JAMES C. CALDWELL

London, Ontario, Canada (open to relocation) · +1-519-636-9979 · James.Caldwell.000@gmail.com
jamescaldwell.net · github.com/James-C-000 · linkedin.com/in/james-caldwell-304235126

SUMMARY

NLP/LLM engineer with end-to-end experience building document AI pipelines, RAG systems, and large-scale text classification workflows. Recently designed and benchmarked a vision-language HTR pipeline for a 571,712-scan historical document corpus using vLLM on H200 GPUs, achieving 85.2% F1 and 94.5% recall against manually verified ground truth. Strengths in LLM inference and serving, prompt engineering, and shipping cross-platform tools from prototype to packaged release.

SKILLS

AI/ML: Large language models (LLMs), vision-language models (VLMs), retrieval-augmented generation (RAG), handwritten text recognition (HTR), supervised text classification, active learning, prompt engineering, data annotation & validation, text mining, model evaluation & benchmarking

Inference & Serving: vLLM, Ollama, llama.cpp

Languages: Python, C#, R, Java, Bash, HTML, CSS

Frameworks & Libraries: PyTorch, scikit-learn, FAISS, Transformers (HuggingFace), pandas, NumPy, BeautifulSoup4, NiceGUI, Pywebview

DevOps & Infrastructure: Docker, Git, GitHub Actions (CI/CD), Jenkins, AWS, Linux system administration

Data: Web scraping, corpus construction, bibliometric analysis

EXPERIENCE

Independent AI/NLP Consultant — London, ON *May 2019 – Present*

Independent practice delivering NLP/ML engineering, document AI, and data pipeline work for university research groups and nonprofit clients. Selected engagements:

ML Engineer – Document AI — Western University *June 2025 – March 2026*

- Built a Python pipeline for Environment Canada (EC) to organize a 6TB archive of 996,602 scanned historical weather observation forms (c. 1840–1960), resolving station Climate IDs against EC inventories, detecting cross-batch duplicates, and sorting the 571,712 in-scope pre-1940 scans by province.
- Developed a single-pass VLM extraction pipeline (Qwen3-VL served with vLLM) to capture handwritten qualitative weather remarks, with checkpointed crash recovery and a streaming design that satisfied EC data-handling requirements by never storing scans on remote GPUs.
- Benchmarked 12 model/resolution configurations (Qwen3-VL 4B/8B/32B at 2–16 MP) against a manually verified 150-file ground-truth set; the recommended configuration, Qwen3-VL-32B on a rented H200 GPU, achieved 85.2% F1 and 94.5% recall at a projected cost of US$687 for the full corpus.
- Authored client-facing project and cost/benefit reports quantifying accuracy, cost, and runtime trade-offs across configurations to support the client's production go/no-go decision.

Software Developer — Western University *Sept. 2024 – April 2025*

- Designed and built the Modular Digital Methodologies Toolkit (MDMT), an open-source, cross-platform desktop application with a NiceGUI + Pywebview interface and a local RAG chatbot (Qwen models + FAISS) (https://github.com/James-C-000/Modular-Digital-Methodologies-Toolkit).
- Packaged and shipped builds for Linux, Windows, and macOS, with releases automated via GitHub Actions CI/CD.

Research Consultant — Aboriginal Sport Circle, Ottawa, ON *Dec. 2023 – March 2024*

- Led a review and comparative analysis of provincial and federal Indigenous sport strategies; validated and analyzed stakeholder workshop data and delivered written reports to the client.

Software Developer — Western University *May 2023 – Sept. 2023*

- Built the Topic Clustering Toolkit, a Dockerized application combining Apache Solr and Carrot2 Workbench for unsupervised topic clustering of unstructured text corpora; authored setup scripts and full documentation and published the package openly.

Earlier Developer Contracts — Western University & UBC Okanagan *Nov. 2021 – May 2023*

- Wrote a Python (BeautifulSoup4) scraping tool to resolve institutional IP address pools and identify all associated Wikipedia edits, enabling provenance analysis of crowd-sourced encyclopedic content.

- Adapted keyword search pipelines to extract and classify targeted content from archived government and administrative records across two separate contracts.

Data Engineer — Western University / UBC *May 2019 – Present (ad hoc)*

- Built a Python pipeline to extract and classify quantitative and qualitative sport and physical activity data from the Indian Affairs Annual Reports (1864–1990), producing structured datasets now publicly available for research use; continue to provide ad hoc data engineering support.

Python Engineer – AI Trainer — DataAnnotation, Remote *Feb. 2026 – Present*

- Train agentic AI systems on software engineering best practices by evaluating model-generated code and reasoning, identifying flawed logic, hallucinations, and opportunities for optimization.
- Author reference solutions that define target model behavior on Python engineering tasks.

Quality Assurance Lead — Mikutech/Joydrop Ltd., London, ON *Sept. 2017 – Aug. 2019*

- Held a cross-functional QA, DevOps, web development, and IT role at a game development studio; built and maintained a Jenkins CI/CD pipeline with custom C# tooling.
- Maintained the company website (HTML, CSS, JavaScript) on AWS; performed manual QA testing and managed bug tracking.

SELECTED PROJECT

Bibliometric ML Classification Pipeline — M.A. thesis research infrastructure *2024 – 2025*

- Trained a supervised classification pipeline (scikit-learn) on ~4,000 labeled examples via active learning to filter a 45–50k-record bibliographic corpus constructed from Web of Science, Scopus, and PubMed with custom Python and R pipelines.
- Built an author name disambiguation pipeline combining n-gram analysis, Levenshtein distance, and phonetic matching.
- Performed collaboration network analysis using Walktrap and Louvain community detection, with burst detection to identify temporal inflection points in the scientific publication record.

PUBLICATIONS

Caldwell, J.C. *Antibiotic Pollution and the Making of a Global Environmental Crisis (1943–2002)*. M.A. thesis, Western University, 2025. hdl.handle.net/20.500.14721/39080

Caldwell, J.C. "'Something in the water': Aquatic Pharmaceutical Pollution in the United States." *Inspiring Minds*, Western University, 2024. ir.lib.uwo.ca/inspiringminds/606

Forsyth, J., Habkirk, E., & Caldwell, J.C. "Indian Affairs Audit." indigenoussporthub.com/indian-affairs-audit-1

SELECTED TALKS

"Generative AI and Data Analysis," Rotman Institute of Philosophy, Western University (2025)
"Generative AI and Teaching," School of Kinesiology, University of British Columbia (2025)
"'If Only There Was an Algorithm': Expediting Historical Research With Computer Programming," North American Society for Sport History, Chicago, IL (2022)

EDUCATION

M.A. in History — Western University, London, ON *2023 – 2025*

- Ivie Cornish Memorial Fellowship (2025). Thesis research involved building the bibliometric ML classification pipeline listed above and conducting bibliometric analysis of publication data.

B.A. in Psychology — Western University, London, ON *2018 – 2023*

- Dean's Honor List (2021–22, 2022–23).

LANGUAGES

English (native) · Japanese (basic conversational, reading, writing)