Anish Panicker

New Jersey, USA | panicker.anish27@gmail.com | Linkedin | My Portfolio | Github | Phone: (201)726-0589 | Open to Relocate |

Skills

- Languages: Python, SQL, R Programing, C++, Javascript, Java
- · Cloud, Systems: AWS (Athena, Bedrock, S3, RDS, EC2, Lambda, SageMaker, Kubernetes), Azure (Databricks, Datalake, Synapse) • Databases: MySQL, MongoDB, Vector DB
- · Libraries, Tools: Keras, PyTorch, TensorFlow, Hugging Face Transformers, OpenCV, Streamlit, Scikit-learn, LangChain, CUDA
- · Developer Tools: VS Code, AWS QuickSight, Apache Airflow, Jupyter Notebook, Power BI, Tableau, Matlab
- Technologies: Hadoop, Snowflake, Talend, Docker, Big Data Pipelines (ETL workflows, data ingestion, real-time stream processing, batch processing)
- Data Science Techniques: KYC/KYB, Feature Engineering, MLOps, CI/CD pipelines

Experience

New Jersey Courts | Machine Learning Engineer Intern

- Fine-tune LayoutLM for guardianship fraud detection using a custom-annotated court document dataset to extract data and flag fraud, targeting **\$10M+ in loss prevention** while ensuring ethical handling of sensitive case information.
- Implement scalable MLOps pipelines using AWS, RDS, and MLflow to streamline model inference, integrate CI/CD, and support deployment of models capable of processing over 50K records monthly.
- Design and maintain a fraud analytics dashboard using Tableau and industry-standard tools to visualize guardianship application data.

Infmobile | Machine Learning Engineer

- Built a churn-prediction pipeline to pinpoint at-risk subscribers and fuel targeted retention campaigns, increasing retention by 12%.
- Developed a custom chatbot using RAG, LangChain, PyTorch, and other frameworks to autonomously resolve customer queries and technical issues, cutting support tickets by 20%.
- Engineered recommendation engines for mobile add-ons and partnered with DevOps to containerize and orchestrate ML pipelines (AWS, Docker, Kubernetes, MLflow), enabling real-time monitoring for over 50 000 subscribers with sub-second inference.

Pittsburgh Technologies | Machine Learning Engineer

- · Conducted in-depth research on AML algorithms to enhance credit card fraud detection, integrating advanced machine learning models into the Nastel Xray platform and leveraging familiarity with KYC/KYB best practices.
- Enhanced Nastel Xray by integrating anomaly detection models and ensuring data integrity in processing sensitive, regulated data; decreased false positives by 15% and improved fraud detection accuracy by 22%.
- Implemented advanced data visualization techniques to transform data presentation and insights, reducing decision-making time by 20% and supporting comprehensive analyses of complex, large-scale datasets.

Projects

Private Document QA Assistant (LangChain, FAISS, Python, Docker, RAG)

- · Orchestrated a privacy-compliant QA assistant leveraging Retrieval-Augmented Generation (RAG) for on-premises PDF/TXT ingestion, sustaining sub-200 ms query latency across a 10 GB document corpus.
- Deployed two LangChain agents, Document Analyzer (auto-flagged 1k+ outdated/duplicate chunks with Excel reporting) and Faculty Onboarding (week-1 checklists & 5-question quizzes)

AI Agents for Code Generation (ReAct Agent Model, Mistral, CodeLlama, RAG, Python)

- Designed a multi-agent system using ReAct with Mistral and CodeLlama, achieving 91.7% code accuracy and 25% better retrieval via a llama index-based vector pipeline.
- Created AI agent tools for API doc retrieval, code reading, and file naming suggestions, integrating Pydantic-Flask post-processing to validate and structure JSON.

Medical Image classification (MLflow, DVC, Docker, TensorFlow, CI/CD Pipeline)

- Applied VGG16 transfer learning on 12,000+ CT scans to classify images into four categories with 90% accuracy, enabling effective medical image interpretation.
- · Configured scalable deployment on AWS using Docker and MLflow, and designed a CI/CD pipeline with DVC and Flask API to automate data versioning and support real-time inference with A/B testing of model metrics.

Multi-agent Customer Support Bot (AI Agents, Langflow, Streamlit)

·Developed a Streamlit-based customer support chatbot using Langflow API with real-time responses from structured data (Csv, JSON, Pdf). Created 3 custom tools for agent interaction and modular deployment.

Certifications

Education

- AWS Certified AI Practitioner
- Google Data Analytics Specialization

- Bloomberg Market Concepts
- IBM Developer Skills Network

New Jersey Institute of Technology Master's, Data Science

Jul 2023 - Jan 2025

Sep 2022 - Jun 2023

Feb 2025

Oct 2024

Mar 2025

Mar 2025

Feb 2025 - Present