, Python* — Jan. 2025 – Apr. 2025

- Built a full-stack web app with Django + PostgreSQL supporting pinning/repinning, search, friend-gated comments, likes, and followstreams; designed a BCNF schema with 12+ tables and triggers enforcing rules (self-repin prevention, redirected likes)
- Developed backend views and REST-style endpoints for auth, board/content management, and keyword search; implemented responsive templates with HTML/CSS, optimized ORM queries for $<$200ms responses, and integrated logging for reliability/debugging.

AD-FTD-CN EEG Classification using Multi-View Transformers | *Python, PyTorch, Transformers, EEG* — Jan. 2025 – May. 2025

- Designed a deep learning pipeline using **Multi-View Transformer (MVT)** with **Graph Attention**, cross-view fusion, and multiscale EEG embeddings integrating time, spectral, and spatial features for classifying neurological conditions using 19-channel EEG data.
- Achieved **98.55% within-subject** and **64.95% cross-subject** accuracy, outperforming SVM, MLP, and ADformer. Generated **scalp heatmaps** and **attention maps** to visualize lobe-specific frequency patterns and enhance interpretability.

Det-IGEN: Detectable Watermarking in Image Generation | *Python, PyTorch, GenAI, Docker* — Sept. 2024 – Dec. 2024

- Implemented invisible tree ring watermarking system into the generation process of Diffusion Models by modifying the initial noise vector in the Fourier domain, incorporating an invisible pattern (key) that remains robust against various image transformations.
- Demonstrated negligible impact on image quality, with minimal changes compared to non-watermarked images, enabling **detection of AI generated content** (**AUC and TPR@1%FPR of 1.000**) and maintained high performance under adversarial conditions.