

Saurav Chennuri

Actively looking for MLE/SDE roles

Boston, MA | +1(484)320-9175 | saurav.varaprasad@outlook.com | [Homepage](#)

Education

Master of Science, **Artificial Intelligence**

Boston University, Boston, MA, (2021 - 2023); **(3.9 / 4.0)**

Bachelor of Technology, **Engineering Sciences**

Indian Institute of Technology (IIT), Hyderabad, India, (2016 - 2020)

Skills

Programing Languages: Python, C++, SQL, Java, Shell Script, C#.

Databases: MySQL, Snowflake, MongoDB, KQL, Elastic Search, GraphQL.

Libraries and Frameworks: Pytorch, TensorFlow, Pytorch-Lightning, OpenCV, XGBoost, Numpy, Pandas, Springboot, Docker, TensorBoard, Wandb.

Work Experience

Fractal Analytics

Software Engineer (**Data Analytics**)

Oct 2020 - Aug 2021

- Worked with a MLOps team focused on providing better data visualization and user experience related to financial data from a multinational bank.
- Developed a web app to provide better dynamic data visualization based on user query improving **query-visualization relevance** by **15%** based on **Mean Reciprocal Ranking** among the top-7 recommended visualizations according to query
- Implemented a **chatbot** that converts natural language to structured data query, providing better abstraction for complex searches, and improving intent understanding by **29%** as compared to traditional search queries.
- Implemented auto suggestion component to recommend users with search data fields based on user query, enhancing user search experience by **40%** based on **A/B testing**.

Center for Brain Recovery, Boston University

Research Assistant

Jan 2023 - May 2023

- Quantified the severity of Aphasia among the patients using **Machine Learning** approaches on the **multimodal data** of structural MRI scans, functional MRI scans, DTI information, and patient's demographic data.
- Introduced better and reliable training process to **mitigate the data leakage** in traditional prediction strategies allowing removal of overlap between training and testing data, and providing more truthful results
- Our new prediction approach through nested cross validation approaches were on par with the best prediction approaches in the area of about **17.51 RMSE** with the ground truth severity scores.

Takenaka Corporation

Machine Learning Engineer Intern

May 2019 - Jul 2019

- Worked with the Information and Communication Technology (ICT) team, focused on developing new technologies to assist in better planning and designing of architecture across multiple landscapes.
- Developed an ensemble model (**Xgboost, SVMs, Random Forests**) to predict daily workforce requirements for construction projects, achieving an error margin of **10 workers/day**.
- Implemented a classification component that predicts the kind of jobs required in a given type of construction with an accuracy of **81%** with respect to the ground truth of the requirement of the job in the construction.

Academic Research

Contrastive Learning to improve text embeddings for low resource language News Articles: (**Natural Language Processing**)

- Inspired by OpenAI's CLIP, I contrastively finetuned **mT5 transformer(small)** to update Chichewa(Nyanja) language news article embeddings to be closer to its parallel english translated new article embeddings.
- Implemented **Mixup** and **NLPAug** based text data augmentation approaches to address class imbalances.
- Final overall news article classification accuracy reached **61.2%** with enhanced embeddings and a **DNN classifier**, while the accuracy varied between **50-60%** for CountVectorizer and **TFIDF embeddings** with **Support Vector Machines, Random Forests, Naive Bayes** and **XGBoost** classifiers..

Multi-Modal Framework for Personalized Short Video Search and Recommendation:

(**Recommendation Systems** | **Multimodal Learning** | **Video Search**)

- Developed an experimental **multi-modal framework** for short video search and recommendation, combining **DistilBERT** for text analysis, **EfficientNet-B4** for visual processing, and a **Temporal Transformer** for user behavior modeling.
- Demonstrated a **28% improvement** in **NDCG@10** over baseline methods on a dataset of 100,000 short videos with associated user interactions.
- Implemented a cross-attention mechanism for **query-video matching** and a **Temporal Transformer** for modeling user search history. Evaluated using offline metrics, attaining a **24%** improvement in **Mean Reciprocal Rank (MRR)** and a **20%** increase in **Mean Average Precision (MAP)** compared to **collaborative filtering** baselines on a held-out test set.
- Created an integrated video understanding module combining NLP techniques for caption analysis, **YOLOv5** for object detection in key frames, and **VGGish** for audio feature extraction. This holistic approach enhanced content diversity in search results by **18%**, as measured by the Intra-List Distance metric.
- Performed **ablation studies** to quantify the contribution of each component, finding that multimodal integration accounted for **40%** of performance gains, temporal user modeling for **35%**, and advanced video understanding for **25%**.

Manipulating SGD for Data Ordering Attacks in Deep Learning approaches: (**Deep learning** | **Computer Vision** | **Adversarial attacks**)

- Implemented a series of adversarial attacks on **ResNets** and **Vision transformers (Vit-b-16)** based on batch reordering and reshuffling methods in 4 attack policies to either reduce the performance of the model or slow down the training process.
- These **blackbox attack** methods exploit the stochasticity of the gradient descent update rule over the epoch, and make it very hard to observe and mitigate the attack.
- **ResNet-18,50** had a classification performance drop of **3-60%** on CIFAR-10,100 benchmarks, while Vit-b-16 had a performance drop of **3-30%** on SVHN benchmark. Surrogate model used is Lenet-5.

Publications

Fusion Approaches to Predict Post-Stroke Aphasia Severity from

Multimodal Neuroimaging Data. [[Github](#) | [Paper](#)]

Published at **ICCV Workshop Computer Vision for Automated Medical Diagnosis, 2023.**

Feature Analysis and Extraction for Post-Stroke Aphasia Recovery

Prediction [[Paper](#)]

Published in **Medical Imaging Understanding Analysis Conference, 2022.**

Towards Fast Crash-Consistent Cluster Checkpointing [[Github](#) | [Paper](#)]

Published in **IEEE conference on High Performance Extreme Computing, 2022.**

Sloshing Noise Classification in Fuel Tanks of Hybrid Vehicles using Convolution Neural Networks. [[Paper](#)]

Published in the **Journal of Acoustic Society of America, 2021.**