

JACOB KURIAKOSE

jkuriak3@asu.edu | +1 (602) 802-6591 | Tempe, AZ | www.linkedin.com/in/jacob-kuriakose

PROFESSIONAL SUMMARY

Data Scientist specializing in NLP, Time Series Forecasting, and multi-agent AI systems, with experience shipping production-grade AI in both research and startup environments. Actively seeking full-time opportunities in data science or ML engineering.

PROFESSIONAL EXPERIENCE

Arizona State University, AI Backend Developer

May 2025 – Present

- Developed and maintained FastAPI backend powering Waterbot, a production RAG chatbot at azwaterbot.org, handling 13 API endpoints including chat, session transcription, ratings, and multi-language translation for 2,000+ campus users.
- Integrated AWS Bedrock Knowledge Base as the RAG retrieval layer, replacing a self-hosted Chroma vector DB; re-indexed embeddings and wired managed retrieval into ECS tasks, improving scalability.
- Provisioned AWS CDK infrastructure (ECS Fargate, ALB, CloudFront, ECR, RDS, DynamoDB, S3, Lambda), cutting manual setup ~40% and deployment downtime ~30%.
- Engineered CI/CD pipeline with GitHub Actions; used ECS CLI redeploys over full stack deploys, reducing release cycle ~85%; configured EC2 IAM role-based auth eliminating static credentials.
- Diagnosed and removed unintended CloudFront basic-auth function; resolved IAM and CDK bootstrap bottlenecks.

Red Hat, Software Quality Engineer Intern

January 2023 - July 2023

- Performed manual and automated testing with GoLang, reducing average test execution time by 30%; crafted detailed test cases in Polarion covering critical software paths.
- Diagnosed software defects and used Docker containers to deploy fixes, enabling reliable and scalable deployments.

BraynixAI, Natural Language Processing Intern

June 2022 - August 2022

- Improved legal document classification accuracy by 10% using Transformer-based NLP techniques.
- Built FastAPI-based ML model APIs to process legal and non-legal documents, streamlining workflows.

PROJECTS

Medical Chatbot ([Source Code](#))

Personal Project, Generative AI Model

- Architected a ReAct (Reasoning + Acting) medical chatbot using LangChain and LangGraph StateGraph, replacing a static RAG pipeline with an explicit Thought → Action → Observation loop that improves reasoning transparency and answer interpretability.
- Built a hybrid retrieval system combining Pinecone vector search (7K embeddings) with CrossEncoder reranking and domain-filtered DuckDuckGo web search, enabling the system to balance verified medical knowledge with real-time clinical updates.
- Evaluated on a 10-question QA benchmark (16 metrics): BERTScore F1 = 0.853, zero hallucination, perfect source grounding and reasoning efficiency (1.0) confirming factually grounded, semantically accurate outputs.
- Integrated a Groq LLM safety classifier as a guardrail on every response and LLM-based conversation summarisation to preserve multi-turn context while reducing token consumption; deployed on Hugging Face Spaces via Docker with automated CI/CD.

Data-Driven Walmart Sales Predictions ([Source Code](#))

Personal Project, Time Series Forecasting

- Analyzed SARIMA, Exponential Smoothing, LSTM, Random Forest, and LightGBM models to forecast store-level sales; achieved 38.92% RMSE improvement with Exponential Smoothing over LSTM, enhancing inventory planning accuracy.
- Implemented MLflow for experiment tracking and DVC for data and model versioning; exposed forecasts via a FastAPI endpoint.

SKILLS

Competencies: Data Cleaning, ML Algorithms (Linear Regression, Logistic Regression, Random Forest, KNN, SVM, etc.), Data Visualization, Exploratory Data Analysis, Data Mining, Generative AI, A/B Testing.

Programming Languages: Python (NumPy, Pandas, Matplotlib, Scikit-learn, Seaborn, SpaCy, NLTK etc.), C++, SQL.

Tools/Frameworks: TensorFlow, Keras, PyTorch, Statsmodels, Git, PySpark, Jupyter Notebook, Tableau, DVC, MLFlow, MS Excel.

EDUCATION

Master of Science in Data Science, GPA: 4.00

August 2024 – May 2026

Arizona State University

Tempe, Arizona

Courses: Software Security, Data Mining, Statistical Machine Learning, Database Management System

Bachelor of Technology in Information Technology, CGPA: 9.0

August 2019 – July 2023

Dr. Akhilesh Das Gupta Institute of Technology and Management

Delhi, India

Published Research Paper: [Monument Tracker: Deep Learning Approach for Indian Heritage](#)

AWARDS & ACCOLADES

Co-founded [Navia](#) out of a hackathon win, leading AI architecture across multi-agent orchestration, Pinecone vector memory, and GPT-5.4 LLM integration, and email/calendar integrations on a stack spanning Next.js, Supabase, and OpenAI; secured non-dilutive funding, launched ASU GATE pilot, and represented Navia at ASU+GSV Summit engaging investors and edtech leaders.