

Dan Huynh

☎ 226-606-9284 | ✉ dan.huynh@uwaterloo.ca | 🔗 LinkedIn | 🐙 GitHub | 📁 Portfolio

OVERVIEW

Languages: C/C++, Java/Scala/Kotlin, Python, TypeScript, GO, PHP, SQL, Swift, MATLAB, VHDL

Technologies: Docker/k8, ArgoCD, Cloud Native, gRPC, REST, React/Angular, ROS2, Statistical ML Libraries

PROFESSIONAL EXPERIENCE

Software Engineer @ Tesla

May 2025 – August 2025

- Architected a cloud-native distributed microservices platform for vehicular crash simulation analysis using **Next.js**, and **Litestar**, implementing 130+ **RESTful** endpoints, 3+ UIs for data manipulation, **gRPC** inter-service communication with daemon threading, and end-to-end type safety via **openapi-typescript** serving 56+ users.
- Designed a hybrid **OLTP/OLAP** database layer (56+ tables) with cross-schema service isolation to centralize vehicle metadata and 1B+ sensor samples annually, using **SQLAlchemy** and **Drizzle ORM** to integrate with a proprietary CLI for automated migrations, and to provide pipelines for real-time and batch analytics.
- Integrated **Azure OIDC** authentication from scratch, with a thread-safe **JWK** caching layer, enhancing security.
- Built CI/CD pipelines for 5+ services using **ArgoCD**, **Kubernetes**, and **GitHub Actions** with static type analysis, linting and conditionally built containers based on code diffs, achieving a 99.96% service uptime.

Software Engineer @ Hevy

January 2025 – April 2025

- Implemented a fuzzy exercise search engine using **fuse.js** and a proprietary recency/frequency-weighted bitap algorithm, achieving a 100% increase in discovery across 12+ locales and reducing search latency by 15%.
- Delivered cross-platform quick actions to start recommended exercises using **Swift/Kotlin** native modules with a **React Native** bridge and **MobX** reactive updates reducing cold-launch workout start times by 35%.

Data Engineer @ theScore

Sept 2024 – Dec 2024

- Designed a configuration-driven **Python** data monitor for rapidly scaling, sport agnostic, odds projection models, validating projection availability, data sparsity across 27+ tables in **BigQuery** and **PostgreSQL**, and anomaly detection using distributional divergence and z-score metrics which trigger alerts to **Slack** and **DataDog**.
- Created an **Argo Workflow** for dynamic table data migration between **BigQuery** and **PostgreSQL** schemas, avoiding **GCP** rate limits for 10M+ row transfers while ensuring row uniqueness and schema consistency.

Software Engineer @ Vivid Seats

Jan 2024 – Apr 2024

- Designed a test-data Backend for Frontend (BFF) with **Spring**, exposing RESTful endpoints for **JPA** entity generation and management that insert regression-agnostic data for end-to-end checkout tests into 4+ vault-authenticated **AWS Aurora** databases, resulting in a 60% build time reduction.
- Developed a distributed subscriber for the test-data BFF which enables on-demand and autonomous data cleanup via client API calls and a proprietary cleanup micro-service, allowing for concurrent data management.
- Integrated **SonarQube** for static code analysis and developed 210+ unit and integration tests using **JUnit5**, **Mockito**, and **Spring** that achieved 96% code coverage for the test-data BFF.

Data Scientist @ PureFacts Financial Solutions

May 2023 – Aug 2023

- Designed and tuned an **Sklearn** bayesian optimized random forest regressor with a mean percentage error of 17.32% that predicts client revenue movements, while providing explanations for model predictions using **SHAP**.
- Built a dashboard using **Plotly Dash** that features dynamic visualizations of investor revenue, AUM, transactions, and customer trends over time for PureFacts clients, enabling data-driven decision making.
- Led development of a **Flask + React** LLM interface integrating proprietary data utilizing **OpenAI APIs** that empower non-technical personnel with accessible information and reduces engineer labor time.

Software Engineer @ Ford Motor Company

Sept 2022 – Dec 2022

- Created **React** components for the fordpro.ca micro-frontend, including a file upload dropzone and a data-model agnostic fuzzy search feature, querying 1000+ **Firestore** records with Regex filtering across 7+ properties.

PROJECTS

Perceptions Lead @ Watonomous — LiDar Object Detection | Github

Sept 2023 – Sept 2024

- Designed a data loader for **OpenPCDet** to process 32-beam Velodyne point clouds into **NumPy** arrays, optimized for **VoxelNeXt**, **TransFusion**, and **PV-RCNN** predictions.
- Wrapped **OpenPCDet** in a **ROS2-humble** node to process point cloud feeds and publish real-time bounding box predictions via **Foxglove** WebSocket for immediate data visualization.
- Collaborated on the design and development of a multi-sensor perceptions pipeline for real-time object tracking.

EDUCATION

University of Waterloo — GPA: 3.9/4.0

2021 – 2026

Honours B.ASc Candidate in Mechatronics Engineering (Option in Software Engineering)

Waterloo, Canada