

# Aleksandr Gyumushyan

647-655-4345 | [alex.gyumushyan@gmail.com](mailto:alex.gyumushyan@gmail.com) | [LinkedIn](#) | [Github](#) | [Personal Website](#)

## EXPERIENCE

---

### Junior Software Developer

Sep. 2025 – Present

*LDA Technologies*

*Mississauga, ON*

- Designed and owned a concurrent, fault-tolerant Python library to interface with server modules over serial connections, utilizing Protobufs and code generation to automate future hardware API integrations
- Collaborated directly with customers to diagnose intricate race conditions, engineering custom stress tests to validate fixes and ensure high system reliability in the field.
- Architected a central REST API layer over a single server (with multiple modules), and a client-facing CLI featuring dynamic, state-aware autocomplete

### Backend Developer Intern

May 2023 – Aug. 2023

*Geotab*

*Oakville, ON*

- Achieved a substantial reduction in Auth API requests (50%+) by re-engineering Geotab's Gateway auth token caching, and using refresh tokens. Stored refresh tokens in a Redis cache and created a new workflow of continuous retrieval and refresh of said tokens
- Deployed a secure, production-grade Redis cache within a microservices environment by leveraging Helm Charts and updating core Kubernetes and Terraform configurations.
- Enhanced developer testing workflows by modernizing the Docker Compose environment with an updated mock API, Redis cache, and profiler, enabling faster feedback cycles and improved debugging of cache performance

### Software Developer Intern

May 2022 – Apr. 2023

*Geotab*

*Oakville, ON*

- Designed and developed a CLI based internal data transformation tool, written in Python, that converts vehicle CANBus data from a DB into either a proprietary C-macro based, or a serialized Protobuf data structures, resulting in an over 95% speedup in vehicle signal upload to a Godevice
- Improved code reliability by designing a comprehensive test suite with over 30 unit and system tests, leveraging mocks, stubs, CI pipelines, static analysis tools and synthetic data to ensure high code coverage and reliability
- Engineered a comprehensive logging system to log tool usage through over 20 parameters to a GCP Big Query table, enhancing error logging, compliance and data accessibility

## PROJECTS

---

### Ultra-fast 1 Billion Row Challenge | [Repo](#) | [Writeup](#) | *Zig, Samply*

Jan. 2026

- Engineered a high-performance data ingestion engine in Zig to parse a 15GB dataset (1 billion records), achieving a single-core processing time of 9.3 seconds
- Slashed execution time by over 1,000% from a naive solution by bypassing standard I/O with zero-copy memory mapping and designing a tuned hash table with linear probing
- Optimized text-parsing throughput by replacing for loops with SWAR (SIMD Within A Register) bitwise operations, utilizing CPU profiling to precisely implement and benchmark this algorithm

### Distributed Key Value Store | [Repo](#) | [Writeup](#) | *Go, Docker, Kubernetes*

Aug. 2025

- Engineered a consensus-based distributed key-value store with fault tolerance and dynamic cluster resizing via a gossip-based membership protocol, deployed locally with Kubernetes
- Implemented Hashicorp's Memberlist for gossip-driven node discovery and Raft for consensus, log replication and leader election, enabling reliable log consensus and automatic failover across nodes in dynamic clusters
- Deployed and documented a local cluster through Kubernetes and Minikube with a StatefulSet, a ConfigMap, a headless local network and publicly exposed HTTP ports, creating a reproducible demo environment

## EDUCATION

---

### Toronto Metropolitan University

Toronto, ON

*Bachelor of Science in Computer Science - 3.92 GPA*

*Sep. 2020 – May 2025*

## TECHNICAL SKILLS

---

**Languages:** Go, Python, Java, C#, TypeScript, Zig

**Databases and Infrastructure:** MongoDB, Redis, RabbitMQ, Kubernetes, Docker, gRPC, Protobuf, Terraform

**Developer Tools and Testing:** Git, Jira, Linux, Pytest, Go testing framework, Pydantic, Flake8