Gayathri Santhanakrishnan Narayanan

Calthorpe Court, 280 Pershore Road, Birmingham B5 7QF • gayathrenarayanan@gmail.com • +44 7810842271
GitHub: https://github.com/GayathriNarayanan1610 LinkedIn: http://www.linkedin.com/in/gayathrisnarayanan

PROFILE

AI Engineer and technical lead with 8+ years of enterprise software engineering, including 3+ years building the production platforms and services that AI and intelligent-automation workflows run on. Currently completing an MSc in Artificial Intelligence and Machine Learning at the University of Birmingham. Designs, builds and deploys AI applications end-to-end LLM agents with tool calling and orchestration (LangGraph), retrieval systems with evaluation harnesses and guard-railed copilots shipped as secure, containerised services running live on Microsoft Azure. Strong background in enterprise integration, cloud-native architecture and leading engineering teams, with a focus on making AI reliable, governed and cost-aware in production.

AI ENGINEERING HIGHLIGHTS

- 8+ years building enterprise software platforms; 3 years delivering AI-enabled enterprise platforms.
- Architected 100+ API-led, event-driven enterprise integrations on Azure and MuleSoft; led migration of 40+ APIs to Azure Integration Services.
- Built enterprise LLM agents with LangGraph and Azure OpenAI, including human-in-the-loop approval, idempotency and full audit logging.
- Built RAG systems with reusable evaluation harnesses (golden sets, hit-rate, MRR, faithfulness) run as regression suites.
- Deployed a live GenAI application to Azure Container Apps containerised, with runtime secret management, scale-to-zero cost control and log-based monitoring.
- Delivered secure AI services with FastAPI and Docker; led engineering teams delivering cloud-native solutions.

TECHNICAL SKILLS

AI Engineering: LLM applications & copilots, agentic workflows (tool use, orchestration, routing, context & memory/state management), LangGraph, LangChain, reusable agent tools incl. MCP (Model Context Protocol), Pydantic schema validation, prompt engineering, structured outputs, function calling, human-in-the-loop

Retrieval & Evaluation: Retrieval-Augmented Generation (RAG), vector search, embeddings, semantic search, grounded answering with source attribution, hallucination reduction & refusal, AI evaluation & observability (golden sets, hit-rate, MRR, faithfulness), regression testing; knowledge graphs & GraphRAG, ontology design, information extraction

Programming & Software Engineering: Python (Pandas, NumPy, scikit-learn), FastAPI, REST APIs, microservices, SQL, SQLite, Git; ML fundamentals model evaluation, feature engineering, transformers/LLM concepts

Cloud, Delivery & Security: Microsoft Azure (Azure OpenAI, Functions, Logic Apps, Service Bus, DevOps), CI/CD (Azure DevOps, Jenkins, Maven), Docker/containers, event-driven architecture, secrets & access controls, responsible-AI guardrails, MuleSoft/enterprise-system integration

Consulting & Delivery: Translating ambiguous requirements into technical plans, stakeholder management, technical documentation, mentoring & technical leadership

AI & MACHINE LEARNING PROJECTS:

Knowledge Graph Evidence-Mapping System - GraphRAG over technical documents (MSc, industry-partnered)

- Building an evidence-mapping system for an industry-partnered MSc project (deep-tech manufacturing) that turns scattered technical documentation into a queryable knowledge graph, linking engineering claims to their supporting evidence, tests, components, risks and versions and surfacing what is proven, assumed, or missing.
- Designing the ontology (claims, metrics, assumptions, evidence, tests, suppliers, risks, validation status) and the document extraction-and-linking pipeline, combining knowledge-graph retrieval with LLM extraction (GraphRAG) so answers stay grounded and source-traceable with machine-extracted links kept separate from human-verified facts.

Agentic Integration Service - LLM agent over enterprise systems (LangGraph + Azure Open AI)

- Built an agentic AI service where an LLM plans and executes multi-step workflows across enterprise systems via tool/function calling (CRM lookup, data enrichment, ticket creation), orchestrated with LangGraph on Azure OpenAI.
- Engineered the controls that make an acting agent safe: human-in-the-loop approval for high-risk actions, idempotency to prevent duplicate side-effects, context management and a structured audit log of every decision and tool call.

RAG Evaluation Service - retrieval pipeline with an evaluation harness

- Built a retrieval-augmented pipeline, chunking, embedding, grounded answering with citations and refusal when retrieved context can't support an answer.
- Built a reusable evaluation harness scoring retrieval and answer quality against a golden test set (hit-rate, MRR, context precision, faithfulness, answer relevance, refusal rate), run as a regression suite so a prompt or model change that breaks behaviour surfaces immediately.

AI Data Copilot - LLM-Powered SQL Assistant (live demo)
Live: copilot.jollyglacier-1f65987e.uksouth.azurecontainerapps.io

- Built a copilot that lets non-technical users query a multi-table SQL database in plain English generating SQL from the live schema, returning results as tables and refusing rather than guessing when a question can't be answered from the data.
- Engineered a deterministic guardrail that validates every generated query and runs only read-only SELECTs, blocking destructive statements; added a self-correction loop that feeds execution errors back to the model to repair failed SQL.
- Delivered as a FastAPI service and Streamlit UI with a deterministic test mode so the full pipeline is validated for reliability without a live model.
- Deployed and running on Azure Container Apps: containerised with Docker, image served from Azure Container Registry, API key injected at runtime from a platform secret (never baked into the image), scale-to-zero for cost control and log-based monitoring.

MCP Tool Connector - reusable tools for AI agents

- Built an MCP (Model Context Protocol) server exposing reusable, authenticated tools (document search, company lookup) to any MCP-compatible agent, with Pydantic-validated inputs and outputs enforcing reliable data contracts.

WORK EXPERIENCE

Accenture Solutions Private Limited

Application Development Team Lead **Dec 2023 - Aug 2025 | Chennai, India**

- Built the platform layer for enterprise AI reusable Python APIs and secure backend services enabling AI workflows, orchestration and data access across multiple enterprise systems.
- Led the migration of 40+ APIs to Azure Integration Services, replacing ageing integration layers that were constraining scale improving performance and reliability across connected systems.
- Shipped production cloud services with authentication, monitoring, logging, CI/CD and governance controls, meeting the security and audit requirements of regulated clients in finance, healthcare and utilities.
- Partnered with solution architects and business stakeholders to identify AI-automation opportunities and turn ambiguous requirements into technical designs that shipped.
- Led and mentored a team of 5 engineers, setting engineering standards through code reviews and knowledge sharing.

Senior Software Engineer **Dec 2020 - Nov 2023 | Chennai, India**

- From 2022, built the secure backend services and data pipelines underpinning intelligent-automation initiatives the platform layers those workflows ran on.
- Architected API-led integrations across Azure and MuleSoft (100+ over tenure), connecting enterprise systems that previously depended on manual data transfer.
- Designed event-driven architectures on Azure (Service Bus, Logic Apps, Functions), replacing batch processes with real-time, high-throughput data flows.
- Built CI/CD pipelines (Azure DevOps, Jenkins, Maven) with automated testing, code review and containerised deployment; ran performance benchmarking and load testing to keep services available under peak load.

Application Development Analyst **Jun 2018 - Nov 2020 | Bengaluru, India**

- Developed API-led integrations (REST and SOAP) on the MuleSoft Anypoint Platform for a global insurance client (QBE), connecting enterprise systems including Salesforce, PeopleSoft and mainframe applications via JMS/Active MQ messaging; supported and debugged production integrations affecting live business data flows.

Application Development Associate **Nov 2016 - May 2018 | Bengaluru, India**

- Built and tested enterprise integration components on MuleSoft and carried out major version migrations (Mule 3.8 through 4.4) replacing deprecated components working to client delivery standards.

EDUCATION

University of Birmingham Birmingham, UK
Masters in Artificial Intelligence and Machine Learning 2025 - 2026
Panimalar Engineering College - Anna University Chennai, India
Bachelor of Engineering in Electrical and Electronics Engineering 2012 - 2016

AWARDS AND ACHIEVEMENT:

- **Masters Campus-Based Internship - University of Birmingham with Severn Trent**
  Worked with multi-disciplinary team and helped bridge data science and environmental engineering by providing insights that informed the sizing and performance evaluation of the proposed nature-based solution.
- 2nd Place in Healthcare AI Hackathon: built an AI-driven solution under time constraint, demonstrating rapid prototyping and problem-solving
- Accenture Celebrates Excellence Award (x2) and Accenture Star Award