

# Daniel Chen

226-821-0808 | [danielchenschw@gmail.com](mailto:danielchenschw@gmail.com) | [linkedin.com/in/daniel-chen-uo](https://linkedin.com/in/daniel-chen-uo) | [github.com/danthepanman](https://github.com/danthepanman)

## EDUCATION

### University of Ottawa

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

Ottawa, ON

Aug. 2023 – Jan. 2028

## EXPERIENCE

### Generative AI developer Intern

Sept. 2026 – present

*Kinaxis*

Ottawa, ON

- Reduced LLM context from **8K to 2K** tokens using LLM-based context compaction, preserving relevant conversation history for downstream inference.
- Integrated **3+ LLM providers** through **Kong** Gateway, routing API requests and normalizing provider-specific request/response formats for consistent downstream consumption

### Incubation Developer Intern

May. 2026 – Aug. 2026

*Nokia*

Ottawa, ON

- Indexed **10K+** log statements in a prototype hybrid search system, combining **BM25**, **vector search**, and **Reciprocal Rank Fusion (RRF)** to semantically locate relevant log sources.
- Optimized a Jira analytics dashboard for **10,000+** concurrent users by refactoring the scheduler to use **server-sent events**, reducing server-side CPU and Disk I/O load by **99%**.
- Implemented the **LogRCA** framework in **PyTorch**, leveraging **Positive-Unlabeled learning** and a **Transformer-based** ranking model for log root cause analysis, improving Hit@5 by **15%**.

### MLOps Engineer Intern

Jan. 2026 – Apr. 2026

*Transport Canada*

Ottawa, ON

- Enhanced F1 score of a fine-tuned **Hugging Face Transformer NLP model** implemented in **PyTorch** by **15%** through hyperparameter tuning and k-fold cross-validation in Databricks.
- Scaled **MLflow** for experiment tracking to **4** production-grade ML workflows, logging **200+** experiments and enabling rapid comparison across multiple model families and hyperparameter grids.
- Automated **Databricks** data pipelines to ingest, clean, and standardize **1M+** waybill records from production, improving data accuracy and enabling scalable downstream analytics.

### Data Engineer Intern

May 2025 – Sept 2025

*Aedo Data Solutions*

Gatineau, QC

- Designed and developed data pipelines with automated ETL workflows processing over **5M+ records daily** using **Azure Data Factory**, **Azure Functions**, **Azure Blob Storage**, and **Dagster**.
- Reduced computing costs by **\$10k/month** by migrating stored procedures and data transformations away from **Snowflake** to **Microsoft SQL Server**.
- Boosted client conversion rate by **50%** through **AI-assisted (Claude Code, Cursor)** prototyping used in sales demos, accelerating iteration cycles and improving stakeholder feedback turnaround

## PROJECTS

### Concerto (Hack the North 2026 Winner) | TypeScript, Next.js, WebSockets, Python, OpenCV, PyAV

- Built Concerto, a real-time concert platform that turns audience phones into synchronized musical instruments, with live section assignment, playback controls, mixing, and emergency muting.
- Implemented with **TypeScript**, **Next.js**, **WebSockets**, **OpenCV**, and **PyAV**, with an optical temporal-encoding pipeline that detected and localized phones from camera footage with **99%** accuracy.

### Dan-AI | Next.js, Tailwind, Python, OpenAI API, Pinecone

- AI-powered** portfolio allowing users to explore projects and experience using natural language.
- Implemented a **RAG** pipeline with semantic search, delivering context-aware and accurate responses.
- Indexed project descriptions and experience in **Pinecone**, allowing the AI to reference specific content.

## TECHNICAL SKILLS

**Front end:** JavaScript/TypeScript, HTML, CSS, React, Tailwind, Jest, Webpack, Prisma ORM, Blazor

**Back end:** Python, Node.js, Express.js, Next.js, Flask, FastAPI, REST API, C#, .NET

**Cloud & Data:** SQL, Snowflake, Firebase, MongoDB, Docker, Dagster, Azure Data Factory/Functions/DevOps

**Machine learning & AI:** Scikit-learn, Matplotlib, Pandas, Databricks, MLflow, PyTorch, Hugging Face, Ollama