

MOUNIKA DADI

mounika.dadi.work@gmail.com - LinkedIn: <https://www.linkedin.com/in/mounika-dadi-a43b7b1a1/>

Education: Computer Science Undergraduate at the University of Illinois at Urbana - Champaign

Product and Operations focused engineer with experience building AI-driven tools, cloud-native systems, and developer platforms. Skilled at taking ideas from 0→1 by combining strong technical execution with product sense, as shown through leading the rollout of an enterprise AI platform at scale and independently creating an **Agentic AI PM tool** with multi-stakeholder dashboards. Comfortable working across engineering, design, and leadership to deliver solutions that are technically sound and drive measurable impact.

WORK EXPERIENCE:

Head of Data Quality & Generative AI Enablement for AIOps LLM Suite at JP Morgan Chase & Co. *(August 2023 - August 2025)*

- Led cross-functional teams such as product, infra, and risk teams to align model capabilities with internal compliance, usability, and business value objectives.
- Designed and deployed agentic workflows for retrieval-augmented generation (RAG), orchestrating LLMs to automate knowledge queries across distributed datasets.
- Balanced product direction and technical execution, shaping the platform's roadmap based on user feedback, adoption data, and risk considerations.
- Built adoption momentum in a high-agency, startup-like environment, often working independently to define scope, rally stakeholders, and deliver impact.
- Led the end-to-end development and enablement of an enterprise-grade internal chatbot, integrating GPT-based models to streamline workflows and boost productivity for over 80,000 AIOps users, resulting in a 44% increase in employee productivity within JPMorgan Chase's LLM Suite.
- Fine-tuned and evaluated GPT-based language models using OpenAI infrastructure, driving significant gains in accuracy, response relevance, and reliability across high-volume use cases.
- Built and scaled Apache Airflow pipelines to fully automate regression testing across model updates, monitoring around 300 quality tests daily to ensure robust performance validation against the firm's comprehensive test suite.
- Architected backend systems to support LLM functionality at scale, including embedding store design, vector management, integration with AWS-hosted data pipelines, and helped ingest thousands of pages of internal documents plus internal Stack Overflow questions for knowledge retrieval.
- Analyzed multi-dimensional chatbot metrics (latency, engagement, resolution rates) to detect performance drift and implement precision-targeted improvements.

Software Engineering Intern at JP Morgan Chase & Co. (June 2022 - December 2022)

- Designed and tested predictive models for risk evaluation, working alongside product managers to interpret results into actionable business strategies.
- Learned quickly in a new industry, adapting technical approaches to insurance-specific data and regulatory requirements.
- Embraced a "build fast" mentality by prototyping solutions and iterating based on feedback.

Software Engineering Intern at State Farm (August 2021 - May 2022)

- Data mining and predictive modeling such as regression, forecasting, and machine learning
- Assessment of risk and financial needs for insurance products using machine learning
- Development of innovative methods and applications to enhance insurance product pricing

Software Engineering Intern at American Family Insurance (May 2021 - August 2021)

- Work on an application that brings homeowners and contractors together for easier communication for home repair assessments and quotes & programming multiple Lambda functions for various application needs in AWS
- Managing API gateway with a team of interns and maintain servers, using basic terraform
- Storing data in database with Lambda functions and using PostgreSQL

Independent Projects (linkedin)

Agentic AI Product Management Tool (2025)

- Built an agentic AI-powered product management assistant that automates key PM workflows, including task tracking, issue detection, and progress reporting.
- Designed custom dashboards for engineers, designers, and executives, each highlighting the insights most relevant to their roles.
- Integrated LLMs with project data sources to generate real-time recommendations, helping teams resolve blockers faster and align on priorities.
- Scoped, prototyped, and launched the tool independently, demonstrating 0→1 execution and strong product sense.
- Shared the project and learnings on LinkedIn, engaging with the developer and AI community around the future of AI-powered collaboration tools.

SKILLS:

Cloud & Security: AWS · GCP · Azure · Kubernetes · Docker · MLOps · Model Deployment · REST APIs · Cloud-Native Security · Containerization · Secure API Design · Cloud Cost Optimization with Grafana

Observability & Tooling: Grafana · Prometheus (familiarity) · Monitoring & Dashboards · AI/ML for Cost Analysis · Cloud Resource Optimization

Programming & Data Tools: Python · Java · SQL · Pandas · NumPy · Scikit-learn · TensorFlow · PyTorch · Linux · Node.js · JavaScript · Hadoop

Business & Strategy: AI/ML Product Adoption Strategy · Data-Driven Storytelling for Product Decisions · Cross-Functional Stakeholder Alignment · Risk & Compliance Narratives for AI/ML

AI & Machine Learning: Machine Learning · Generative AI · Large Language Models (LLMs) · Agentic AI · Natural Language Processing (NLP) · Predictive Analytics · Feature Engineering · Model Optimization · Responsible AI · Prompt Engineering · MCP (Model Context Protocol) · Governance (NIST, ISO, EU/U)

Collaboration & Leadership: Cross-functional Collaboration · Mentorship · Technical Leadership · Early-Stage Product Development · Ambiguity Management · Team Guidance · Stakeholder Alignment · Business-Technology Bridge