Fei Lin

f34lin@uwaterloo.ca | github.com/Flin42 | linkedin.com/in/feilin/

Education

University of Waterloo *2020 - 2025*
Bachelors of Software Engineering *90.2% CGPA*
Courses: Algorithms, Data Structures, Databases, Operating Systems, Concurrency, Networks, User Interfaces

Skills

Languages: Python, C++, Go, TypeScript/JavaScript, SQL, Scala, Java, Kotlin
Technologies: Docker, Spark, GCP, AWS, Kubernetes, Node.js, React, Flask, Airflow, Redis, Grafana, MongoDB, gRPC

Experience

Software Engineering Intern *Toronto, Ontario*
Cohere *May 2024 - Present*
- Designed an **asynchronous** query solution using **Python** and **BigQuery** to speedup the retrieval and analysis of a data exceeding **hundreds of TBs** in size, resulting in a **70% reduction** in processing time
- Refactored **Spark** jobs to run on **preemptible / spot VMs**, cutting compute cost for key data pipelines stages by over **60%** with near-zero performance loss
- Implemented a **Depth First Search** algorithm to detect and group individual data points into connected components to capture all relevant context in a dataset, resulting in higher quality **LLM** pretraining documents
- Optimized performance of **Dataproc clusters** with over **1000+ vCPUs** used for distributed workloads by fine-tuning instance types and resources to achieve improved processing times and cost-effectiveness
- Wrote a **Python** job to download and transfer data into **GCS**, using **multithreading** and **buffering** to achieve a throughput of over **1 GB/s**

Software Engineering Intern *Toronto, Ontario*
Intact Financial Corporation *Jan 2024 - Apr 2024*
- Added **Kinesis Firehose** to a search service for analytics ingestion and storage into **S3**, reducing data latency by **40%** and enhancing data processing efficiency
- Refactored a **Python** microservice to read from **Kafka** for event-driven processing, resulting in significant reductions to downstream errors rates
- Optimized a rule-based test case generator by eliminating redundant test paths, speeding up system tests by **20%**

Software Engineering Intern *Toronto, Ontario*
Swan Cloud Inc. *May 2023 - Aug 2023*
- Developed **Flask** APIs in **Python** for the orchestration and monitoring of computing jobs assigned to distributed computing nodes, ensuring robust communication and coordination
- Optimized the download of files by parallelizing **REST API** requests, reducing total download time by over **3x**
- Used **Redis** to cache tasks in a workload scheduler, reducing excess database queries and improving matching times

Software Engineering Intern *Toronto, Ontario*
Index Exchange *Sep 2022 - Dec 2022*
- Worked on a sampling solution in **Go** to efficiently collect and measure metrics from **millions** of ad impressions
- Developed data pipelines in **Scala** and **Spark** to write **millions** of billing records into a **Vertica** data warehouse
- Utilized **Go** and **Gin** to create API endpoints for retrieving billing info and payment history for hundreds of accounts

Data Scientist *Toronto, Ontario*
Wish *Jan 2022 - Apr 2022*
- Wrote an automated reporting job on delivery metrics for core business teams using **Python**, **SQL**, and **Airflow**
- Created **ETL** pipelines that continuously ingested live raw data to update dynamic datasets with **millions** of records

Projects

Sundial
- AI scheduler that automated calendar generation, made using **Go**, **TypeScript**, **React** and hosted on **AWS**
- Integrated Cohere's **Command R** API to generate calendars for users after parsing their provided PDF schedules
- Created pages and UI components with **React**, **Tailwind**, and **TypeScript**