

Tejas Wilkhoo

✉ twilkhoo@gmail.com | 🏠 twilkhoo.com | 📷 twilkhoo | 🌐 twilkhoo

Education

University of Waterloo

Bachelor of Computer Science, Management Minor, Dean's Honors

Waterloo, ON

Sept 2021 - April 2026

- **Courses:** Distributed Systems, Artificial Intelligence, Operating Systems, Networks, Algorithms, Data Structures. **GPA: 3.9.**

Work Experience

Decagon

Member of Technical Staff

New York, NY

Jun 2026 - Present

- Building AI Agents to deliver concierge customer service experiences.

Bloomberg

Software Engineering Intern

New York, NY

Sep 2025 - Dec 2025

- Architected a graph-based semantic modeling system using **RDF4J** to map relationships across **30+ production services**, enabling engineers to explore and visualize complex schema intersections within the **Realtime Enterprise Data Services** ecosystem.
- Engineered an automated ingestion pipeline using **Python** to parse **Bloomberg API schemas** and enrich **100,000+ nodes** with business metadata, while providing structured access to the graph database using a **SPARQL** service.

Databricks

Software Engineering Intern

Bellevue, WA

May 2025 - Aug 2025

- Launched Schema Evolution in **Apache Spark** Structured Streaming, enabling **Fortune 500** clients with **multi- $\$XXM$** data pipelines seamlessly propagate schema changes from ingestion sources (Kafka, Delta, etc.) to existing stateful streaming workflows.
- Migrated the **RocksDB** intermediate state format to utilize **Apache Avro** encoding to leverage native type evolution capabilities while maintaining **production-grade throughput**, using **Scala** and **Bazel**.

Bloomberg

Software Engineering Intern

New York, NY

Sep 2024 - Dec 2024

- Led the design and implementation of **Apache Kafka** ETL pipelines facilitating **hundreds of thousands of fixed income** post-trade records daily, for a new **Apache Cassandra** database cluster featuring a refined node partitioning strategy.
- Significantly optimized **multithreaded database queries** to minimize latency, using **Apache Spark** with **Python** (PySpark) and **DataStax's C++ Cassandra** driver, leading to a decrease from **5 hours to 8 seconds** (99.94% faster) for client-side data acquisition.

Google

Software Engineering Intern

New York, NY

May 2024 - Aug 2024

- Architected **global time-monitoring tools** with real-time alerts for atomic/GPS clock faults, improving uptime for distributed time servers, critical for **ultra low-latency systems** and **databases** using **Golang** and **gRPC**.
- Eliminated the time taken to disable a machine manually by automating the process of tracking time disagreements and GPIO errors on servers, using **Python** and **SQL-like queries** with Google's monitoring database.

Google

STEP Intern

Sunnyvale, CA

May 2023 - Aug 2023

- Launched an automated service to programmatically identify malicious URLs hosted on Google Compute Engine virtual machines, yielding a **10x increase** in hostile Compute Engine users addressed and disabled.
- Deployed **C++** services and **Golang** scripts for DNS/IP lookups using **gRPC** and **Protobuf**, and reprovisioned the existing lookup pipeline for the WHOIS service to incur a **99.8% load decrease** by redirecting traffic to a feature cache.

Google

STEP Intern

Waterloo, ON

May 2022 - Aug 2022

- Crafted ingestion pipelines and a new feature generator for an ML-based bug triaging tool that reduces SWE hours spent debugging integration tests by dynamically routing bugs, improving clustering accuracy by **21%**.
- Developed a **custom Go server** using **Pub/Sub queues**, **gRPC**, and **Protobuf** for inter-process communication, minimizing latency to **3700ms** (75% decrease from previous) per workflow execution.

Skills

Languages Python, C++, C, Go, Java, SQL, GraphQL

Technologies PostgreSQL, MongoDB, MySQL, gRPC, Protobuf, Git, Linux/Unix, Docker, Kubernetes

Services AWS (EC2, S3, DynamoDB), Google Cloud (Storage, Compute, Firebase), Apache (Cassandra, Kafka, Spark)