

Mikhail Wilson

+1 (905) 514-6244 | mwils338@uwo.ca | linkedin.com/in/mikhail-wilson | mikhailwilson.ca

EDUCATION

Western University

Bachelor of Science, Pure Mathematics — Minor in Computer Science

London, ON

Sept. 2024 – May 2027

EXPERIENCE

Software Engineer

University of Toronto

Jan. 2026 – Apr. 2026

Toronto, ON

- Architected and deployed a serverless event registration platform on AWS Lambda + API Gateway, handling 500+ registrations with zero downtime and cutting administrative overhead by 60%.
- Designed a real-time attendee pipeline syncing Lambda event triggers to Google Sheets, enabling live capacity tracking and automated transactional emails without persistent backend infrastructure.
- Implemented idempotent ticket generation with UUID-keyed deduplication to prevent double-registration on retry; caught a silent data corruption edge case during load testing before production launch.

Data Engineer & Co-Founder

DLICIO | dlicio.com

Feb. 2024 – Dec. 2025

Waterloo, ON

- Designed event-driven microservice architecture on AWS using Apache Kafka, processing 2.1M user interaction events/month across discovery, engagement, and feed ranking pipelines; owned schema evolution and consumer group lag monitoring.
- Built and maintained React Native client and REST API layer serving 55k requests/day at p95 ~180ms; profiled and resolved N+1 query bottlenecks in feed endpoints degrading latency under concurrent load.
- Engineered video delivery pipeline on S3 with signed URL expiration and CDN cache headers to reduce egress cost and prevent content leakage; catalog scaled to 80+ clips across 1,000+ pilot users.
- Instrumented Kafka consumer lag and API error rates with CloudWatch dashboards and PagerDuty alerting; reduced mean time to detection on feed outages from hours to under 8 minutes.

Data Engineer

Bruce Power

May 2025 – Jan. 2026

Tiverton, ON

- Spearheaded end-to-end migration of 35,000+ nuclear licensing documents from a brittle VBA/Access system to a normalized SQL database; engineered Python ETL pipelines achieving 95%+ reduction in query latency and eliminating a class of data integrity errors caused by Access locking.
- Built an automated ingestion engine using Python (OpenCV, spaCy, GPT-4) to parse and classify 40,000+ legacy scanned records; implemented computer vision pipeline for checkbox detection on degraded scans, achieving production-grade accuracy on a dataset with no ground-truth labels.
- Designed a confidence-scoring layer on the GPT-4 extraction pipeline that flagged low-certainty records for human review rather than silent insertion, reducing erroneous records in the SQL database to under 0.3% and making the system auditable for a regulated nuclear environment.
- Delivered Power BI audit dashboards for CNSC regulatory workflows with Row-Level Security; automated Dev/Test/Prod promotion via Azure REST API CI/CD, eliminating environment drift and manual deployment overhead.

PROJECTS

Unchained — Supply Chain Disruption Prediction | *PyTorch Geometric, Next.js, deck.gl, Mapbox* | [GitHub](#)

- Designed and trained LifelineGNN, a spatiotemporal Graph Attention Network + GRU on a 500-node supply chain graph; achieved 17.7% RMSE improvement over persistence baseline by capturing spatial neighbor dependencies and temporal shock propagation simultaneously.
- Formulated a physics-informed loss function combining shortage-weighted MSE with a differentiable sigmoid-gated conservation penalty, enforcing inventory conservation as a hard structural constraint rather than relying on emergent model behavior.
- Ablated the GAT attention mechanism against a vanilla GCN baseline and documented per-layer attention weight distributions, confirming the model correctly upweighted high-degree hub nodes during shock propagation — interpretability evidence beyond raw RMSE.
- Engineered a Next.js / React frontend with deck.gl and Mapbox animating 365-day disruption cascades across all 500 nodes in real-time, synchronized with per-node inventory state and forecast line charts.