

Qiyang He

☎ (765)337-0786 | @ qiyanghe1998@outlook.com | 🔗 LinkedIn | 📄 Google Scholar | 🏠 Homepage

EDUCATION

Purdue University

M.Sc. in Computer Science; GPA: 3.91/4.00

West Lafayette, IN, USA

Jan. 2021 – Dec. 2023

Southern University of Science and Technology (SUSTech)

B.Eng. in Computer Science and Engineering; GPA: 3.86/4.00; Rank: 3/146

Shenzhen, Guangdong, China

Sep. 2016 – Jul. 2020

RESEARCH INTERESTS

Database Systems; Query Optimization; Data Management; Distributed Systems; Data Infrastructure; Metadata Management; LLM Systems

EXPERIENCE

Snowflake – Metadata Team, Software Engineer

Feb. 2024 – Present

Worked on scalable metadata management systems for petabyte-scale analytical queries

- **Led metadata file clustering**, enabling compact storage of metadata based on customer-defined clustering keys. This significantly improved pruning efficiency and reduced query compilation time by **30%** on targeted workloads; the feature was adopted by major customers including **OpenAI**.
- **Built adaptive metadata file clustering**, a background optimization that intelligently balances performance and resource usage, cutting resource consumption by **40%** while maintaining similar pruning benefits.
- **Developed a hidden-table testing framework** to safely validate metadata/data-change features against live workloads, enabling faster, lower-risk releases.
- **Shipped new metadata APIs** that support approximate metadata requests, allowing the compiler to bypass precision-heavy lookups when safe; improved planning/compile latency and throughput on large **Iceberg** workloads.

Cockroach Labs – SQL Query Team, Research Intern

May 2023 – Aug. 2023

- **Developed a workload-aware index recommender** that analyzes real query traces to suggest optimal indexes. Delivered **5× throughput improvement** on analytical workloads (TPC-H) and **15% latency reduction** on transactional workloads (TPC-C) over manually tuned baselines.
- **Designed a cost-based index merging algorithm**, consolidating overlapping indexes to reduce storage footprint and ongoing maintenance overhead.

Pinterest – Distributed Storage System Team, Research Intern

May 2022 – Aug. 2022

- **Built a space-efficient incremental backup** using Multi-Paxos that stored only delta updates instead of full files, achieving a **10× speedup** in backup and restore performance.

Purdue University – Programming language Group, Research Assistant

Jan. 2021 – May 2023

- **Developed a C code generation backend for SparkSQL** using (LMS)-based generative programming to support incremental view maintenance.
- Identified inefficiencies in DBToaster and **designed tree- and hash-based index structures**, improving correlated nested aggregate query performance by up to **1000×**

SUSTech – Database Group, Research Assistant

Dec. 2019 – Jul. 2020

- **Built a high-dimensional vector search engine** using product/residual vector quantization; implemented custom CUDA kernels for batched inner-product computation and on-GPU sorting to generate top-k results, delivering low-latency, high-throughput retrieval while reducing memory footprint vs. exact search.

PUBLICATIONS

SIGMOD 2022

Efficient Incrementalization of Correlated Nested Aggregate Queries using Relative Partial Aggregate Indexes (RPAI)

OOPSLA 2021

Reachability types: tracking aliasing and separation in higher-order functional programs

ACM CSUR 2025

A Survey of Learned Indexes for the Multi-dimensional Space

SKILLS & ACHIEVEMENTS

Competitive Programming: ACM-ICPC Regional (Gold/Silver); China National Olympiad in Informatics (Bronze)

Programming Languages: C++, Python, Go, Java, C, Scala, Rust, SQL, Bash.

Systems & Infra: Linux, Git, AWS S3

Data & DB Systems: Metadata management; query optimization; incremental view maintenance; indexing; distributed data systems; RocksDB; SparkSQL; Apache Iceberg