Arjun Parasuram Prasad

+1-9143360944 | arjun.prasad@nyu.edu | linkedin.com/in/arjun-p-prasad | prasadboi.github.io

Education

NYU, Courant Institute of Mathematical Sciences — **New York, United States**
Master of Science in Computer Science; GPA: 3.83/4 — **Aug. 2024 – May 2026**
Machine Learning, Computer Vision, Big Data & ML Systems, GPU Prog., Algorithms, OS

BITS, Pilani - Hyderabad campus — **Hyderabad, India**
Bachelor of Engineering - Computer Science; GPA: 8.65/10 — **Aug. 2019 – June 2023**
Deep Learning, NLP, Data Structures & Algo., Database Systems, Probability & Statistics

Technical Skills

ML Engineering: Search, Ranking, RAG, ML Systems, Vector DB, Natural Language Processing, Computer Vision, LLMs, Distributed Training & Inference, PyTorch, Huggingface, Weights & Biases, Langchain, vLLM
Data Engineering: SQL, Vector Databases, Postgres, Snowflake, Azure Data Factory, ETL, HDFS, Power BI, CI/CD
Software Engineering: Java, C/C++, Python, Docker, Kubernetes, Singularity, HPC, AWS, Distributed Systems, Rest API, Fast API

Work Experience

Data Analytics Engineer — June 2023 – Aug. 2024
Providence Health Care — *Hyderabad, India*

- Engineered ETL pipelines using Azure Data Factory for in-house data migration tool (Ultrasonic), handling 10M+ row datasets across diverse source/target systems
- Collaborated with marketing team stakeholders to develop Snowflake-based analytics framework, integrating advanced stored procedures with Power BI dashboards to internalize campaign marketing KPIs, eliminating outsourcing costs by $2M+
- Shipped CI/CD pipeline to automate SQL script deployments, reducing manual effort by ∼60 hours/month
- Stack: Azure Data Factory, Snowflake, Python, SQL, Power BI, Azure Dev Ops

Machine Learning Intern — Jan. 2023 – June 2023
SuperPe.in (Formerly Mewt.in) — *Bengaluru, India*

- Constructed analytical dashboards for understanding sales team performance - tracking merchant onboarding, and financial KPIs and helped identify and eliminate bottlenecks that reduced drop-offs by 40%
- Developed an OCR pipeline that extracted and validated merchant business document numbers with 97% accuracy using Google Vision APIs + OpenCV, eliminating manual entry requirements. Deployed using AWS Lambda and Amazon EFS
- Created a RAG-based conversational agent using LangChain and Milvus vector database to automate customer query resolution via agent-tool workflows
- Stack: Postgres, Fast API, OpenCV, Langchain, Milvus, AWS, Google Cloud Vision

Research Experience

LLMs as Event Forecasters
Mengye Ren, Agentic AI Learning Lab — Feb. 2025 - Current

- My role here as been to develop the RAG pipeline infrastructure and conduct evaluations of various retrieval mechanisms on their ability to improve LLM event-forecasting performance over daily news data.
- Stack: PyTorch, vLLM, Huggingface, Milvus, Singularity, SLURM, Git

Context Aware Behavioral Fingerprinting of IoT Devices
Barsha Mitra, BITS-Pilani — Aug. 2022 - Dec. 2022

- Collected/analyzed 600K+ packets using Wireshark; engineered 18 packet-level features (TCP window size, payload entropy) achieving 94% accuracy via Random Forest classifier.
- Designed IoT network with Raspberry Pi gateway + 8 NodeMCU/sensor nodes, reducing per-device costs by 90% compared to prior project setups.
- Published at SECRYPT 2023 with benchmark F1-score of 94.66% for device identification in dynamic operating contexts.
- Stack: Python, C, Raspberry Pi, Node MCU, Scikit-learn, Wireshark

Empirical Software Engineering
N L Bhanumurthy, BITS Pilani | Nucleus Software Exports Ltd. — Jan. 2022 - Jun. 2022

- Implemented Python scripts to extract 18 source code metrics (method coupling, complexity) via CKJM tool and 7 code-change metrics (LOC added/deleted) from 835+ SVN diffs.
- Developed ML models using scikit-learn to identify bug-prone Java files with 91%+ accuracy, informing Agile policy updates that reduced post-release defects.
- Automated metric extraction pipeline resolving multi-collinearity in feature sets for improved model reliability.
- Stack: Java, Python, Scikit-Learn, SVN

Publications

Prasad, A.; Biju, K.; Somani, S. and Mitra, B. (2023). **Context-Aware Behavioral Fingerprinting of IoT Devices via Network Traffic Analysis.** 20th International Conference on Security and Cryptography, SECRYPT 2023. (Link: https://dblp.org/rec/conf/secrypt/PrasadBSM23)

Prescod, Paul, Ben Feuer, Andrii Hladkyi, Sean Paulk and Arjun Prasad.**Auto-Markup BenchMark: Towards an Industry-standard Benchmark for Evaluating Automatic Document Markup.** Balisage: The Markup Conference 2023. (Link: https://www.balisage.net/Proceedings/vol28/html/Prescod01/BalisageVol28-Prescod01.html)