

AI & GenAI Engineering Experience

A curated summary of hands-on AI/GenAI/ML work — from deploying and evaluating LLM API endpoints to building GenAI workflows in production and advising on ML/MLOps practice — plus personal practice building with Claude as an agentic development partner. Ordered most recent first.

PROFESSIONAL ENGAGEMENTS

Deloitte consulting engagements with a direct AI/ML/GenAI component. Client names withheld per standard consulting confidentiality practice.

Australian State Government Insurance Company — Data Analyst / Engineer — GenAI-Powered Legacy Migration Mar 2025 – Jun 2025

A large-scale migration of on-premise data into a new Snowflake Enterprise Data Warehouse required converting a huge volume of legacy SAS scripts into SQL — a slow, error-prone task by hand.

- Built a GenAI workflow using Anthropic's Claude 3.7 to convert legacy SAS scripts into Snowflake SQL at scale
- Owned PII-masking across hundreds of SAS files, designing a Python Regex-based solution to identify and mask/map sensitive data
- Supported the ETL team writing SQL stored procedures across 3 staging layers and 1 target layer to migrate data into production

Impact — Delivered productionised code in 50% less time than planned and a 50% reduction in resourcing cost — efficient enough that Bill was moved onto the ETL stream to support further activities off the back of it.

Claude 3.7 · Python (Regex) · Snowflake SQL · SAS

Big 4 Australian Bank — Business Analyst — Generative AI Risk & Compliance MVP Aug 2024 – Sep 2024

16,000+ compliance incidents are reported annually at this bank; 75% require a compliance assessment taking 2.5–5 hours to complete by hand. The client needed a GenAI copilot to cut that time down.

- Gathered requirements and translated them into delivery, bridging gaps between data scientists, solution architects, and Risk & Compliance business stakeholders
- Ran end-to-end testing and drove sign-off on production-gating artefacts (Test Plan, Test Approach, Detailed Technical Design, Test Completion Report)
- Ran daily standups and scrum ceremonies — backlog grooming, RAID tracking — to keep the build on schedule

Impact — Shipped the first of 5 pioneer GenAI use cases piloted into production at the bank, establishing a repeatable GenAI solution pattern the Data Insights function could reuse for future use cases.

GenAI Copilot · Confluence / Jira · Agile / Scrum

Big 4 Australian Bank — Gen AI Engineer — Policy Simplification POC Jul 2024 – Aug 2024

The bank needed to summarise long, dense policy documents without losing their original intent or meaning — a task where getting the fidelity/compression trade-off wrong is costly.

- Deployed LLM API endpoints in Azure ML Studio to summarise lengthy policy documents
- Built a chunking pipeline to break large PDF documents into inputs the endpoint could consume
- Designed evaluation metrics and acceptance criteria to benchmark GPT-4o, GPT-4, GPT-3.5, and Llama 3 on summarisation quality and length reduction

Impact — Delivered a working POC alongside a model-agnostic evaluation framework the bank could reuse to benchmark future LLMs on document-compression tasks.

Azure ML Studio · GPT-4o / GPT-4 / GPT-3.5 · Llama 3 · Python

The client's Data Science team had built predictive ML models but needed help improving, communicating, and deriving ongoing value from them.

- Evaluated the client's existing predictive ML modelling and MLOps processes
- Outlined recommendations to streamline modelling and MLOps workflows in Python
- Delivered a roadmap for more effective stakeholder communication and migration away from legacy systems

Impact — Delivered a technical report and stakeholder presentation that gave the client's data science function a concrete roadmap for maturing its ML practice.

Python · MLOps Advisory · Predictive Modelling

PERSONAL AI PRACTICE

Separate from client delivery — how Bill builds with AI on his own time.

Agentic Development with Claude Code

Ongoing

Bill uses Claude Code as a daily agentic pair-programmer across every personal project he ships — not as an autocomplete tool, but as an agent that plans, edits, runs builds/tests, and drives git/CI end-to-end. This site is itself a live example: it was designed, built, content-audited, and deployed through extended agentic sessions with Claude Code — multi-step planning, tool-orchestrated edits across a whole codebase, automated verification (lint, build, Playwright-driven visual/responsiveness audits), and GitHub Actions deploys, all driven conversationally. The same practice is behind RiftCompare (riftcompare.com), a live Riftbound TCG price-comparison platform with 10,000+ monthly visits — Claude Code has been a hands-on development partner across its scraping pipeline, UI, and deployment, with Bill directing the architecture and reviewing every change.

- Comfortable directing multi-step agentic workflows: planning, execution, and verification loops rather than single-shot prompts
- Hands-on with tool-use/orchestration patterns — file edits, shell commands, browser automation, CI/CD — chained together to ship real, deployed product
- RiftCompare (riftcompare.com) — a real, live product with real traffic, not a demo — built and iterated on with Claude Code as a development partner

Claude Code · Agentic workflows · CI/CD · Playwright