

EDUCATION

New York University(Courant Institute of Mathematical Sciences)

New York, NY

*Master of Science in Computer Science; CGPA: 3.93/4.00**Sept 2022 - May 2024***Coursework:** *Deep Learning, Big Data & ML Systems, Cloud and Machine Learning, NLP, Large Language and Vision Models.***Current Work:** *1. Graduate Teaching Assistant for Intro to Computer Vision Course (Taught by Prof. Jean Ponce)**2. Graduate Teaching Assistant for Deep Learning Course (Taught by Prof. Alfredo Canziani and Prof. Yann Lecun)***National Institute of Technology, Tiruchirappalli**

Trichy, India

*Bachelor of Technology in Computer Science and Engineering**July 2016 - June 2020**Minor Degree in Management; CGPA: 8.6/10***Coursework:** *Natural Language Processing, ML, Probability Theory, Statistics, Linear Algebra, Optimization Techniques*

RESEARCH AND WORK EXPERIENCE

Data Science Intern, Product & Insights | Audible, Amazon | Newark, NJ**Sept 2023 - Present**

- Designed a big data ETL pipeline using **Spark SQL** and **Scala** to extract a unified analytics dataset.
- Developed an end-to-end **MXNet** based ML model pipeline with SageMaker Step Functions to predict customer propensity to purchase Audible original content. Operationalized model training, evaluation, and deployment on AWS cloud infrastructure.
- Built a model to rank buying offers on Amazon PDP for individual customers, optimizing conversion rates and incremental lifetime value. Executed end-to-end model development, including training and inference using XGBoost on AWS SageMaker.

Deep Learning Research Affiliate, Neural Data Science Lab | Georgia Tech | Atlanta, GA July 2023 - Oct 2023

- Worked on a domain adaptation method for complex time series that can handle both feature and label shifts.
- Investigated the impact of methods like Optimal Transport, CORAL, and Sinkhorn loss on aligning time and frequency features, to enhance the generalization performance within an unknown target domain.

Graduate Research Assistant, Prof. Vasant Dhar | New York University | NYC, NY**May 2023 - Sept 2023**

- Conducted a series of experiments to predict future returns using OHLC charts derived from real-time stock data, employing deep learning models including lightweight CNN, Vision Transformer and Auto-Encoders. Established a baseline F1-Score of **51%** with raw time-series features.

Software Engineer (Machine Learning) | Qualcomm | Bangalore, India**June 2020 - Aug 2022**

- Designed and optimized several positioning algorithms to deliver solutions to automotive and mobile-based customers.
- Implemented denoising auto-encoders to improve signal acquisition sensitivity and super-resolution autoencoders to improve multipath performance, reducing power consumption by **1.5dB** and improving CEP95 by **20%**.
- Prototyped a proximity detection model that preserves user privacy using two GNSS satellites + Angle of Arrival (AoA) based DBSCAN clustering technique.
- Improved position accuracy by enhancing weak signal context detection, fine-tuning initialization parameters in Kalman filtering-based position algorithm, and applying code enhancements in E-911 sessions leading to a **30%** improvement in an open-sky environment and submitted this work for a **US patent**.

Software Engineer Intern | Qualcomm | Bangalore, India**May 2019 - July 2019**

- Performed a thorough trade-off study to mitigate multipath issue seen in GNSS by comparing various robust estimators such as RANSAC, L-medians, and M-estimator. L-medians emerged as the most effective, reducing error to **30m** and tolerating up to $N/2$ outliers.

SELECTED PROJECTS

Hate-LLaMA: An Instruction-tuned Audio-Visual Language Model for Hate Content Detection - [[Code](#)][[Paper](#)]

- Fine-tuned Video-LLaMA - that effectively processes video and audio information multimodally - on a hate video dataset over 4 RTX8000 GPUs, achieving an F1-score of **73%** on the unseen test dataset.
- Released a **benchmark** dataset comprising **300** labeled hate/non-hate videos & deployed a demo for Hate-LLaMA(Multimodal LLM) on HuggingFace Spaces using **Gradio**.

Multi-Image Fusion and Semantic Segmentation for Power Substation Delineation | *Pytorch, TorchGeo, HPC*

- Collaborating with **Transition Zero**, a leading clean energy technology company to develop a system leveraging deep learning models to fuse multiple revisits of low-resolution **Sentinel-2 L1C imagery** and detect power substation from the super-resolved image.
- Fine-tuned a U-Net model with a ResNet50 encoder using pre-trained weights from PyTorch's **TorchGeo** library, achieving an Intersection over Union (IOU) score of **43.8%**.

Integrated Content Management Platform - New York Public Library (NYPL) | *Pytorch, HuggingFace, OpenCV*

- Deployed **Vision Transformer** based classifier for asset type detection and enabled automated metadata tagging using **Flan-T5** for enhanced content search accessibility of NYPL's collection items.
- Integrated **crowd-sourcing** techniques and enabled periodic model re-training, leading to a significant improvement in user experience and a more efficient content management process.

Video Frame Prediction and Semantic Segmentation with Self-Supervised Learning - [[Code](#)] | *Pytorch, Wandb, HPC*

- Implemented a fully CNN architecture - **SimVP** - for video frame prediction and trained a U-Net for Semantic Segmentation on a diverse synthetic dataset consisting of video clips of 3D moving objects, achieving a Jaccard Index of **0.251** with minimal fine-tuning.

Semantic Role Labeling on NomBank Dataset - [[Code](#)] | *Pytorch*

- Developed a GCN + Bi-LSTM model to accurately identify argument-predicate relationships within sentences by utilizing syntactic and contextual features, including dependency relations, POS tags, word and lemma embeddings that yielded a F1-score of **0.83**.

PROGRAMMING SKILLS

Python, C++, C, Matlab, Java, R, Scala, Spark SQL, Pytorch, Tensorflow, Keras, Docker, Kubernetes, AWS, GCP, PySpark, Cuda, Tableau, Scipy, Scikit-learn, Pandas, Django, Git, MySQL, PostgreSQL, MongoDB, Oracle DB