

Haoqian (Frank) Zhang

(+1) 437-971-8869 | haoqian.zhang@mail.utoronto.ca | [LinkedIn](#) | [GitHub](#)

EDUCATION

University of Toronto

Toronto, Canada

Honours Bachelor of Science in Computer Science (Co-op)

Sep. 2021 – Apr. 2026

- Software Engineering Stream, Minor in Statistics
- Overall GPA: 3.55/4.00; Major GPA: 3.75/4.00
- Core Courses: Engineering Large Software Systems, Computer Networks, Compilers & Interpreters, Operating Systems, Algorithm Design and Analysis, Data Structures, Computer Organization, Intro to Databases, Intro to Machine Learning and Data Mining, Intro to Software Engineering, Regression Analysis
- Honors: Dean's List

TECHNICAL SKILLS

Languages: Java, C/C++, Python, R, JavaScript, TypeScript, HTML, CSS, SQL

Databases: PostgreSQL, Milvus, MySQL, MongoDB, MarkLogic, Firebase

Frameworks & Tools: React, LangChain, LangGraph, Spring Boot, Node.js, Angular, Express, Docker, Git, Postman, CI/CD

EXPERIENCE

Scholars Portal

Toronto, Canada

Software Development Engineer Intern

Sep. 2023 – Sep. 2024

- Built an **academic search platform** using **Angular** with query-parameter-based navigation, enabling **dynamic filtering, faceted search, and paginated exploration** of large-scale document collections.
- Designed and implemented a **scalable data ingestion pipeline** using **Apache NiFi** to transform and index data from **PostgreSQL** into **MarkLogic**, supporting downstream **search and RAG-based applications**.
- Developed reusable **NiFi templates** and multi-table join workflows to convert relational data into hierarchical **XML/JSON** structures, improving pipeline efficiency by **78%**.
- Developed **Java-based data validation and recovery tools** with dynamically generated **XQuery** to reconcile inconsistencies across heterogeneous publisher datasets.
- Implemented **ISSN-based metadata matching** to automate title recovery and improve data consistency across ingestion workflows.

PROJECTS

Humanly | *React, TypeScript, Node.js, PostgreSQL, Redis, Docker, GCP* | [Live Demo](#)

Sep. 2025 – Mar. 2026

- Built an **event-driven writing activity tracking system** that captures fine-grained typing events and streams them to backend services for near real-time analysis and storage. Designed **modular AI workflows** for multi-step query processing using **LLMs** with contextual input and response generation.
- Deployed a containerized full-stack application on **Google Cloud Platform** and implemented **CI/CD pipelines** using **GitHub Actions**.

UniLocker | *MongoDB, Express, React, Node.js*

Sep. 2025 – Dec. 2025

- Built and integrated **RESTful APIs** between the React frontend and Express/MongoDB backend, streamlining item listing and retrieval workflows and reducing frontend-backend integration issues by **20%**.
- Implemented **Google Maps API** features for locker visualization and navigation, reducing average locker lookup time by **~30%** based on user interaction flow improvements.
- Collaborated with a team of 7 to integrate new features and enhance platform performance.

Apache Superset Open Source Contributor | *JavaScript, TypeScript* | [PR Link](#)

- Implemented **filtered CSV/Excel export functionality**, ensuring exported data respects frontend query states and user-applied filters. for table visualizations, ensuring exports respect frontend search filters.
- Modified Superset table chart export pipeline across the React frontend and backend export logic.
- Contributed production-ready code to Apache Superset through collaborative code reviews.