

Aditya Sankranthi — AI/ML Engineer

(530) 562-5883 | sankranthiadiyaa@gmail.com

PROFESSIONAL SUMMARY:

- Accomplished AI/ML Engineer with 5 years of experience in designing and deploying robust AI systems.
- Proficiently develops software projects emphasizing AI solutions using Python in VSCode/Pycharm.
- Expert in deploying and maintaining scalable AI systems across Databricks and Microsoft Azure services.
- Skilled in leveraging Azure Kubernetes Service (AKS) for containerized AI application deployment in production.
- Maintains and extends internal APIs, SDKs, and clients crucial for supporting various project initiatives.
- Actively participates in architecting sophisticated software and cloud solutions on Azure and Databricks.
- Strong experience with GitHub for version control, commits, pull requests, and CI/CD integrations.
- Familiar with Databricks, including Unity Catalog, MLflow, and collaborative notebooks for MLOps.
- Proven ability to lead technical walkthroughs and clearly articulate complex design decisions to stakeholders.
- Hands-on experience deploying applications, servers, and ML models effectively in a production setting.
- Comfortable thinking in terms of robust engineering systems, moving beyond isolated scripting efforts.
- Expert in creating CI/CD workflows in GitHub to streamline deployment processes for AI solutions.

EDUCATION:

Master of Science in Computer Science @ University of Wisconsin- Milwaukee

WORK EXPERIENCE:

AI/ML Engineer @ HCA Healthcare | Nashville, TN

Sep 2024 – Present

- Led the design and development of an end-to-end AI system to predict patient readmission risk, utilizing Python.
- Implemented robust data cleaning, preprocessing, and feature engineering for Electronic Health Record (EHR) data on Databricks.
- Developed advanced deep learning classification models using TensorFlow, deployed as scalable AI systems on Azure.
- Optimized model performance and reliability through rigorous hyperparameter tuning and cross-validation, ensuring production readiness.
- Utilized MLflow for comprehensive experiment tracking, model versioning, and reproducibility across iterations within Azure.
- Deployed high-availability inference pipelines using Docker containers on Azure Kubernetes Service (AKS) for scalable model serving.
- Implemented proactive model monitoring to detect data drift and automate periodic retraining within the Azure environment effectively.
- Contributed to architecting robust cloud solutions on Microsoft Azure, integrating with various enterprise services and data sources.
- Collaborated effectively with clinicians and stakeholders in Agile sprints, communicating technical design decisions and managing expectations.

Technologies Used: Python, TensorFlow, Pandas, NumPy, MLflow, Docker, Azure Kubernetes Service, Microsoft Azure, SQL, GitHub, JIRA

AI Engineer @ Citigroup | New York, NY

Aug 2022 – Jul 2023

- Built and deployed sophisticated AI systems for fraudulent credit card transaction detection using Python and Scikit-learn.
- Performed extensive exploratory data analysis and preprocessing on high-volume transactional data within Databricks notebooks.
- Engineered impactful features including transaction velocity and spending patterns to enhance fraud detection model accuracy.
- Developed highly effective classification models, addressing class imbalance and optimizing performance for production deployment.
- Managed the end-to-end model lifecycle, including experiment tracking and versioning, using MLflow on Azure Machine Learning.

- Integrated trained ML models into Azure ML inference pipelines for efficient, batch fraud scoring in a production setting.
- Actively participated in architecting robust cloud solutions on Microsoft Azure, contributing to system design.
- Worked closely with fraud analysts, validating AI system performance and improving detection quality through clear communication.
- Utilized GitHub for version control, code reviews, and managing pull requests within a structured software development lifecycle.

Technologies Used: Python, Scikit-learn, Pandas, NumPy, Spark, MLflow, Azure Machine Learning, SQL, GitHub, JIRA

Machine Learning Engineer @ Target | Minneapolis, MN

Apr 2020 – Jul 2022

- Developed a cutting-edge recommendation engine to personalize product suggestions for online retail customers.
- Collected, cleaned, and transformed diverse customer interaction and purchase history datasets using Pandas and NumPy.
- Performed advanced feature engineering based on customer behavior, browsing history, and product metadata.
- Built robust recommendation models leveraging deep learning techniques with PyTorch on GCP AI Platform.
- Optimized recommendation quality through systematic hyperparameter tuning and comprehensive offline evaluation metrics.
- Exposed prediction services through scalable FastAPI endpoints for seamless integration with retail applications.
- Assisted in continuous monitoring of recommendation quality and periodically retrained models with recent customer data.
- Collaborated effectively with data scientists, software engineers, and product teams using Agile methodologies.
- Utilized GitHub for version control and documentation, ensuring maintainability and knowledge transfer of AI systems.

Technologies Used: Python, PyTorch, Pandas, NumPy, FastAPI, GCP AI Platform, PostgreSQL, GitHub, JIRA

TECHNICAL SKILLS:

- **Programming Languages:** Python, SQL, JavaScript
- **Cloud Platforms:** Microsoft Azure (Azure Kubernetes Service, Azure Machine Learning, Azure DevOps), AWS SageMaker, Google Cloud AI Platform
- **AI/ML Frameworks:** TensorFlow, PyTorch, Scikit-learn
- **MLOps & Deployment:** Databricks (Unity Catalog, MLflow, Notebooks), Docker, FastAPI, Flask, REST APIs
- **Version Control & Collaboration:** GitHub (Commits, Pull Requests, Integrations), Git, JIRA, Confluence, Azure DevOps
- **AI/ML Concepts:** RAG Architecture, Vector Embeddings, Semantic Search, Agentic Orchestration, Deep Learning, Supervised & Unsupervised Learning
- **Data Processing:** Pandas, NumPy, Apache Spark, Feature Engineering, EDA
- **Databases:** PostgreSQL, MySQL, SQL Server, MongoDB, Snowflake, ServiceNow Data Models
- **CI/CD & Automation:** GitHub Actions, Azure DevOps Pipelines
- **APIs & Services:** REST APIs, SDKs, Internal Clients