

Jerry Zhu

<https://jerryzhu.org>

Email: j25zhu@uwaterloo.ca

Devpost: devpost.com/Bobliuuu

Github: github.com/Bobliuuu

EDUCATION

University Of Waterloo

Computer Science, Pure Math, Computational Math, Combinatorics with Co-op, 3.8 GPA
AWS Captain, CSC, DSC, GDSC, Tech+, TedxUW, Mathsoc, MEF, WAT.AI, WARG, UTMIST, Varsity Ultimate

Sept. '22 – Apr. '27

Waterloo, ON

SKILLS

Languages: Go, Python, C++, Java, TypeScript, SQL, Rust, Solidity

ML/Data: PyTorch, TensorFlow, LLMs, LangChain, DBT, Snowflake

Infra: AWS, GCP, Kubernetes, Terraform, Docker, Bash

EXPERIENCE

StackAdapt, Software Engineer

May '25 – Aug '25

- Built internal dashboards workflows used by analytics teams to target large-scale segment data
- Optimized GraphQL queries across Go and Ruby on Rails services, reducing fetch latency by **35%**
- Designed backend APIs and data models for million record audience segmentation pipelines
- Contributed to internal Chrome extension with **60%** engineer adoption, improving code ownership visibility

Super.com, Data Engineer

May '25 – Aug '25

- Built/audited 50+ DBT models powering Looker and Redash dashboards, increasing data integrity by **30%**
- Led migration of ETL ingestion from Fivetran to Estuary, Snowpipe, and Airflow, improving reliability, maintainability, and contributing to a **10%** MRR lift
- Reduced analytics and warehouse costs by **20%** through Snowflake sharding strategies and HighTouch and Amplitude optimization

Aview International, Full Stack Software Engineer

Sept '23 – Dec '24

- Pioneered an agentic e2e video distribution pipeline on AWS and GCP, increasing throughput by **8x**
- Trained a neural TTS dubbing model using NestJS, supporting 30+ languages and serving 100+ users
- Reduced North Star KPI costs by **30%** through AI editing and automated review flows
- Optimized frontend renders through Next.js SSR and CI/CD with Docker and Kubernetes, reducing LCP by **60%** and TTFB by **20%** with SSE streaming for progressive DOM updates

OpenLane, DevOps/SRE Software Engineer

May '23 – August '23

- Automated **10+** deploy pipelines with **Retool**, **ArgoCD**, **Terraform**, and **Azure DevOps** CI/CD
- Designed an access-elevation system on EKS using Helm and Istio, reducing Orca security alerts by **10%** and increasing RBAC test coverage by **30%**
- Operated Kubernetes backed infra and observability stacks with Splunk, PagerDuty, Honeycomb, and Grafana, resolving 100+ critical incidents, increasing system observability coverage by **40%**

Aviato, Founding Machine Learning Engineer

Dec '22 – May '23

- Developed a recommendation system using NLP models to predict acquisition metrics with a **85%** accuracy
- Built a semantic search and summarization engine over startup datasets using multi-shot text-to-SQL LLMs with LangChain and Pinecone, improving query latency by **4x**
- Pioneered a scalable VC data ingestion pipeline aggregating 20+ sources, increasing valid ingested data by 250% using DynamoDB, OpenSearch, Kubernetes scraper fleets, and Rayobyte proxies

Sysconverge, Software Engineer Lead

July '22 – Dec '22

- Built automated energy reporting pipelines using EnergyPlus SDK and PyTorch, leading to **6x** speeds
- Applied SAHI segmentation, CV, and OCR to classify architectural artifacts with **90%** accuracy
- Led development of a full-stack web application using TypeScript, Tailwind CSS, MongoDB, and C#, automating Ontario Building Code 2D-to-3D modeling workflows
- Delivered 100+ automated modeling projects using custom Ladybug XML and custom document intelligence driven workflows

PUBLICATIONS

Talk to your data: Enhancing Business Intelligence with LLM Text-to-SQL

IEEE Link

- LLM-driven semantic parsing + text-to-SQL for BI systems. Presented at DATA '23

Leveraging AI for Enhanced Synthetic Generation for Data Imbalance

IEEE Link

- GAN/SMOTE/VAE ensemble methods for binary classification data augmentation.
- Presented at IATMSI '24

PROJECTS

Shopify Monolith

Bobliuuu/shopify-monolith

- Created an e2e data pipeline to sync Shopify and local data to distributed DB with query resolvers
- Used SDV schema inferencing to deploy Airflow scripts to sharded tables and export BI tables with **4x** manual speeds
- Used custom MCPs with Ruby on Rails to optimize insights with **90%** data coverage

CSC Cloud

Documentation

- Act as sysadmin for CSC's cloud services with **1000+** active members servicing **10+** machines
- Manage server mirror infra using **Kerberos** and **LDAP**, fixing **100+** critical traffic issues

ShallowSearch

Bobliuuu/ShallowSearch

- Evaluated LLM synthetic data generation on DeepSeek R1 70B and Groq LLaMa 70B distilled on Ollama
- Achieved **91%** semantic similarity on Apache error logs winning Rootly's data challenge

Good Samaritan

Bobliuuu/Good Samaritan

- Developed a cross-platform mobile app using **Kotlin** & **Jetpack** to keep students safe on campus
- Integrated real time alerting and tracking for **100+** users with **30%** faster report time