

Farid Ariai

PhD Candidate in Data Science
Brisbane, QLD, Australia

+61 466 770 657
f.ariai@uq.edu.au
farid-ariai.github.io
linkedin.com/in/farid-ariai
github.com/farid-ariai
Google Scholar

PROFILE

PhD candidate in Data Science at The University of Queensland, specialising in Legal NLP, LLM evaluation and efficient inference. Research spans legal retrieval-augmented generation, parameter-efficient fine-tuning, quantisation and human–AI interaction. Combines controlled experimentation with research software development, GPU/HPC workflows and human-in-the-loop evaluation.

EDUCATION

PhD in Data Science 2024 – Present

The University of Queensland

Supervisors: Gianluca Demartini and Joel Mackenzie, School of Electrical Engineering and Computer Science.

Master of Software Engineering 2021 – 2023

University of Tehran

Ranked first in cohort.

Bachelor of Software Engineering 2012 – 2017

Ferdowsi University of Mashhad

RESEARCH EXPERIENCE

PhD Researcher — Legal NLP and LLM Evaluation April 2024 – Present

The University of Queensland, Brisbane, Australia

- Conduct research on legal retrieval, parameter-efficient model adaptation and post-training quantisation, with an emphasis on controlled evaluation and reproducible experiments.
- Develop research software spanning dataset preparation, model training and inference, statistical evaluation and GPU/HPC execution.
- Co-authored a survey of Legal NLP published in *ACM Computing Surveys*.

Research Collaborator — Digital Deliberative Democracy (D3) April 2024 – Present

The University of Queensland / University of Zurich

- Contribute to research on NLP and human–AI interaction in civic participation through the D3 project, supported by SNSF Grant 205975.

SELECTED RESEARCH PROJECTS

Legal RAG: Retrieval Cues and Graph Structure

- Co-developed a Legal RAG pipeline combining LLM-based information extraction, Neo4j, dense/BM25 retrieval, reranking and controlled tests of retrieval cues versus graph structure.
- Contributed reproducible retrieval and evaluation infrastructure covering a 772,001-passage legal corpus and 4,584 HousingQA questions across 47 jurisdictions.

MJL-NER: Legal Entity Supervision in LLMs

- Co-developed controlled experiments on legal-entity supervision using QLoRA fine-tuning, vLLM inference, token-matched training and legal-citation counterfactuals.
- Contributed to evaluation across CaseHOLD and MultiEURLEX, including two 7B model families, three seeds and a reconstructed 197-document legal-entity resource.

Quantisation Limits of LLMs for Legal NLP

- Built a configurable evaluation framework for six open LLM configurations from 7B to 72B parameters, comparing post-training quantisation from FP8 to sub-2-bit variants.
- Executed and analysed 159 full-test result cells across LEDGAR, CaseHOLD and EURLEX, with paired statistical testing, structured-output audits, checkpoint-footprint measurement and backend controls.

LLM-Assisted Swiss Consultation Drafting

- Developed an LLM-assisted consultation-drafting backend with FastAPI, PostgreSQL and SSE streaming, including participant-linked transcripts and usage telemetry.
- Engineered transactional study-budget controls and recovery-aware multi-VM accounting, alongside representative multi-turn cost benchmarks and streaming load tests.

AI-Guided Online Deliberation

- Built a reproducible analysis pipeline for a four-arm randomised online-deliberation experiment covering 336 participants, 8,118 comments and 411,515 cleaned behavioural events.
- Implemented argument-tree compliance analysis, pre-registered statistical testing and LLM-based discourse annotation with blinded human validation and H100/Slurm inference.

LLM Code-Style Assessment and Human Evaluation

- Developed a Python code-style assessment pipeline integrating assignment context, Pylint and schema-validated feedback across OpenAI, Anthropic and local vLLM workflows.
- Built a shared 26-criterion human–AI annotation framework with source-integrity checks, reviewer adjudication and exact/overlap-aware evaluation of line-level findings.

PUBLICATIONS

Journal articles

1. **Ariai, F.**, Mackenzie, J., & Demartini, G. (2025). Natural Language Processing for the Legal Domain: A Survey of Tasks, Datasets, Models, and Challenges. *ACM Computing Surveys*, 58(6), Art. 163, 1–37. doi:10.1145/3777009
2. **Ariai, F.**, Tayefeh Mahmoudi, M., & Moeini, A. (2024). Enhancing Aspect-Based Sentiment Analysis with ParsBERT in Persian Language. *Journal of AI & Data Mining*, 12(1), 1–14. doi:10.22044/jadm.2023.13666.2482
3. Aryai, V., Kharazi, M., & **Ariai, F.** (2017). “Nearest Zero-point” Algorithm for Cooperative Robotic Search Missions. *IAES International Journal of Robotics & Automation*, 6(1), 49–58. doi:10.11591/ijra.v6i1.pp49-58

Conference proceedings

4. Aryai, V., Mahmoodian, M., Vahidi Ferdowsi, N., & **Ariai, F.** (2018). Structural Reliability Analysis of Corroding Steel Bridges Using Random-Field Representation. *Maintenance, Safety, Risk, Management and Life Cycle Performance of Bridges*. 9th International Conference on Bridge Maintenance, Safety & Management (IABMAS 2018), Melbourne.

TEACHING EXPERIENCE

Tutor — CSSE7030: Introduction to Programming

July 2025 – Present

The University of Queensland, Brisbane

- Tutor postgraduate students in Python programming, algorithmic thinking and software design.

INDUSTRY EXPERIENCE

Software Developer

September 2015 – January 2024

Neyestan Rayaneh, Mashhad, Iran

- Designed and maintained enterprise software systems.
- Managed the full software-development lifecycle: requirements, implementation, testing and deployment.

TECHNICAL SKILLS

- **Programming:** Python, SQL, Bash, JavaScript, C++
- **LLM adaptation and inference:** PyTorch, Hugging Face Transformers, PEFT/LoRA, QLoRA, vLLM, llama.cpp
- **Retrieval and evaluation:** Dense retrieval, BM25, reranking, Neo4j/Cypher; FP8, AWQ and GGUF evaluation; structured-output and checkpoint-footprint analysis
- **Applications and data systems:** FastAPI, PostgreSQL, SQLite, SSE streaming, OpenAI and Anthropic API integration
- **Experimental methods:** Controlled comparisons, counterfactual evaluation, paired bootstrap intervals, multiple-testing correction, human annotation and agreement analysis
- **Research infrastructure:** Slurm, GPU/HPC workflows, tensor parallelism, Docker, resumable experiments and automated testing
- **Languages:** Persian (native); English (professional / academic proficiency)

ACADEMIC REFEREES

Professor Gianluca Demartini — Principal Supervisor, School of EECS, UQ.

Dr Joel Mackenzie — Associate Supervisor, School of EECS, UQ.