

Truman Chan

647-616-3752 | trumanzhe@gmail.com | linkedin.com/in/trumanchan | github.com/tmanzhe | tmanzhe.com

EDUCATION

York University & Schulich School of Business

BSc (Honours), Computer Science & Master in Management (Double Degree)

Toronto, ON

Sept. 2024 – Apr. 2028

EXPERIENCE

Software Engineer Intern

Kepler Communications

Toronto, ON

Sept 2025 – Apr 2026

- Contributed to design reviews for a **v2 telemetry architecture** scoped after the existing pipeline failed under a new satellite tranche, tiering **~500K records / 10 min** off a single **TimescaleDB** instance into **Apache Iceberg** tables on **S3** with Timescale retained as a hot cache.
- Rebuilt observability for the satellite telemetry pipeline, migrating CSV ingester logs to structured **NDJSON** with 15+ queryable fields and wiring **Elasticsearch** and **Grafana** aggregations, making ingest failures queryable by on-call engineers for the first time.
- Architected and built a **versioned database migration engine** from scratch, replacing a manual rollout process that bottlenecked a **live satellite launch**, with a **pluggable adapter interface** unifying **MySQL**, **TimescaleDB**, and **ZODB** and **strict linear execution** that rejects out-of-sequence migrations.
- Consolidated and **baselined 139 historical migrations** across **447 files** into a centralized **schema history table** with **SHA-256 checksum validation** and preflight health checks, embedding the engine in **CI/CD** to make deployments **idempotent, auditable, and recoverable mid-run**.
- Automated idle-transaction cleanup on live production clusters with a **Kubernetes CronJob**: Python probes detect and terminate long-running sessions holding resources, containerized with **Docker** and wired to **Slack alerts** so the team sees every action it takes.

Software Engineer Intern

AATENG Group

Richmond Hill, ON

May 2025 – Aug 2025

- Led a four-engineer team building and shipping an esports tournament platform on **AWS (Next.js, Spring Boot)** that ran live events for **1,000+ players** from a **10,000-member community**.
- Owned backend development across **18 JPA entities**, 9 REST controllers, and 14 services, securing every endpoint with **Spring Security**, a custom **JWT filter**, and **Discord OAuth2**, and shipping incrementally behind **server-side feature flags**.

PROJECTS

LLM Extraction Eval Harness | *Python, LangGraph, DuckDB, Parquet, FastAPI*

GitHub Repo

- Designed the measurement methodology: **per-field scorers** sort predictions into four outcome classes, **hallucination rate is reported separately**, and **Wilson confidence intervals** keep results from being claimed below the sample's **detectable difference**.
- Ran a **power analysis before labeling** so undersized slices were oversampled rather than reported as findings, over a pipeline parsing **1,995 job postings** into a validated schema, served from **Parquet** via embedded **DuckDB** behind **FastAPI**.

Clash of Plans | *FastAPI, LangChain, Cohere, Next.js, Supabase, Docker*

GitHub Repo

- Built an **AI agent** that plans student schedules: classifies request intent, then returns **typed, schema-validated objects** rather than free text, on a **FastAPI/Supabase** backend with **Google Calendar** sync over **OAuth 2.0**.

Homelab | *RKE2, Ansible, MinIO, MetalLB, ArgoCD, Tailscale*

- Self-hosted bare-metal **RKE2** Kubernetes cluster provisioned by idempotent **Ansible** roles, running **MinIO**, **MetalLB**, self-hosted **CI/CD**, and **Prometheus**, with DNS and VPN hosted outside the cluster so a cluster failure cannot take down the tools needed to fix it.

TECHNICAL SKILLS

Languages: Python, Java, TypeScript/JavaScript, SQL, C

Backend & Web: Spring Boot, Spring Security (JWT), OAuth 2.0, JPA/Hibernate, REST APIs, FastAPI, Node.js

Data & Cloud: PostgreSQL, TimescaleDB, MySQL, DuckDB, Parquet, Elasticsearch, AWS (S3, RDS, EC2), GCP

Infrastructure & Tooling: Kubernetes, Docker, Ansible, ArgoCD, GitHub Actions, Grafana, pytest, LLM/agent tooling