

Ankush K. Mulkar

Engineering Manager | AI Product Engineering | LLM & Cloud Systems

ankushmulkar@gmail.com | +91 9028621757 | linkedin.com/in/ankushmulkar | ankushmulkar.com | github.com/AnkushMulkar

PROFESSIONAL SUMMARY

Engineering Manager with 4+ years of experience across AI/ML and NLP, including more than 2 years building AI-powered products at AllyNerds. Progressed from AI Engineer to Technical Lead to Engineering Manager, leading a five-person engineering team with direct reports. Owned hiring, performance management, engineering roadmap, technical direction, and production operations while working closely with the Founder and CEO. Hands-on background in LLM applications, RAG, AI agents, Python/FastAPI, cloud infrastructure, and end-to-end product delivery.

PROFESSIONAL EXPERIENCE

Engineering Manager | AllyNerds | Feb 2026 – Present

- Led a five-person engineering team with direct reports, overseeing hiring, performance management, coaching, technical execution, and day-to-day delivery.
- Owned the engineering roadmap with the Founder and CEO, translating business priorities into technical plans, scoped releases, and execution priorities.
- Directed architecture and delivery across AI services, backend APIs, data stores, AWS infrastructure, deployment, security, monitoring, and production operations.
- Owned production reliability and incident response, using monitoring and operational feedback to prioritize fixes, improve stability, and guide release decisions.
- Evaluated LLM and infrastructure approaches across response latency, operating cost, reliability, scalability, and engineering effort to make practical architecture decisions.
- Partnered with product and growth functions on product analytics, user behaviour insights, experimentation, and initiatives targeting the US university market.
- Delivered and operated AI-powered product capabilities from technical design through deployment, monitoring, and ongoing optimization.

Technical Lead | AllyNerds | Dec 2024 – Jan 2026

- Led engineering execution across AI, backend, frontend, infrastructure, and deployment, translating product requirements into implementation plans and technical solutions.
- Designed and delivered LLM-powered capabilities and retrieval workflows using Python, FastAPI, model APIs, RAG, and vector search.
- Architected and operated AWS infrastructure for production AI services, using EC2, Lambda, S3, RDS, CloudWatch, Docker, and PostgreSQL.
- Established CI/CD, testing, monitoring, and security practices that improved release reliability and supported production delivery.
- Worked across the full product lifecycle, from technical discovery and architecture through implementation, deployment, debugging, and production support.

AI Engineer | AllyNerds | Feb 2024 – Dec 2024

- Built AI features and FastAPI backend services for an AI-powered career platform, including LLM applications, RAG pipelines, and automation workflows.
- Integrated model providers with PostgreSQL and Redis to support AI-driven application workflows and user experiences.

- Deployed containerized AI workloads on AWS and improved reliability and performance through testing, monitoring, and production debugging.
- Collaborated across product and engineering to translate user and business requirements into production-ready AI capabilities.

NLP Engineer | TechAxenic | Oct 2022 – Oct 2023

- Built NLP solutions for document intelligence, chatbots, classification, and enterprise AI assistants.
- Developed BERT-based pipelines for information extraction, semantic search, and language understanding.
- Worked across model development, preprocessing, evaluation, and integration of NLP capabilities into application workflows.

Machine Learning Intern | Bhaskar Solutions | Jan 2022 – May 2022

- Built machine learning models and preprocessing workflows for business use cases and evaluated model performance and deployment approaches.

TECHNICAL SKILLS

AI & ML: LLM Applications, RAG, AI Agents, NLP, LLM Integration, Prompt Engineering, BERT, Semantic Search

Languages & Frameworks: Python, JavaScript, TypeScript, SQL, FastAPI, Node.js, React, Next.js

Cloud & Infrastructure: AWS, EC2, Lambda, S3, RDS, CloudWatch, Docker, GitHub Actions

Data & AI Infrastructure: PostgreSQL, Redis, ChromaDB, Vector Search

Engineering: REST APIs, CI/CD, Testing, Monitoring, Security, Performance Optimization, Production Operations

CONTRIBUTIONS

AI Research: Collaborating with a small research group on observation-driven AI research, developing and proposing original observations, practical insights, and research directions aimed at contributing meaningful perspectives to the AI industry and future publications.

<https://github.com/AnkushMulkar/ai-research-lab>

Vertical Combinator: Founded a community-driven initiative focused on making students, graduates, professionals, and aspiring builders' industry-ready through practical learning, real-world projects, AI, technology, and entrepreneurship.

<https://www.verticalcombinator.com/>

EDUCATION

Master of Science, Artificial Intelligence and Machine Learning (NLP)

Liverpool John Moores University | 2023

Executive PG Program, Artificial Intelligence and Machine Learning

IIIT Bangalore | 2021 - 2023

Bachelor of Engineering, Electrical Engineering

Nagpur University | 2011 - 2014