Brandon Zhou

Fourth-year Computer Science Student
github.com/brandonzhou2002 | linkedin.com/in/zihaozhou

Waterloo, ON
Mobile: 437-788-2669
Email: zihao.zhou@uwaterloo.ca

Technical Skills

Languages: Python, Go, SQL, GraphQL, C, C++, Java, Scala, HTML, CSS, TypeScript, R

Tools: GCP, Azure, AWS, Docker, FastAPI, Uvicorn, Alembic, Apache Airflow, Terraform, Hugging Face, Git, GitHub, GitHub Actions, Linux, Supabase, Firebase, Pinecone, LangChain, React

Working Experience

Software Developer Co-op Jan 2025 – Dec 2025

Dawson Partners - a private equity firm offering structured liquidity solutions Toronto, ON

- Re-engineered the core backend services for the General Partner Portal, optimizing the financial ETL pipeline via **batch API processing** and consolidated **SQLAlchemy ORM query logic** to cut latency by **66%**, while implementing an automated conflict resolution algorithm that reduced manual data review by **70%**.
- Architected a testing framework utilizing **Pytest** and **SQLAlchemy runtime introspection** to auto-generate test data in **isolated environments**, integrated into the **CI/CD pipeline** to reduce production bugs by **35%**.
- Owned the Budget Management application, implementing **temporal versioning** and refactoring **ORM query patterns** to decouple monolithic tables, reducing redundancy by **30%** and eliminating **N+1 bottlenecks**.
- Built a payment orchestration service managing transaction lifecycles by triggering invoice submissions via the **Kyriba API** and consuming **webhooks** for real-time synchronization, reducing turnaround time by **40%**.
- Refactored the **GraphQL base layer** to consolidate logic for processing complex nested data structures, reducing code redundancy by **40%** while implementing **MD5-based hashing** to ensure deterministic ID generation.

Research Assistant [arXiv:2507.06430] Sep 2024 – Dec 2024

University of Waterloo - a research-intensive institution driving innovation Waterloo, ON

- Developed a reproducible data ingestion pipeline to process raw **network traffic logs**, parsing unstructured JSON into structured data frames to ensure schema consistency with benchmark research datasets.
- Derived features from **packet headers**, including payload length, directionality, and inter-arrival times, to optimize input dimensionality for machine learning classifiers.
- Trained and benchmarked a suite of models including **kNN**, **XGBoost**, and **Transformers**, achieving **60% accuracy** in activity classification by evaluating model generalization on temporally distinct datasets.

Data Engineer Co-op Sep 2023 – Dec 2023

JANA Corporation - a natural gas firm offering pipeline consulting services Aurora, ON

- Built ETL pipelines using **Azure Data Factory** and **Apache Airflow** on **Docker**, reducing job execution time by **11%** while orchestrating automated monthly workflows.
- Codified cloud infrastructure using **Terraform** modules to enforce **Infrastructure as Code (IaC)**, eliminating **configuration drift** and reducing environment provisioning time by **50%**.
- Implemented the **Data Vault 2.0** architecture (**Hub, Link, Satellite**) to decouple ingestion from business logic, allowing for seamless **schema evolution** while improving integration throughput by **20%**.
- Engineered a secure development workflow via **Google CLI** and **Cloud Build** that enables local development without persistent credentials, guaranteeing **100%** data confidentiality and minimizing manual intervention.

Data Scientist Co-op Sep 2022 – Apr 2023

Septodont - a pharmaceutical company specializing in dental and medical products Cambridge, ON

- Researched and evaluated **10+ machine learning algorithms** for customer retention and sales forecasting, developed prototypes, and optimized solutions, aligning outcomes with business objectives of increasing revenues.
- Deployed an **XGBoost model** on **Azure** to predict purchase probabilities, boosting target customer identification accuracy from **23% to 28%** across **170k+ customers** and outperforming the **K-means** model in production.
- Automated the **model development pipeline**, from feature engineering to deployment, increasing scalability across various products and boosting model iteration speed by **40%**.

Education

University of Waterloo Waterloo, ON

Candidate for Bachelor of Computer Science (cGPA: 92%) Sep 2021 – Apr 2026