

Jaylen Callahan

📞 (513) 555-0184

✉️ jaylen.callahan.ai@example.com

🌐 linkedin.com/example/jaylencallahan

📍 Cincinnati, OH 45202

SUMMARY

Mid-level AI Engineer with 4+ years delivering production machine learning systems and AI-powered applications for manufacturing, supply chain, and operations analytics. Builds Python services, forecasting and **anomaly detection** models, **LLM applications**, **MCP servers**, and **REST APIs** across AWS and Databricks environments. Owns work from requirements discovery and feature engineering through training, evaluation, deployment, monitoring, and iterative improvement. Converts ambiguous operational needs into prioritized product backlogs, model success criteria, and maintainable designs. Partners with BI developers, data scientists, platform engineers, domain experts, and executives, explaining tradeoffs, confidence intervals, limitations, and business value in plain language. Brings practical focus on responsible AI, reliable delivery, and measurable decisions that save time and improve operational performance.

EXPERIENCE

AI Engineer

January 2025 - Present

Northstar Manufacturing Analytics

Cincinnati, OH

Leads production AI product delivery for plant scheduling, aftermarket operations, commercial workflows, and maintenance intelligence. Guides engineers across AWS, Databricks, Python, LLM applications, APIs, MLOps, and responsible AI, connecting technical decisions with faster operational decisions and dependable user adoption.

- Translated planner interviews and executive goals into a prioritized AI/ML backlog with acceptance criteria, delivery estimates, and business measures for scheduling and aftermarket decisions.
- Built a production RAG assistant with Python, FastAPI, OpenSearch vector search, and AWS services, giving service analysts cited contract and maintenance guidance with clear confidence limits.
- Created an MCP server for approved inventory, work-order, and forecast capabilities, adding authentication, versioning, **rate limiting**, and API documentation for reusable enterprise integration.
- Implemented MLflow and **Weights & Biases** tracking, registry promotion, automated evaluation, and **A/B testing** for forecasting and LLM releases across development and production.
- Designed Databricks medallion pipelines for demand, supplier, and work-order data, adding freshness checks, feature validation, **feature engineering**, and drift alerts for dependable daily outputs.
- Mentored four engineers on **prompt engineering**, SHAP, LIME, secure coding, CI/CD, automated testing, and **responsible AI** reviews; presented model tradeoffs and uncertainty to business leaders.

Machine Learning Engineer

May 2023 - December 2024

BlueRidge Industrial Software

Louisville, KY

Developed and deployed machine learning products for equipment service, parts demand, repair volume, and operational dashboards. Combined modeling, API engineering, deployment controls, and stakeholder discovery so manufacturing and supply chain teams could act on predictions with greater confidence and less manual analysis.

- Developed supervised classification and anomaly detection models for equipment service events with Python, scikit-learn, Databricks, and structured feature engineering, supporting faster maintenance triage.
- Delivered **time-series forecasting** pipelines for parts demand and repair volume, comparing seasonal baselines with gradient-boosting methods and documenting accuracy by product family.
- Exposed inference through versioned Flask and **FastAPI** REST APIs with token authentication, request validation, latency logging, OpenAPI documentation, and controlled rollback procedures.
- Established MLflow model registry conventions and **GitHub Actions** deployment checks covering unit tests, data-contract validation, container security scans, and AWS release approvals.
- Partnered with **BI developers** on predictive recommendations and natural-language search inside operational dashboards, using feedback to revise prompts and improve retrieval quality.
- Presented **proof-of-concepts** to manufacturing and **supply chain** stakeholders, linking prediction quality, response time, and adoption with labor planning and inventory decisions.

Associate Data Scientist

June 2022 - April 2023

Riverbend Operations Lab

Indianapolis, IN

Supported early AI/ML product discovery for production, supplier, and maintenance operations. Built reproducible analyses, data-quality controls, semantic search prototypes, and model comparisons that helped domain experts select credible use cases and move promising concepts into production-readiness review.

- Analyzed production, supplier, and maintenance datasets to frame **anomaly detection** and demand forecasting use cases, documenting assumptions, data gaps, and expected operational value.
- Built reproducible **Python** notebooks and Databricks jobs for **supervised learning**, clustering, time-series features, and model comparison with controlled train-test splits.

- Added data-quality checks for missing timestamps, duplicate work orders, inconsistent equipment identifiers, and stale source records before training data reached analysts.
- Prototyped **semantic search** with embedded maintenance documents and a vector database, evaluating retrieval precision against labeled technician questions.
- Used **SHAP** explanations and error analysis to identify influential features and communicate model limitations, confidence intervals, and false-positive patterns to domain experts.
- Supported Agile planning, technical documentation, peer code reviews, and live demonstrations that moved two **proof-of-concepts** into formal production-readiness assessments.

PROJECTS

Operations RAG Knowledge Service 📅 2025

Created a reusable retrieval-augmented generation service for maintenance and contract knowledge, combining document ingestion, vector search, prompt templates, source citations, response evaluation, and API access controls.

Supply Chain Forecast Evaluation Lab 📅 2024

Built a comparative forecasting framework for parts demand and repair volume, pairing seasonal baselines with machine learning models, experiment tracking, error analysis, and business-facing accuracy reports.

Predictive Maintenance Search Prototype 📅 2023

Developed a semantic search prototype for technician questions using embedded maintenance documents, vector database retrieval, labeled evaluation data, and analysis of precision, latency, and response relevance.

LEADERSHIP & AWARDS

- Applied AI Capstone Showcase, Best Operations Analytics Project, 2022
- Cincinnati Data Science Challenge, Finalist, 2021
- College of Engineering Dean's List, 2021 and 2022

EDUCATION

Bachelor's Degree in Computer Science

2022

University of Cincinnati GPA: 3.8

Cincinnati, OH

Coursework: Machine Learning, Data Structures, Database Systems, Software Engineering, Statistics

CERTIFICATIONS

- AWS Certified Machine Learning - Specialty 📅 2025
- Databricks Certified Machine Learning Associate 📅 2024
- DeepLearning.AI Generative AI with Large Language Models 📅 2024

TECHNICAL SKILLS

- **Programming Languages:** Python, Java, TypeScript
- **API Development:** FastAPI, Flask, REST APIs
- **Cloud Platforms:** AWS, Amazon S3, AWS Lambda
- **Data Platforms:** Databricks, Delta Lake, SQL
- **ML Frameworks:** scikit-learn, pandas, NumPy
- **LLM Applications:** Prompt engineering, MCP servers, intelligent agents
- **Retrieval Systems:** RAG, OpenSearch, semantic search
- **MLOps Tools:** MLflow, Weights & Biases, model registries
- **Delivery Engineering:** GitHub, GitHub Actions, CI/CD
- **Modeling Methods:** Classification, clustering, optimization
- **Forecasting Methods:** Time-series forecasting, seasonal baselines, gradient boosting
- **Model Quality:** A/B testing, automated evaluation, train-test splits
- **Observability:** Model monitoring, data drift, prediction latency
- **Responsible AI:** SHAP, LIME, bias detection
- **Software Practices:** Lean, Agile, XP, automated testing

SKILLS

- Python
- FastAPI
- Databricks
- Vector databases

- Java
- TypeScript
- SQL
- Flask
- REST APIs
- AWS
- MLflow
- Prompt engineering
- RAG
- Anomaly detection
- A/B testing

PROFESSIONAL AFFILIATIONS

- AI and Data Science Student Association, Technical Workshop Lead, 2021 - 2022
- Cincinnati Open Source Hackathon, Manufacturing Analytics Mentor, 2022
- Computer Science Peer Tutor, University of Cincinnati, 2020 - 2022

LANGUAGES

- English (Native)
- Spanish (Intermediate)

ADDITIONAL INFORMATION

Work Status : Authorized to work in United States. No sponsorship required.

REFERENCES

AVAILABLE ON REQUEST