

Email: ranjana.rajendran@gmail.com

Ph: [Mobile] (831) 332 9892

Github: ranjanarajendran

Portfolio , LinkedIn Profile

Work Authorization: Eligible to work for any U.S. employer without sponsorship

EDUCATION

MS, Computer Science, *University of California Santa Cruz*, September '10 - March '13

Course Work: Analysis of Algorithms, Advanced Operating Systems, Security in Computer Systems, Computer Graphics, Programming Languages, Data Mining, Machine Learning, Mathematical Statistics, Design & Implementation of Database Systems

BTech, Computer Science & Engineering, *Cochin University of Science and Technology, India*, September '02 - May '06 (First Class with Distinction & Honours)

PUBLICATIONS

"Detecting localized homogeneous anomalies over spatio-temporal data", (Data Mining and Knowledge Discovery, Volume 28), *Internship project with IBM India Research Lab, 2012*

"Horus: Fine-Grained Encryption-Based Security for High Performance Petascale Storage", *Proceedings of the 6th Parallel Data Storage Workshop (PDSW '11)*, Nov '11

PATENT

US-20250053985-A1 : Method and System for Cryptocurrency Fraud Detection, result of a hackathon while at Mastercard. Pub. Date: Feb 13, 2025

AWARDS & RECOGNITIONS

Professional: Spot Bonus Award at Cloudera for team performance (Q1 2014)

Academic: UC Regents Graduate Fellowship (2010-2012) ; **District Merit Scholarship**, 2000-2006 by Department of Collegiate Education, Kerala, India ; GRE: 800/800 Quantitative

WORK EXPERIENCE

Independent ML Engineer & Author

August '23 - Present

- Authored Engineering Autonomous AI (in progress) – comprehensive technical guide to building production agentic systems covering LLM deployment, multi-agent orchestration, and MLOps best practices
- Built and deployed Tech Stack Advisor – production-grade multi-agent AI system featuring 5 specialized agents, LangGraph orchestration, Qdrant vector store for conversational memory, full CI/CD pipeline with Prometheus/Grafana monitoring, JWT authentication, and Railway cloud deployment
- Developed MLOps infrastructure and production AI systems focusing on LLM deployment challenges: context management, retrieval systems (RAG), vector databases (FAISS, Qdrant), and infrastructure for serving models at scale
- Implemented deep learning projects including computer vision (Group Emotion Recognition with RetinaFace/DeepFace on Google Vertex AI), NLP (spoiler detection with BERT/Longformer achieving 71.5% accuracy), and generative models (Variational Autoencoder for Fashion MNIST)
- Executed traditional ML projects including time series regression (volcano eruption prediction with 99.7% R² using RidgeCV), graph ML (social network clustering with Node2Vec achieving 0.73 silhouette score), and comprehensive EDA/feature engineering workflows

Lead Software Engineer, *Ekata-Mastercard (Identity Verification)*, Seattle

August '22 - August '23

- Engineered PII suppression and identity workflow systems powering ML-based fraud scoring for enterprise customers, ensuring compliance with privacy regulations while maintaining less than 100ms latency
- Built distributed ETL pipelines processing terabytes of identity data using Spark on Databricks & EMR, orchestrated via Airflow, with automated data quality checks and observability dashboards
- Designed and managed Redis-backed identity graph with 50M+ entities, enabling real-time ML fraud scoring with dynamic suppression and stateful updates
- Co-invented patented fraud detection system (US-20250053985-A1) integrating multi-signal reasoning and real-time anomaly detection with production ML pipelines

System Development Engineer, *Amazon Web Services*, Seattle

July '17 - August '22

- Owned monthly build and release cycles for *Elastic Map Reduce* (EMR), maintaining 99.9% service uptime across ML and big data frameworks (Hive, Spark, Presto, Flink). Developed automated Java-based deployment tools reducing release time by 40%
- Built and maintained integration testing framework with 200+ automated tests covering all EMR applications, reducing release-blocking bugs by 60% and ensuring reliability for downstream ML pipelines
- Resolved 100+ critical integration issues between ML frameworks (Spark MLlib, TensorFlow on Spark) and the underlying OS, reducing customer-impacting incidents by 40%
- Designed & integrated new EC2 compute types (GPU instances, Graviton processors) and applications into EMR, enabling distributed deep learning training and improving ML workload performance by up to 50%

- Participated in 24/7 on-call rotation, triaging and resolving production incidents for ML and big data workloads, maintaining high service availability for enterprise customers
- Authored technical documentation and collaborated with product managers to launch new EMR features, directly supporting data scientists and ML engineers at Fortune 500 companies

Senior Developer Support Engineer, Qubole, Santa Clara

April '16 - May '17

Optimized large-scale data and feature engineering pipelines on Qubole's Hadoop-as-a-Service platform using MapReduce, Spark, Hive, and SparkSQL, enabling downstream ML model training and analytics. Achieved up to 70% faster runtimes and 90% fewer pipeline failures, improving reliability of ML workflows. Collaborated with engineering and product teams to drive platform enhancements based on real-world data science and ML challenges.

Hadoop Engineer, Altiscale, Palo Alto

February '15 - March '16

Handled customer feature requests and support tickets for production *Map-Reduce / Spark jobs*, ML data pipelines, ETL, and security (*Kerberos*). Specifically, the activities included:

- Troubleshoot and optimized customer Hive, Spark, and Spark SQL applications, achieving up to 50% query performance improvements through partition tuning, memory configuration, and query plan optimization
- Performed Java performance and memory profiling using Eclipse Memory Analyzer to diagnose and resolve OOM errors and bottlenecks in production MapReduce and Spark jobs processing terabytes of data
- Translated customer ML pipeline requirements into product feature requests and authored 20+ knowledge-base articles documenting best practices for big data and ML workloads
- Configured and troubleshoot Kerberos authentication and security in multi-tenant Hadoop environments, ensuring secure access control for enterprise customers

Solution Architect, Cloudera, San Francisco

September '13 - January '15

ML Ops Engineer – Hadoop Ecosystem Solutions

- Deployed and managed CDH clusters (up to 45 nodes) on AWS EC2 and private cloud infrastructure, supporting enterprise customers running large-scale ML training pipelines and data analytics workloads processing hundreds of terabytes
- Configured MySQL backend with master-slave replication for Cloudera Manager and CDH metadata services, ensuring high availability and fault tolerance for cluster management infrastructure
- Architected Kerberos authentication for CDH with multiple integration patterns: local KDC with backup KDC, one-way cross-realm trust to Active Directory, and direct AD integration for enterprise SSO
- Designed and implemented batch and real-time data pipelines using Spark, Hive, and Impala on HDFS, enabling automated data ingestion and transformation for ML model training with 99.5% SLA adherence
- Integrated real-time data sources using Flume and RDBMS ingestion via Sqoop, building data warehouses supporting continuous ML model retraining for enterprise customers
- Automated end-to-end ML pipeline orchestration using Oozie with integrated monitoring (Tableau dashboards), reducing pipeline failures by 40% and enabling self-service analytics for data science teams
- Implemented role-based access control (Sentry) for secure multi-tenant access to Hadoop resources

Certifications: Cloudera Certified Developer (CCDH), Administrator (CCAH), HBase Specialist (CCSHB)

Software Engineer, Gap Inc via Nisum Technologies, San Francisco June '13 - September '13

- Developing testing frameworks in Ruby using *Watir & Taza frameworks*.
- Release testing and debugging for an order management system implemented in Java with several tools including SoapUI.

Research Intern - IBM India Research Lab, Bangalore,

July '12 - December '12

Researched, developed, and implemented an ML-based anomaly detection system (Java, R, Unix shell) for spatial and spatio-temporal datasets, resulting in a publication in *Data Mining and Knowledge Discovery* (Vol. 28).

Graduate Student Researcher, Storage System Research Centre, UC Santa Cruz

January '11 -

December '11

Researched, developed, and implemented an encryption model for the Linux based distributed file system (Ceph), which resulted in a publication at a prestigious workshop.

Teaching Assistant, UC Santa Cruz

Fall '10, Fall '11, Spring '12

Teaching assistant for undergraduate courses in Computer Programming & Unix systems, Mathematics (Precalculus), and Business Information Systems. Examination grader and invigilator for graduate and undergraduate courses. Assessed students' comprehension, prepared examinations accordingly, gave feedback, and evaluated their performance.

MTS - Software Engineer, Lucent Technologies, R & D Centre, Bangalore

Dec '06 - Dec '08

Performed duties as module leader, engineered innovative solutions to meet customer requirements, enhancing system efficiency by 30% of an OA & M GUI for Radio Access Networks, and trained team members in various J2EE technologies.

Software Engineer, Wipro Technologies, Cochin

August '06 - December '06

Designed and implemented fixes for middleware components in digital television set-top boxes.

TECHNICAL
SKILLS

Machine Learning & MLOps Infrastructure:

- *MLOps & Orchestration:* MLflow (experiment tracking, model registry), Airflow, Databricks, Model versioning, A/B testing, Canary deployments, Drift detection, Feature stores
- *ML Frameworks:* PyTorch, TensorFlow, Scikit-Learn, PyTorch Lightning, Spark MLlib, Sentence-Transformers
- *Model Deployment & Serving:* Docker, Kubernetes, FastAPI, gRPC (Protocol Buffers), GraphQL (Strawberry), Hypercorn (HTTP/2), AWS (EMR, ECS, SageMaker), Google Vertex AI, Railway, Model serving APIs, Production inference optimization
- *Monitoring & Observability:* Prometheus, Grafana, Model performance dashboards, Latency tracking (p50/p95/p99), Prediction monitoring, Data quality checks, Custom metrics collection
- *Performance Testing:* Locust (load testing), Benchmark analysis, Throughput optimization, Protocol comparison (REST/gRPC/GraphQL)
- *LLM & GenAI:* Langchain, LangGraph, OpenAI, Anthropic Claude, HuggingFace Transformers, BERT, Longformer, RAG pipelines, Multi-agent orchestration
- *Vector Databases:* FAISS, Qdrant (production RAG systems)
- *Specialized ML:* RetinaFace/DeepFace (computer vision), Node2Vec (graph embeddings), NetworkX, tsfresh (time series), Pydantic

Big Data & Distributed Systems:

- *Data Processing:* Spark, Hadoop, MapReduce, Hive, Presto, Impala, HDFS, Oozie, PIG
- *Distributed Systems:* Multithreaded Java programming, Concurrency management, Distributed file systems (HDFS, Ceph), Thread pool executors
- *Message Queues & Streaming:* Apache Kafka, RabbitMQ, Amazon SQS/SNS, Redis Pub/Sub

Cloud & DevOps:

- *Cloud Platforms:* AWS (EC2, EMR, ECS, CodeBuild, S3, Lambda), Google Cloud (Vertex AI), Railway
- *CI/CD & Automation:* Jenkins, GitHub Actions, Docker, Kubernetes, Docker Compose, Infrastructure as Code

Databases & Storage:

- *SQL Databases:* MySQL, PostgreSQL, SQLite, OracleDB, Spark SQL
- *NoSQL & Caching:* Redis, HBase, Cassandra

Programming & Software Engineering:

- *Languages:* Python, Java, C, C++, Scala, R, Ruby, SQL, UNIX shell scripting
- *Architecture & Design:* OOAD, Design Patterns, Microservices, SOA, Agile, BDD methodologies
- *Testing & Quality:* pytest, JUnit, MRUnit, mypy, ruff, structlog, Cucumber, Apache Bigtop

Web Development & APIs:

- *Frameworks:* FastAPI, Flask, Spring MVC, REST APIs, gRPC (Protocol Buffers), GraphQL (Strawberry)
- *API Design:* Swagger/OpenAPI, SOAP, HTTP/2, JSP, Apache Tomcat, NginX, ASGI servers

Data Visualization: Streamlit, Plotly, Dash, Matplotlib, Tableau

Security: Kerberos, Sentry, JWT, OAuth 2.0, bcrypt, Rate Limiting, Firewalls (IPTable rules)

Version Control & Tools: Git, SVN, Vim, Visual Studio Code (with Claude Code), PyCharm, RStudio, Eclipse, Jupyter

Documentation & Typesetting: $\text{T}_{\text{E}}\text{X}$ ($\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$, $\text{BIB}\text{T}_{\text{E}}\text{X}$), Markdown, HTML, CSS, Technical Writing

LANGUAGE
SKILLS

English : Fluent in reading, writing, and speaking. My education was in English. (TOEFL: 112/120)
Hindi : Fluent in reading, writing, and speaking.
Malayalam (native speaker) : Fluent in reading, writing, and speaking.

WORK
AUTHORIZATION

United States : Permanently authorized to reside and work for any employer without sponsorship.
India : Permanently authorized to reside and work without visa.
European Union : Derived eligibility to reside and work in member states under applicable provisions.