

ANJALI MUDGAL

+12023912733 ◊ mudgal.anjali.am@gmail.com ◊ github.com/ianjali ◊ linkedin.com/in/anjali-mudgal ◊ ianjali.github.io

OBJECTIVE

Machine Learning Engineer specializing in LLM development and optimization, with expertise in fine-tuning techniques such as LoRA, QLoRA, and PEFT. Skilled in building AI agents and search-oriented systems using frameworks like LangGraph and CrewAI. Experienced in deploying production-ready models and optimizing processes for seamless scalability. Excited to contribute by creating and deploying innovative technologies to enhance delivery experiences and drive customer satisfaction.

RELEVANT WORK EXPERIENCE

Data Scientist

Enquire AI

Mar 2023 - Present

Washington, DC

- **AI Assistant Speed Optimization using LangGraph:** Developed customized AI agents with multithreaded document search, migrating from AWS Bedrock to LangGraph architecture, resulting in 30% faster response time.
- **Predictive Modeling:** Implemented automated LLM model fine-tuning pipeline using cronjobs on EC2 instances.
- **Optimizing Search Speed using Vector Database and Redis:** Accelerated profile search speed by 40% through efficient document retrieval, integrating pgvector on Postgres for BERT embeddings to enhance user experience.
- **Advance SQL and Statistics:** Proposed and implemented a novel approach using Mean Average Precision with Common Table Expressions in SQL, boosting correctly matched profiles from 50% to 70%.

Data Scientist

World Bank

Mar 2024 - May 2024

Washington, DC

- **Interactive Coding Assessment:** Designed assessment tools for World Bank's Data Science skill certification program.

Software Engineer (Machine Learning, Mobile)

Samsung Research & Development

Jan 2019 - July 2022

Noida, India

- **Real-time Speed Optimization:** Reduced sensor-to-application **image frame latency by 100ms** for **image processing pipeline** in phones, through optimization techniques, significantly enhancing operational efficiency.
- **Custom Mobile Net Architecture:** Reduced RAM usage from 500MB to 80MB by developing space-efficient MobileNet and QLoRA for real-time image segmentation, enabling bokeh effect in cameras suitable for low-budget smartphones.
- **Mentor, Research and Development:** Actively guided 5 interns, leading to their successful patent filings.

PATENT AND PUBLICATIONS

- Mudgal, A., Kush, U., Kumar, A. et al. Multimodal fusion: advancing medical visual question-answering. *Neural Comput & Applic* (2024). [Link](#) | [View](#) | [GitHub](#)
Developed a multimodal architecture(vision-language) that interprets radiology images and answers user questions specific to that image, helping patients understand medical findings from their scans.
 - Boosted BLEU and ROUGE scores from 0.12/0.15 to 0.37/0.40 by fine-tuning BLIP and its subcomponents on a medical dataset. Designed custom Multimodal Architecture, achieving BLEU and ROUGE scores of 0.41 and 0.44.
- Authored A2 grade patent with Samsung, "Enhanced Optical Flow Video Stabilization for Multimedia Content", application number 202211062410.

TECHNICAL PROJECTS

- **Mindful Tweet Hub** Accelerated real-time **search speed to 10ms** by creating **vector similarity index using Hierarchical Navigable Small World(Faiss HNSW)** on **BERT embeddings** for similar tweet matching. [GitHub](#)
- Efficient LLM fine-tuning with LoRA and QLoRA. [Medium](#) | [GitHub](#)

EDUCATION

Master of Science in Data Science, The George Washington University (GPA: 4.00/4.00) Aug 2022 - May 2024
Relevant Coursework: Machine Learning, Deep Learning, NLP, Linux for DevOps, Time Series Forecasting

Bachelor of Technology in Communication & Computer Science, LNMIIT

Aug 2015 - May 2019

TECHNICAL SKILLS

Programming Languages

Python, C++, R, JAVA, MATLAB, Streamlit

Libraries & Frameworks

LangGraph, TensorFlow, Pytorch, Pandas, Transformer, pgvector, Matplotlib, Sklearn

Technologies & Tools

GitHub, Huggingface, Tableau, Docker

Cloud Services

AWS (Elastic Beanstalk, S3 buckets, EC2, SageMaker, Bedrock, Glue), Groq

Database

SQL, PostgreSQL, Redis, OpenSearch, MongoDB

Development & Operations

Agile Development, Software Architecture, RESTful APIs, Kubernetes

CERTIFICATION

- AWS Certified Cloud Practitioner (CLF-C02).