

Zhiru Zhu

CONTACT INFORMATION

E-mail: zhiru@uchicago.edu

Homepage: <https://zhiru.netlify.app>

EDUCATION

University of Chicago, Chicago, IL, USA

PhD student in Computer Science

September 2020 - Present

- Advisor: Prof. Raul Castro Fernandez
- Research area: Data management and privacy

Franklin & Marshall College, Lancaster, PA, USA

B.A., Computer Science, Magna Cum Laude

August 2013 - May 2017

- GPA: 3.77 / 4.0. Phi Beta Kappa Honor Society

University of Edinburgh, Edinburgh, UK

Visiting student in School of Informatics

September 2015 - May 2016

- GPA: 3.73 / 4.0. IFSA-Butler Academic Achievement List

RESEARCH EXPERIENCE

University of Chicago, Chicago, IL, USA

September 2020 - Present

- Currently working on project *Towards Easing the Deployment of Differential Privacy* (Paper in submission).
- Worked on a data discovery system, *Ver*, that identifies project-join views over large repositories of tables that do not contain join path information. This work is published in ICDE 2023.
- Worked on *Data Station*, a data escrow that enables the formation of data-sharing consortia through delegated, trustworthy, and auditable computation. This work is published in VLDB 2022.

New York University, New York, NY, USA

July 2017 - May 2018

- Advisor: Prof. Azza Abouzied, Prof. Alexandra Meliou (UMass Amherst)
- Research on techniques for efficient evaluation of incremental package query that satisfies constraints collectively in a large-scale database system.

Franklin & Marshall College, Lancaster, PA, USA

May 2016 - May 2017

- Advisor: Prof. Erin Talvitie
- Research on model-based reinforcement learning algorithms in the Atari 2600 domain and explored sparse methods to address scaling challenges.

WORKING EXPERIENCE

ZILLIZ, Shanghai, China

Software Engineer

February 2019 - May 2020

- Core developer of the open-source vector database, Milvus (<https://github.com/milvus-io/milvus>). Designed and implemented multiple components of the system, including query indexing methods, memory manager, meta manager, SDK, tracing, gRPC server, and overall program compilation process.
- Developed query optimizer of a GPU-accelerated OLAP database based on PostgreSQL and increased query performance by two orders of magnitude over CPU database.

PUBLICATIONS

Yue Gong*, **Zhiru Zhu***, Sainyam Galhotra, and Raul Castro Fernandez. *Ver: View Discovery in the Wild*. **ICDE 2023** (* equal contribution)

Siyuan Xia, **Zhiru Zhu**, Chris Zhu, Jinjin Zhao, Kyle Chard, Aaron Elmore, Ian Foster, Michael Franklin, Sanjay Krishnan, Raul Castro Fernandez. *Data Station: Delegated, Trustworthy, and Auditable Computation to Enable Data-Sharing Consortia with a Data Escrow*. **VLDB 2022**