Curtis Murray, PhD

+61 413 453 827 | curtis@curtismurray.tech

Available immediately

Profile

AI engineer shipping production LLM systems — end-to-end agentic workflows with custom MCP servers, multi-step tool orchestration, and RAG pipelines grounded in verified entities to reduce hallucination. Built neuro-symbolic tooling that grounds LLM reasoning in Neo4j knowledge graphs, and a clinical entity-linking RAG pipeline that outperformed ClinicalBERT and direct LLM generation by constraining outputs to real ontology candidates. Daily Claude Code user. PhD in applied mathematics (NLP focus) with Q1 publications. Strong foundations in Bayesian inference and statistics. Looking to join a team solving real problems with practical AI engineering.

Technical Skills

Agentic AI & LLMs: Custom MCP server development, agent orchestration and multi-agent workflows, RAG architectures, knowledge graph integration (Neo4j), Claude Code, AWS Bedrock, query reformulation, LLM-based entity extraction.

NLP & ML: Transformer fine-tuning (BERT variants), HuggingFace, PyTorch, TensorFlow, topic modelling (hSBM, LDA), text classification.

Programming: Python (5+ years), R, MATLAB, Git, Bash, SQL.

Infrastructure & Cloud: Docker & docker-compose, REST APIs, AWS (Bedrock), GCP, Slurm HPC.

Statistics & Visualisation: Bayesian inference, probabilistic generative models, uncertainty quantification, network analysis, R (tidyverse, ggplot2, plotly, Keras).

Recent Projects

Spoken Tamil Morphology Analyser – Neuro-Symbolic Agentic AI 2025 – Present

- Built custom MCP server and agent workflows that ground LLM outputs to a ∼1,000-node Neo4j knowledge graph encoding spoken Tamil grammar, morphemes, and constructions.
- Designed neuro-symbolic RAG loop: LLM proposes a segmentation hypothesis → graph queries validate against defined entities → LLM refines on structured feedback → iterate until grounded. Context maintained across the trajectory, with IR and tool-calls tuned to maximise relevant retrieval and minimise output token lengths through model-adjudicated summarisation. Reduces hallucination by constraining outputs to verified graph entities.
- Multi-agent orchestration for graph construction: agents proposed entries from OCR'd reference material with deduplication and adjudication for edge cases, producing the grammar graph itself.
- First working computational morphological analyser for spoken Tamil — literary Tamil has established analysers, spoken has had none — in a setting with no canonical rules to unpack surface forms, where direct LLM generation is unreliable.

Tamil Language Learning Platform – Full-Stack EdTech 2025 – Present

- Shipping EdTech platform for Tamil diaspora learners: click-to-analyse words with morphological breakdown via the MCP server above, spaced repetition, and learner progression.
- Stack: React, Express, PostgreSQL, Neo4j. Pre-launch.

Experience

Clinical NLP Engineer & Research Fellow Feb 2024 – July 2025

RMIT University · Honorary appointment, The University of Melbourne *Melbourne, AU*

- Built RAG pipeline for clinical entity linking: embedded ontology entities (DrugBank, RxNorm) into vector space, retrieved candidates via similarity search, then LLM-adjudicated. Outperformed string matching, ClinicalBERT, and direct LLM generation on the target task — direct LLM generation was unreliable due to hallucinated ontology codes, which the retrieval-based adjudication pattern prevents by constraining outputs to real candidates.

- Implemented LLM-driven query reformulation: when initial retrieval returned poor candidates, the LLM rewrote the query (e.g., brand names to generic drug names) and retried, measurably improving hit rates on hard cases.
- Fine-tuned transformer models (ClinicalBERT, VetBERT) for document classification on HPC clusters with Slurm job arrays.
- Worked in a regulated clinical data environment with governance and privacy constraints; presented outputs to clinical stakeholders at MCBK 2024.

PhD Researcher – Applied Mathematics & NLP Feb 2020 – May 2025
University of Adelaide *Adelaide, AU*

- Thesis: *NLP Reveals Patient Reported Experiences* — built Bayesian network-based topic modelling pipelines (hierarchical stochastic block models) for analysing patient health narratives at scale.
- Extracted early COVID-19 symptom signals from social media, automatically detecting then-novel symptoms (anosmia, ageusia) before widespread clinical recognition.
- Published in *Artificial Intelligence in Medicine* (IF 6.2, Q1) and at ICWSM (AAAI CORE Rank A).
- Deployed analyses on university HPC (Slurm) with Bash automation and parallel job arrays.

Solutions Engineer Nov 2022 – Jun 2023
Energy Exemplar *Adelaide, AU*

- Pre-sales technical role building PLEXOS (Gurobi backend) optimisation models to demonstrate product fit for prospective customers. Led a green/blue hydrogen model that closed a $600k software licence sale to a government customer.
- Creative mathematical reframes of hard customer problems — e.g. solving a daily constraint problem as graph shortest-path. Presented models directly to enterprise and government stakeholders.

Research Assistant – Health Surveillance NLP Jan 2021 – Mar 2022
University of Adelaide *Adelaide, AU*

- Applied topic-model time-series analysis to adverse device-report narratives for early signal detection.

Mathematics Tutor & Teaching Assistant 2019 – 2022
University of Adelaide *Adelaide, AU*

- Walk-in help-room tutor and teaching assistant across undergraduate mathematics courses; explained technical concepts to diverse audiences, provided structured feedback on student work.

Education

PhD, Applied Mathematics 2020 – 2025
University of Adelaide

Bachelor of Mathematical Sciences (Honours), Statistics 2019
University of Adelaide *First Class Honours · University Medal · GPA: 7.0*
Thesis: Uncertainty in Deep Learning

Bachelor of Mathematical Sciences, Pure Mathematics 2015 – 2018
University of Adelaide

Selected Publications

Murray, C., Mitchell, L., Tuke, J., & Mackay, M. (2025). Probabilistic emotion and sentiment modelling of patient-reported experiences. *Artificial Intelligence in Medicine* (IF 6.2, Q1), 139, 103178.

Murray, C., Mitchell, L., Tuke, J., & Mackay, M. (2024). Revealing patient-reported experiences in healthcare from social media using the DAPMAV framework. *Digital Health*, 10.

Murray, C., Mitchell, L., Tuke, S., & Mackay, M. (2021). Symptom extraction from the narratives of personal experiences with COVID-19 on Reddit. *Proc. ICWSM* (AAAI CORE Rank A).

Honours & Awards

University Medal, University of Adelaide (2019)

Queensland AI Hub Health Hackathon – Winning team ($20K AUD) (2020)

Dean's Commendation for Academic Excellence (2018)

References available on request.