Lokesh Karanam

(+1) 657-252-8201 lokeshkaranam3@gmail.com linkedin.com/in/lokesh-karanam lokeshkaranam.me

Summary

Software Engineer with 5 years of experience specializing in ML model deployment, MLOps, and big data technologies in cloud environments. Proficient in tools and frameworks such as Hadoop, Spark, and NoSQL databases, with deep expertise in designing scalable ETL pipelines and data workflows on AWS and Snowflake. Skilled at blending cloud, data engineering, and machine learning technologies to deliver robust end-to-end solutions, working with cross-functional teams to optimize data-driven systems and accelerate business impact.

Technical Skills

Languages: C++, Python, Java, PHP, SQL, JavaScript, HTML/CSS, LaTeX
Databases & Data Warehousing: Snowflake, Databricks, MySQL, SQLite, MongoDB
Cloud Platforms: Amazon Web Services, Microsoft Azure
Data Science & ML Libraries: Scikit-Learn, NumPy, Pandas, SciPy, NLTK, Matplotlib, Statsmodels
Big Data & Processing Frameworks: Apache Spark, PySpark, Snowpark, Hadoop
Machine Learning & AI: PyTorch, Keras, LangChain, Generative AI (LLMs, RAG Applications)
Web Development: React, Angular, Bootstrap, RESTful APIs
Developer Tools: Git, VSCode, PyCharm, IntelliJ, MATLAB, Jupyter Notebook, PowerShell

Experience

Staff Software Engineer — May 2023 – Present

RR Donnelley — *Austin, TX*

- Developed a Spark application to efficiently parse and anonymize 50 TB of S3 data, removing PII such as names, emails, phone numbers, and addresses based on customer deletion requests
- Led the implementation and usage of Snowflake across multiple teams, providing guidance and best practices to ensure effective utilization of the platform, resulting in improved data management and cost savings
- Served as Tech Lead in streamlining engineering workflows by selecting and integrating an AI coding assistant, reducing code review times and accelerating feature delivery.
- Designed and developed end-to-end data processing services utilizing C# microservices and Snowflake stored procedures to ingest, clean, transform, and model customer data, delivering actionable insights to clients
- Collaborated with cross-functional teams to identify performance bottlenecks and refactor Spark code, significantly improving resource utilization and reducing SageMaker infrastructure expenses by 40% for the ML engine project
- Designed a serverless architecture to capture LLM API usage data via SQS and Lambda, ingesting prompt responses, models, and token counts into Snowflake for granular cost analysis
- Implemented S3 lifecycle configurations to transition objects to Intelligent-Tiering and delete non-current objects, reducing S3 storage costs by $120K annually across development, QA, and production accounts

Software Development Engineer — Jul. 2022 – Apr. 2023

Amazon Web Services — *Seattle, WA*

- Designed a serverless application using Lambda, Kinesis, SQS, and SNS to process high velocity (4k records/sec) real-time data
- Served as an on-call engineer for the team, responsible for responding to incidents and ensuring service availability
- Worked with stakeholders to define and implement Service Level Objectives (SLOs) and Service Level Agreements (SLAs), ensuring that our team service performance meets the needs of customers
- Actively participated in post-incident reviews and root cause analyses, identifying areas for improvement and implementing corrective actions to prevent similar incidents from occurring in the future

Data Scientist — Teaching Assistant — Jan. 2021 – Jul. 2022

MU Institute for Data Science and Informatics — *Columbia, MO*

- Predicted Acute Kidney Injury (AKI) in ICU patients using a novel Residual Time-Gated LSTM architecture, improving prediction accuracy by 7% over Vanilla LSTM
- Developed an Ensembled SVM model with probabilistic estimates, boosting classification accuracy by 6%
- Designed a system using Detectron2 for object detection on labeled street-view images to analyze neighborhood characteristics
- Built and optimized scalable ETL pipelines, including data extraction, transformation (aggregation, filtering, normalization), and validation for large datasets
- Assisted in course delivery for Data Mining & Information Retrieval, Database Analytics, and Communication Network Analytics

Technical Analyst **July. 2020 – Dec. 2020**

Axtria Ingenious Insights *Noida, India*

- Optimized the overall testing process by automating the workflow thereby reducing the human effort by around 60 hours/month
- Involved in standardized documentation creation, storage, and maintenance process at central repository reducing transition efforts by 50%
- Worked as a Java engineer, specializing in automation testing, using Java-based frameworks such as Selenium and TestNG
- Developed and maintained test suites and frameworks, enabling efficient and effective testing of applications

Publications & Research

Intrusion Detection Mechanism for Large Scale Networks using CNN-LSTM **2020**

Authored research paper published in IEEE Xplore focusing on network security and intrusion detection *IEEE Xplore*

Classifying Cover Crop Residue from RGB Images: SVM vs SVM Ensemble **2021**

Developed and compared machine learning approaches for agricultural image classification *IEEE SSCI*

Continuous Anticipation of AKI in the ICU using Time-Gated LSTMs **2022**

Research on predictive healthcare analytics using deep learning for critical care applications *AMIA CIC*

Understanding Personal, Social Network, and Social Network Contributors to Firearm Risks **2022**

Total project funding Received: 1,700,000 USD *CDC*

Certifications

Databricks Certified Generative AI Engineer Associate **2024**

Hands on Deep Learning by PadhAI **2023**

Machine Learning by Coursera **2022**

Education

University of Missouri - Columbia **Columbia, MO**

Master of Science in Computer Science; GPA: 3.9/4.0 *Jan. 2021 – Jul. 2022*

- **Course Work**: Cloud Computing, Machine Learning and Pattern Recognition, Unsupervised Learning, Advanced Web Development, Computational Intelligence, Principles of BigData Management, Advanced Natural Language Processing, Artificial Intelligence 2

Atal Bihari Vajpayee-Indian Institute of Information Technology and Management **Gwalior, India**

Integrated Post Graduation in Information Technology; GPA: 7.2/10 *Jul. 2015 – Jun. 2020*

- Selected for Merit-cum-Means scholarship, which is awarded to Top 10 students in the university.
- **Course Work**: Computer Vision, Data Structures and Algorithms, Object Oriented Programming, Computer Networks, Operating Systems, Computer Organization and Architecture, Database Management Systems, Data Mining and Data Warehousing, Natural Language Processing.