

Yiqing (William) Sun

Available: May 11, 2026 – Aug 21, 2026 (14 weeks)

+1 (412) 616-5461 ✉ yiqings2@andrew.cmu.edu [linkedin.com/in/williamodinson](https://www.linkedin.com/in/williamodinson) github.com/williamodinson

Education

Carnegie Mellon University

Master of Information Systems Management, BIDA Track.
Heinz College Dean's List for Fall 2025. (GPA: 3.89/4.0)

Pittsburgh, PA

August 2025 – December 2026

Tianjin University

Bachelor's, Information Management and Information Systems.

Tianjin, China

September 2021 – July 2025

Professional Experience

Carnegie Mellon University, School of Computer Science

Teaching Assistant for 15-619 Cloud Computing S26

Pittsburgh, PA

December 2025 – Present

- Designed the semester project for 130+ students: a high-performance tweet recommendation service requiring gRPC interservice communication, TB-scale ETL on GCP Dataproc, and production deployment on AWS EKS with Aurora and ElastiCache.
- Held weekly office hours mentoring students on cloud system design, Kubernetes deployment, and performance debugging.

Shanghai Mingchen Mould & Plastic Technology Co., Ltd

Software Development Intern

Shanghai, China

July 2024 – October 2024

- Conducted requirements analysis across 3 business workflows; derived a unified database schema (40+ fields) with data integrity constraints to support the company's digital transformation, and collaborated with cross-functional teams to build backend services in JavaScript for CRUD, visualization, and platform integration.
- Built an automated migration pipeline to transform and load 3,000+ legacy records from scattered spreadsheets and manual documents into the new schema, handling data cleaning, format normalization, and edge-case validation.

Projects

High-Performance Blockchain Mining Service

Team Project

Pittsburgh, PA

September 2025 – December 2025

- Engineered a blockchain microservice in Golang, deployed on AWS EKS; optimized for maximum throughput per m6g.large instance (2 vCPU) under a fixed cost budget, sustaining 9,000+ RPS in 30-minute stress tests.
- Conducted iterative profiling with pprof across the full request path; identified per-request zlib allocations as the highest-impact bottleneck and implemented sync.Pool-based compressor reuse, cutting heap allocations by ~40% and P99 latency by 34 ms.
- Containerized services with Docker, orchestrated deployments on Kubernetes via Terraform (EKS, ALB), and automated CI/CD with GitHub Actions (ECR image builds, Helm rollouts), enabling reproducible environments across dev and production.

Gastronome: Business Recommendation Platform

Sole Developer, [Repository](#)

Tianjin, China

January 2025 – July 2025

- Built a business recommendation platform on the Yelp Open Dataset (150K businesses, 7M reviews) with Django, PostgreSQL, Redis, OpenSearch; automated CI/CD with GitHub Actions; load-tested with Locust at 200 concurrent users (P95 < 100 ms).
- Generated business similarity embeddings with GNN and indexed them in FAISS for nearest-neighbor retrieval; designed a multi-tier fallback strategy (personalized scores, similar businesses, regional hotlist) so the system always returns relevant results regardless of user history.
- Decoupled model serving from the Django request path via a gRPC service, and ran dedicated Celery workers via RabbitMQ for async inference and background tasks; monitored with Sentry.

Twitter Graph Analytics and PageRank Computation

Individual Project

Pittsburgh, PA

September 2025 – October 2025

- Implemented PageRank on a Twitter follower graph (Spark DataFrame API, Scala); optimized data skew through user-based repartitioning and reduced redundant shuffles by pre-joining static tables, cutting end-to-end runtime from 1h 20m to 20 min.
- Benchmarked on AWS EMR vs Azure Databricks; Databricks halved runtime via Adaptive Query Execution optimizing joins.

Skills

Languages: Go (Golang), Python, Java, Scala, SQL, JavaScript, HTML/CSS

Technologies: Django, Gin, gRPC, PostgreSQL, Redis, RabbitMQ, OpenSearch, Apache Spark, Celery, Docker, Kubernetes, Terraform, AWS (EKS, ECR, RDS, ElastiCache, EMR), GCP (Dataproc), FAISS, GitHub Actions, Locust, Sentry, Linux, Git