

Elambharathi Thangavel

📍 Bengaluru, India ✉ elambharathi.pt@gmail.com 📞 +91 9677407190 🌐 github.com/elampt

in linkedin.com/in/elambharathithangavel 🔗 elambharathi.com

Career Objective

Python Engineer building production-grade AI agents and LLM applications with LangGraph, RAG, and MCP. Built an autonomous Jira-to-PR coding agent with self-healing test loops, Playwright visual verification via MCP, and human-in-the-loop approval – with full LangFuse observability. Also built an agentic RAG chatbot with two-stage retrieval and config-driven architecture. Won PhonePe's AI Immersion Hackathon. Backed by production experience – automation pipelines on HDFS, on-call across 6 microservices, and FastAPI backends with PostgreSQL.

Education

Bachelor's Amrita Vishwa Vidyapeetham University, Coimbatore, India

08/2020 – 05/2024

- B.Tech in Computer Science and Engineering (**CGPA: 7.29**)

Skills

AI/GenAI: LangChain, LangGraph, RAG, FAISS, Agentic AI, Sentence Transformers, Cross-Encoder Reranking, Prompt Engineering, LLM APIs (Groq, OpenAI)

Languages: Python, SQL, Shell Scripting

Backend: FastAPI, REST API Design, Async Python, PostgreSQL, Redis, Pydantic, JWT Authentication, Rate Limiting

Data & Infra: PySpark, Apache Hive, HDFS, Airflow, Cron, Docker, Git, Postman

Work Experience

PhonePe [🔗](#), Product Solution Engineer

Bengaluru

India's leading digital payments platform – Quick-Commerce division (Pincode)

07/2025 – Present

- **Won the PhonePe “AI Immersion Hackathon”** – built a complete end-to-end appointment booking system (merchant-facing: salons, spas, etc.) in 5 days using AI-assisted development. Led architecture design, API contract definition, and backend integration; judged on production-readiness and code quality.
- Built an **end-to-end automation pipeline** (Python + PySpark + Hive) replacing a paid third-party platform – extracts merchant data from HDFS, applies business logic, and creates tasks via internal REST APIs. Runs **twice daily via cron** with 99%+ uptime and zero manual intervention.
- Provided **L1/L2 on-call support across 6 distributed microservices** (Catalogue, Transaction, Fulfillment, Search, Buyer/Seller Apps). Achieved best-in-team metrics: **MTTA < 1hr, MTTR < 13hrs** through systematic log analysis and cross-service debugging.
- Served as **primary technical POC for merchant API integrations** (Order, Inventory, Catalog APIs) – led end-to-end lifecycle from consultation to production sign-off with **100% success rate**. Debugged API request/response flows using Postman and log analysis tools.
- Built an **automated reporting pipeline**: Google Sheets API → HTML templating → PDF generation → SMTP email delivery, eliminating daily manual reporting work.

Jobshine [🔗](#), Python Developer | Chatbot Engineer

Singapore (Hybrid)

Singapore-based startup – blue-collar recruitment automation

05/2024 – 05/2025

- Built and maintained **AI chatbots across WhatsApp, Messenger, and Telegram** using Python and Make (Integromat); automated candidate engagement workflows and job application pipelines for messaging platforms.
- Developed **in-house Python web crawlers** (Selenium, BeautifulSoup, Requests) replacing a paid scraping tool, saving **100%** in tool costs and reducing manual data collection by **70%**.
- Performed **EDA on engagement data** using Pandas to improve lead quality and recruitment conversion outcomes.

Projects

Jira Coding Agent – AI-Powered Code Automation (04/2026) [Project link](#) [🔗](#)

- Built an **autonomous AI coding agent** using LangGraph that watches Jira for tickets, modifies a React codebase via RAG-powered code search (FAISS + grep), runs tests with a **self-healing retry loop**, and creates PRs with **before/after screenshots via Playwright MCP**.

- Includes human-in-the-loop approval for high-risk changes and full LLM observability via LangFuse.

Bubble – Production-Grade RAG Chatbot (04/2025 - Present) [Project link](#)


- Built an **agentic RAG chatbot** using LangChain, LangGraph, and FAISS with a **two-stage retrieval pipeline**: bi-encoder (FAISS) for fast recall → cross-encoder reranking for precision, with configurable relevance thresholds.
- Designed **config-driven architecture** with Pydantic validation and YAML config – switching LLM providers, embedding models, or retrieval parameters requires zero code changes. Improved chunking from 322 noisy fragments to 82 high-quality chunks.


FinSight – Stock Portfolio Tracker & Sentiment Analyzer (05/2025) [Project link](#)

- Full-stack application with **FastAPI backend** (async endpoints, JWT auth, dependency injection), **PostgreSQL** database (Neon), real-time P&L tracking via yFinance, and news sentiment analysis using FinBERT. Deployed with Streamlit UI.

Certificates

Unsupervised Learning, Recommenders, Reinforcement Learning (06/2025) [DeepLearning.AI](#) | [Coursera](#) 

Advanced Learning Algorithms (05/2025) [DeepLearning.AI](#) | [Coursera](#) 

GenAI Intensive Course (04/2025) [Google](#) | [Kaggle](#) 

Supervised Machine Learning: Regression and Classification (02/2023) [DeepLearning.AI](#) | [Coursera](#) 

Publications

A Multi-Round Zero Knowledge Proof Algorithm for Secure IoT and Blockchain Environments [DOI](#) 

Published on International Journal of Safety and Security Engineering