

Sai Varun Reddy Mullangi

☎ 4804428803 | ✉ smullan2@asu.edu | in | 🌐

EDUCATION

Arizona State University

Arizona, AZ

Master of Science in Computer Science; GPA: 4.0/4.0

May 2024

- Relevant Courses: Data Structures and Algorithms, NLP, Advance Operating Systems, Cloud Computing

Birla Institute of Technology and Science (BITS), Pilani - Hyderabad

Hyderabad, India

M.Sc. (Hons.), Mathematics and B.E. (Hons.); GPA: 7.4/10

May 2019

TECHNICAL SKILLS

Programming Languages: Python, Java, C/C++, JavaScript, SQL

Frameworks, Libraries and Tools: PyTorch, TensorFlow, REST, Flask, Redis, MongoDB, GDB, Git, Kubernetes, Docker

Professional Attributes: Problem-solving skills, Effective communication, Leadership ability

WORK EXPERIENCE

Arizona State University

Arizona, AZ

Graduate Research Assistant

Jan 2023 – present

- As a Graduate Research Assistant with Professor Tejaswi Gowda, I'm developing IMUNet++, an efficient deep learning architecture for accurate position estimation from raw IMU measurements. This innovative solution is designed for edge devices and implemented using PyTorch.
- As a Graduate Student Assistant at the College of Health Solutions under Professor Yi Zhou, I'm currently developing an interactive Face Pong game using JavaScript, Three.js and MediaPipe for real-time face detection. Researchers can configure game settings, and all data is stored in MongoDB. I'm also creating a data analysis dashboard with Plotly for pattern identification.

ZS Associates

Bangalore, India

Data Scientist

Jan 2021 – May 2022

- Part of a team that designed a positive-unlabeled machine learning algorithm to identify ovarian cancer patients with missing insurance claims for surgery, biomarkers, and PARP (poly ADP-ribose polymerase inhibitors) which resulted in identification of \$40 million unaddressed opportunities for the client.
- Developed a Seq2Seq LSTM Encoder-Decoder model to identify anomalous ovarian cancer patient journeys and used cross-entropy loss to evaluate the anomalous patient cluster.

Quantiphi Analytics

Mumbai, India

Software Development Engineer

July 2019 – Jan 2021

- Implemented containerization of an existing solution and deployed it on a Google Kubernetes Orchestration Cluster, while also building a highly scalable video processing framework using Pub/Sub, resulting in a 158% improvement in the client's product efficiency
- Part of a team that designed and implemented a binary CNN based AI model for top US medical firm to classify images of a SARS-CoV-2 Antigen test cassette for the COVID-19 virus, achieving exceptional sensitivity and specificity of 0.99 and 0.98 respectively
- Developed and implemented a data retrieval model that fetches information from two separate unstructured databases for a given user query. Utilized Universal Sentence Encoder to create embeddings, weighting column names and values, and integrated DistilBERT to find the exact answer from the top 3 scored columns based on their Cosine Similarity

ACADEMIC PROJECTS

OS Boot Process Optimization and Memory Management Enhancement

Spring 23

- Developed CPU-specific Boot ROM and optimized bootloader for QEMU RISC-V system. Enhanced OS detection logic, significantly improving boot process speed. Conducted extensive testing using diverse test cases and employed GDB for meticulous debugging.
- Implemented on-demand paging and copy-on-write techniques using page faults. Dynamically allocated process pages, optimizing memory usage. Developed a robust page fault handler, enhancing system responsiveness. Implemented efficient page swapping for optimal memory utilization.

Image Recognition service using AWS

Spring 23

- To develop an image recognition service using AWS cloud resources such as EC2 - to run the application, S3 - to store the input image and output results and SQS - for load balancing, autoscaling

- Users will upload their images to a URL. The query images are then buffered in an SQS queue till it reaches the app tier where the pre-trained deep learning model is used to predict the output labels. These labels are then queued back into SQS buffer that leads its way to the web tier and finally to the user

Relationship Extraction tackling Large Input Large Output (LILLO) problem

Fall 2022

- Utilized context-aware summarization to improve the performance of relation extraction on Electronic Health Records (EHR) which are known to be long and pose a challenge for state-of-the-art NLP systems. Utilized the 2010 Relations Dataset for the research

End to end learning of steering commands for self driving

Fall 2022

- Developed a self-driving car simulation using 28k images from cameras mounted on the vehicle, utilizing Udacity's simulator. Trained a CNN model with multi-task learning to predict angle and throttle based on the simulated images.