

SUMMARY

Ex-Chartered Accountant turned GenAI Data Scientist shipping production AI at Commonwealth Bank, from a customer-facing disputes agent reaching 7M app customers to streaming guardrails that cut inference latency approx. 37x. Monash Master of Data Science, three accelerated promotions at CBA in two years, with depth across agents, guardrails, fraud, scams, risk, and full-stack AI products.

EXPERIENCE

Commonwealth Bank of Australia

Melbourne, VIC
2023 - Present

- Data Scientist, GenAI

DISPUTES AGENT CHATBOT

- Built CBA's first customer-facing GenAI disputes agent end-to-end, now serving 7M app customers and 200+ disputes per day, and a finalist in the 2025 CBA Excellence Awards.
- Cut dispute rework by 12% by designing a structured agent workflow that asked more precise questions upfront and captured downstream-ready information.
- Hit 96% intent-classification, 94% information-extraction, and 88% dispute-description evaluation accuracy across the agent's key components.
- Generalised the agent into a reusable dialogue template now used by 3 other teams, with a no-code layer letting non-technical staff express business logic without engineering effort.
- Owned the underlying LangGraph agent logic, DAG dialogue modelling, BFS traversal, backend APIs, database design, evaluation, and risk/governance controls.

FRAUDS AND SCAMS AGENTIC WORKFLOW

- Built a production frauds-and-scams agentic workflow used by 200+ fraud analysts, synthesising customer, crypto, fraud-check, and ML signals into structured reports and recommended next actions.
- Cut analyst review effort by approx. 30 minutes per case by automating signal synthesis, case summarisation, and recommendation generation.
- Took the system from initial steel thread (ADK) to production, owning agent logic design and translating analyst workflows into a deployed GenAI tool with fraud stakeholders.

STREAMING TEXT-BASED GUARDRAILS

- Cut guardrail inference latency approx. 37x, from approx. 1,800ms (LLM-based evaluation) to 49ms, by fine-tuning a purpose-built model, unblocking risk/governance approval for voice-agent productionisation.
- Fine-tuned open-source foundation models with custom architecture and loss functions, reaching 0.88 F1 / 0.997 AUROC at final-sequence level and 0.77 F1 at streaming-policy level.
- Led the applied research and pilot productionisation for streaming financial-advice guardrails end-to-end.

Pitcher Partners

Melbourne, VIC

- Audit, Taxation and Advisory Manager

Mar 2016 - Mar 2021

- Managed audit, taxation, and advisory engagements across an approx. AUD 450K p.a. book, earning 4 accelerated promotions in five years.
- Built a bank-transaction classification algorithm for forecasting and analysis, saving 40+ hours of manual work.

PROJECTS

Culina

- AI-Powered Nutrition and Macro Tracking App - useculina.com Production
- Designed, built, and shipped a live AI nutrition app solo end-to-end (React/Vite, FastAPI, PostgreSQL, Redis), featuring LLM-based food logging from text and photos and semantic search via pgvector nearest-neighbour retrieval on image-derived embeddings.

EDUCATION**Monash University**Melbourne, VIC
Mar 2021 - Dec 2022

- Master of Data Science, Minor Thesis
- **Dux of Master of Data Science cohort** GPA: 4.00 / 4.00; Dean's Achievement List; Research candidate; Industry Director, WIRED Society.

Chartered Accountants Australia and New Zealand

- Graduate Diploma of Chartered Accounting Mar 2017 - Dec 2019
- Candidate of the Module, Victoria - Taxation, 2018.

CERTIFICATIONS

- AWS Certified AI Practitioner.
- Deep Learning Specialization.

SKILLS

GenAI & Agents: LLMs, agentic workflows, LangGraph, ADK, PydanticAI, CrewAI, LiveKit, prompt engineering, LLM evaluation, guardrails, synthetic data, fine-tuning.

Machine Learning: NLP, classification, information extraction, time-series forecasting, fraud/risk modelling, embeddings, pgvector, nearest-neighbour search.

Backend & Data: Python, SQL, R, FastAPI, PostgreSQL, Redis, SQLAlchemy, APIs, database design, ERDs, indexing, SQL optimisation, rate limiting.

Infrastructure: Docker, Linux, Vercel, Railway, Linode, Logfire, AWS, production deployment.

Domains: Banking, GenAI risk, fraud, scams, disputes, financial crime, audit, tax, advisory, expected credit loss.