

# Meghana Ganapa

meghanaganapa1@gmail.com | +61431541467 | [linkedin.com/in/meghana-ganapa](https://www.linkedin.com/in/meghana-ganapa) | [github.com/meghanaganapa](https://github.com/meghanaganapa) | [meghanaganapa.github.io/meghana-portfolio](https://meghanaganapa.github.io/meghana-portfolio)  
| Available to work immediately | Open to relocation |

## Summary

---

Data Science graduate currently working as an Agentic AI Engineer, with hands-on experience building GenAI applications, AI agents, RAG systems, LLM evaluation workflows, automation pipelines, dashboards, and machine learning solutions. Skilled in Python, SQL, AWS, APIs, Power BI, LLMs, vector search, and workflow automation, with a strong focus on practical business decision support. Experienced in LLM-as-a-judge evaluation, model benchmarking, prompt engineering, and cloud-based AI deployment. Curious, passionate, and a fast learner, with a strong interest in building intelligent AI solutions that solve real-world problems and create measurable impact.

## Education

---

### The University of Melbourne

Feb 2023 – Dec 2024

Masters in Data Science

- Data Science Peer-to-Peer Mentor - provided academic guidance and support on course-specific challenges.

## Relevant Experience

---

### Agentic AI Engineer

May 2026 – Present

Enspyr – Melbourne, Victoria

- Built and tested model-as-a-judge evaluation workflows using Claude, OpenAI, and Gemini APIs to assess model agreement, reasoning quality, escalation decisions, latency, and cost-performance tradeoffs across coding and reasoning tasks.
- Experimented with HumanEval and BBH benchmarking harnesses, resolved bugs in the HumanEval evaluation dataset workflow, and improved JSONL logging, parser/scoring validation, failure diagnostics, and reproducible LLM benchmarking pipelines.

### Research Intern

Feb 2023 – May 2023

Melbourne Data Analytics Platform – Melbourne, Victoria

- Conducted comparative analysis of ImageJ and Pixel Annotation Tool to support biomedical image annotation, data validation, and CNN-based model development for gonadal tissue analysis.
- Implemented and evaluated a Python-based CNN image classification workflow, improving tissue classification accuracy by 15% through preprocessing, model testing, and performance evaluation.

### Associate Software Engineer

Sep 2022 – Dec 2022

Qentelli Solutions Private Limited – Hyderabad, Telangana

- Optimised SQL queries and automated REST API workflows, reducing manual processing by 30% and improving data retrieval efficiency for production systems.
- Collaborated with cross-functional teams to translate business requirements into scalable backend components, API integrations, and data-driven automation solutions.

### Software Engineer Intern

Feb 2022 – Aug 2022

Accolite Digital – Hyderabad, Telangana

- Developed and deployed backend modules using Spring Boot, MongoDB, and Model-View-Controller architecture to improve application scalability, maintainability, and system reliability.
- Contributed to API development, database integration, debugging, and software testing workflows, supporting a 30% increase in development velocity and a 25% improvement in system reliability.

## Projects

---

### GridFlex AI: Renewable Energy Optimisation Agent

Jun 2026

*Amazon Bedrock, AWS Lambda, DynamoDB, EventBridge Scheduler, Python*

- Built an AWS-based simulated real-time GenAI system to optimise renewable energy usage by recommending EV charging, flexible appliance scheduling, and home battery usage during clean-energy surplus periods.
- Developed Lambda-based telemetry and decision workflows to generate simulated household energy data, detect solar surplus and price signals, and store traceable recommendations in DynamoDB.
- Integrated Amazon Bedrock and EventBridge Scheduler to automate explainable energy recommendations, demonstrating a real-time-ready architecture for renewable energy load shifting and future IoT integration.

### CareerPilot: AI-Powered Career Intelligence Platform

Apr 2026

*Python, FastAPI, PostgreSQL, Playwright, AWS, REST APIs*

- Built an AI-powered workflow platform to automate job data collection, screening, ranking, and decision support using Python, FastAPI, Playwright, PostgreSQL, REST APIs, and LLM-based evaluation.

- Developed automated ingestion pipelines with scheduling, deduplication, structured logging, monitoring, and failure-handling mechanisms to improve reliability, scalability, and operational stability.
- Optimised ranking and data processing workflows to improve role relevance, reduce manual job search and screening effort by approximately 30%, and support faster career decision-making.

### Retrieval-Augmented Generation (RAG) System for Medical Textbooks

Apr 2025 – Jun 2025

*Python, SentenceTransformers, Vector Databases, CrossEncoder, LLMs*

- Engineered a Retrieval-Augmented Generation pipeline for AI-powered question answering over large-scale medical textbooks using document chunking, embeddings, vector search, reranking, and LLM response generation.
- Implemented SentenceTransformers-based embedding generation, vector similarity search, and CrossEncoder reranking to improve retrieval accuracy, answer relevance, and grounded response quality.
- Designed evaluation workflows, retrieval benchmarks, and quality testing pipelines to reduce hallucinations, improve trust in generated answers, and reduce manual document search time by approximately 40%.

### Fraud Claim Detection Agent

May 2026

*Python, Machine Learning, Docker, AWS Lambda, Amazon ECR, Amazon Bedrock Agent*

- Built an end-to-end fraud claim detection agent to screen insurance claims and return fraud prediction, probability score, risk level, warning flags, and manual review recommendations.
- Trained a majority-voting ensemble model using Logistic Regression, XGBoost, LightGBM, and CatBoost, with rule-based explainability flags to support transparent fraud risk assessment.
- Packaged the trained model into a Docker container, deployed it through AWS Lambda using Amazon ECR, and integrated it with Amazon Bedrock Agent for natural-language claim assessment and consistent triage support.

### AI-Driven Polyp Extraction and Colonoscopy Recommendation System

Jan 2024 – Dec 2024

*Python, Machine Learning, Ollama, Random Forest, Flask, PDF Processing*

- Experimented with multiple machine learning and NLP approaches to extract structured polyp data from unstructured histology reports, including size, histology, dysplasia, and clinical risk features.
- Implemented an Ollama-based extraction workflow to convert histology report content into structured JSON outputs, improving usability of unstructured clinical text for downstream clinical decision support.
- Built a Random Forest colonoscopy recommendation system using extracted polyp data and clinical guideline logic, achieving approximately 83% accuracy, and deployed it through a Flask app to reduce manual PDF review effort by approximately 30%.

## Technical Skills

---

**Programming Languages:** Python, SQL, R, Java, TypeScript

**AI, GenAI & Automation:** Large Language Models (LLMs), AI Agents, Retrieval-Augmented Generation (RAG), Prompt Engineering, LLM Evaluation, LLM-as-a-Judge, Vector Search, Workflow Automation, OpenAI APIs, Amazon Bedrock

**Machine Learning & Data Science:** Supervised Learning, Classification, Regression, Random Forest, XGBoost, Model Evaluation, Feature Engineering, NLP, Time-Series Analysis

**Data Engineering & Analytics Engineering:** ETL/ELT Pipelines, Data Modelling, Star Schema, Data Validation, Data Quality Automation, Data Governance, Workflow Optimisation, CI/CD Concepts

**Cloud & DevOps:** AWS (Bedrock, Lambda, S3, Glue, Redshift, EC2, DynamoDB, EventBridge, ECR), Azure, Docker, GitHub Actions, Apache Airflow, Containerisation

**Business Intelligence & Analytics:** Power BI, DAX, Power Query, Tableau, Dashboard Development, Data Analysis, Reporting, Semantic Data Modelling

**Backend, APIs & Tools:** Flask, FastAPI, REST APIs, API Integration, Git, GitHub, Jupyter Notebook, VS Code

**Databases:** PostgreSQL, MySQL, MongoDB, Neo4j, Qdrant, Vector Databases, Relational Databases

**Soft Skills:** Stakeholder Communication, Problem Solving, Cross-functional Collaboration, Technical Documentation

## Achievements

---

- Won the Pitching Challenge at Watt the Hack for *WattWise*, an AI-powered mobile app helping buyers and sellers of older homes identify energy-efficiency upgrades aligned with NatHERS 7-star standards; awarded \$10,000 USD in Codex credits and 1 year of Codex Pro.
- Published two peer-reviewed research papers in Springer on machine learning applications in healthcare analytics: *"Cognitive Baby Care Solutions for Smart Parenting"* and *"Effective Prediction of Brain Tumor Using Machine Learning Algorithms"*.
- Participated in Smart India Hackathon-2018, I-Innovate, and IBM hackathons, showcasing innovative solutions to technical challenges.

## Certifications

---

- AWS Certified Cloud Practitioner
- AWS Certified Generative AI Developer – In Progress