

# Rui Liu

Department of Computer Science  
University of Chicago  
E-mail: [rui Liu@cs.uchicago.edu](mailto:rui Liu@cs.uchicago.edu) || Web: [csrui Liu.github.io](http://csrui Liu.github.io)

## RESEARCH INTERESTS

---

My research interests are grounded in data systems and bridge to other important areas, currently centering on building and optimizing data-intensive systems for various applications, such as data science and artificial intelligence. My previous research focused on building systems for data discovery, blockchain, mobile and ubiquitous computing.

## EDUCATION

---

- Ph.D., Computer Science** Sept 2017 – Dec 2023
- University of Chicago, with full scholarship
  - Advised by Aaron J. Elmore and Michael J. Franklin
  - Close collaboration with Sanjay Krishnan
- M.Phil., Computer Science** Oct 2012 – Dec 2015
- Hong Kong Polytechnic University, with full scholarship
  - Advised by Jiannong Cao
- B.Eng., Software Engineering** Sept 2008 – June 2012
- Northeastern University (China), Graduation with Honors

## PROFESSIONAL EXPERIENCE

---

- Microsoft** June 2022 – Sept 2022
- Research Intern at Gray Systems Lab
  - Working on feature engineering optimization for machine learning workloads using data provenance
  - Mentors: Kwanghyun Park, Jesús Camacho Rodríguez, Yuanyuan Tian
- DocuSign** July 2021 – Sept 2021
- Data Science Intern at AI Engineering Team
  - Working on machine learning infrastructure and optimizing the training pipeline
  - Mentor: David Wong
- National University of Singapore** Aug 2016 – July 2017
- Research Assistant at Data System Research Group
  - Working on blockchain, data integration, distributed data management
  - Mentor: Beng Chin Ooi
- Chinese University of Hong Kong** Mar 2016 – Aug 2016
- Research Assistant at Department of Information Engineering
  - Working on data sharing in mobile sensing systems
  - Mentor: Kehuan Zhang

## PROJECTS

---

Selected projects in reverse-chronological order.

### 1. Resource-adaptive Query Execution for Cloud-native Databases

**Role:** Leader. **Affiliation:** University of Chicago.

- ▶ Designing and developing a resource-adaptive query execution framework, Riveter, to enable adaptive query suspension and resumption at various granularities when resources become unavailable or impractical during execution.
- ▶ We built a prototype system that can support various suspension strategies: redo strategy, pipeline-level strategy (by modifying DuckDB), and process-level strategy (by building on top of a Linux checkpoint and restore tool).
- ▶ We devised a cost model to determine the suspension strategy adaptively and evaluated Riveter using TPC-H benchmark to show the adaptivity and effectiveness of Riveter.

### 2. Optimizing Data Pipelines for Machine Learning in Feature Stores

**Role:** Core Member. **Affiliation:** Gray Systems Lab@Microsoft.

- ▶ Proposing and developing optimizations for machine learning data pipelines in feature stores, including feature computation reuse, incremental feature refresh, scheduling of feature maintenance, and compute engine selection.
- ▶ I implemented the optimizations targeted in point-in-time join by modifying Feathr, a state-of-the-art open-source feature store widely used in production at LinkedIn. I also evaluated them on use cases from the TPCx-AI benchmark and real-world online retail applications. Our experimental analysis shows that our optimizations can accelerate data pipelines by up to 3x over state-of-the-art baselines. The prototype system is under further development and potential integration with Microsoft Azure.

### 3. Resource-arbitration for Progressive Iterative Analytics

**Role:** Leader. **Affiliation:** University of Chicago.

- ▶ Proposing a resource arbitration framework for progressive iterative analytic jobs to continuously prioritize the jobs and determine if/when to preempt and reallocate resources.
- ▶ We implemented two resource arbitration systems: (1) Rotary-AQP, based on the modification of Apache Spark, for approximate query processing, and (2) Rotary-DLT, an extension to TensorFlow, for deep learning training.
- ▶ We evaluated the two systems using the TPC-H benchmark and a survey-based deep learning workload, respectively. The evaluation results show that both significantly outperform the state-of-the-art and other commonly used baselines in terms of various metrics such as the number of attained jobs in a workload and attainment progress of each job.

### 4. Accelerating Container-based Deep Learning Hyperparameter Optimization Workloads

**Role:** Core Member. **Affiliation:** AI Engineering Team@DocuSign.

- ▶ Designing and developing, Relish, a system specifically to accelerate Hyperparameter Optimization (HPO) workloads in container-based environments (e.g., Docker).
- ▶ I implemented an end-to-end system to realize Relish. It can efficiently segment long-running and memory-intensive HPO jobs into multiple sub-jobs in terms of trials and assign them to various devices for efficient parallel execution. I evaluate Relish using a synthetic workload derived from a three-month trace of a multi-tenant cluster at DocuSign. The results show that Relish can significantly improve GPU utilization and accelerate the workload.

### 5. Optimizing Neural Network Training for Hyper-Parameter Tuning

**Role:** Leader. **Affiliation:** University of Chicago.

- ▶ We built a prototype system on top of TensorFlow for sharing training resources among deep learning models with a cost model for balancing the benefits and limitations of packing models for simultaneous training.
- ▶ We conducted an ablation study to identify some trade-offs when sharing limited resources for training multiple neural network models. We also extended Hyperband, an extensively employed hyper-parameters tuning algorithm, and accelerated the performance up to 2.7x through the packing operation.

### 6. Draining the Data Swamp

**Role:** Core Member. **Affiliation:** University of Chicago.

- ▶ Building data discovery and management tools for data swamps where data lakes have overwhelming heterogeneous and continuous data, lax data model design, poor metadata management, and dynamic user requirements.
- ▶ I built a schema completion tool by re-implementing Google's WebTables to identify data with similar schemas in two large datasets—Google WebTables dataset consisting of 1B+ data records and a CDIAC dataset consisting of 0.5M files

that contain environmental science data. The results show that the schema completion tool cannot efficiently identify the correct synonyms, even in relaxed cases, such as Top-20. The tool is subsequently integrated into our proposed framework of data similarity that combines file similarity specification and human-in-the-loop interaction.

## 7. Benchmarking Private Blockchains

**Role:** Core Member. **Affiliation:** National University of Singapore.

- ▶ Building BlockBench, the world-first framework for evaluating the performance of private blockchains.
- ▶ I built the Hyperledger Fabric (v0.6.0 and v1.0.0) in a private network, implementing APIs and evaluating the performance. I also conduct a comprehensive technical survey on private blockchains, focusing on Ethereum and Hyperledger.

## PUBLICATIONS

---

### Conference Papers

- [1] **Rui Liu**, Aaron J. Elmore, Michael J. Franklin, Sanjay Krishnan, “Riveter: Adaptive Query Suspension and Resumption Framework for Cloud Native Databases”. IEEE International Conference on Data Engineering (**ICDE**), Utrecht, Netherlands. 2024.
- [2] **Rui Liu**, Jun Hyuk Chang, Riki Otaki, Zhe Heng Eng, Aaron J. Elmore, Michael J. Franklin, Sanjay Krishnan, “Towards Resource-adaptive Query Execution in Cloud Native Databases”, Conference on Innovative Data Systems Research (**CIDR**), Santa Cruz, USA. 2024.
- [3] **Rui Liu**, Kwanghyun Park, Fotis Psallidas, Xiaoyong Zhu, Jinghui Mo, Rathijit Sen, Matteo Interlandi, Konstantinos Karanasos, Yuanyuan Tian, Jesús Camacho-Rodríguez, “Optimizing Data Pipelines for Machine Learning in Feature Stores”, VLDB Endowment, vol. 16, no. 13, pp. 4230–4239. Presentation at International Conference on Very Large Data Bases (**VLDB**), Guangzhou, China. 2024.
- [4] **Rui Liu**, Aaron J. Elmore, Michael J. Franklin, Sanjay Krishnan, “Rotary: A Resource Arbitration Framework for Progressive Iterative Analytics”, IEEE International Conference on Data Engineering (**ICDE**), Anaheim, USA. 2023.
- [5] Minzhao Liu, Junyu Liu, **Rui Liu**, Henry Makhanov, Danylo Lykov, Anuj Apte, Yuri Alexeev, “Embedding Learning in Hybrid Quantum-Classical Neural Networks”, IEEE International Conference on Quantum Computing and Engineering (**QCE**), Broomfield, USA. 2022.
- [6] **Rui Liu**, David Wong, Dave Lange, Patrik Larsson, Vinay Jethava, Qing Zheng, “Accelerating Container-based Deep Learning Hyperparameter Optimization Workloads”, Workshop on Data Management for End-to-End Machine Learning (**DEEM@SIGMOD**), Philadelphia, USA. 2022.
- [7] Heng Zhang, Lingda Li, Hang Liu, Donglin Zhuang, **Rui Liu**, Chengying Huan, Shuang Song, Dingwen Tao, Yongchao Liu, Charles He, Yanjun Wu, Shuaiwen Leon Song, “Bring Orders into Uncertainty: Enabling Efficient Uncertain Graph Processing via Novel Path Sampling on Multi-Accelerator System”, ACM International Conference on Supercomputing (**ICS**), Virtual. 2022.
- [8] **Rui Liu**, Sanjay Krishnan, Aaron J. Elmore, Michael J. Franklin, “Understanding and Optimizing Packed Neural Network Training for Hyper-Parameter Tuning”, Workshop on Data Management for End-to-End Machine Learning (**DEEM@SIGMOD**), Virtual. 2021.
- [9] Will Brackenbury, **Rui Liu**, Mainack Mondal, Aaron J. Elmore, Blase Ur, Kyle Chard, Michael J. Franklin, “Draining the Data Swamp: A Similarity-based Approach”, Workshop on Human-In-the-Loop Data Analytics (**HILDA@SIGMOD**), Houston, USA. 2018.
- [10] Tien Tuan Anh Dinh, Ji Wang, Gang Chen, **Rui Liu**, Beng Chin Ooi, Kian-Lee Tan, “BLOCKBENCH: A Framework for Analyzing Private Blockchains”, ACM International Conference on Management of Data (**SIGMOD**), Chicago, USA. 2017.
- [11] Zhe Zhou, Wenrui Diao, Xiangyu Liu, Zhou Li, Kehuan Zhang, **Rui Liu**, “Vulnerable GPU Memory Management: Towards Recovering Raw Data from GPU”, Privacy Enhancing Technologies Symposium (**PETS**), Minneapolis, USA. 2017.
- [12] **Rui Liu**, Jiannong Cao, Sebastian VanSyckel, Wenyu Gao, “PriMe: Human-centric Privacy Measurement based on User Preferences towards Data Sharing in Mobile Participatory Sensing Systems”, IEEE International Conference on Pervasive Computing and Communications (**PerCom**), Sydney, Australia. 2016.

- [13] **Rui Liu**, Jiannong Cao, Lei Yang, Kehuan Zhang, “PriWe: Recommendation for Privacy Settings of Mobile Apps Based on Crowdsourced Users’ Expectations”, IEEE International Conference on Mobile Services (**MS**), New York, USA. 2015.

### Journal Articles

- [1] Sanjay Krishnan, Aaron J. Elmore, Michael J. Franklin, John Paparrizos, Zechao Shang, Adam Dziedzic, **Rui Liu**, “Artificial Intelligence in Resource-Constrained and Shared Environments”, ACM SIGOPS Operating Systems Review (**OSR**), vol. 53, no. 1, pp. 1-6, 2019.
- [2] Ruiyun Yu, Jiannong Cao, **Rui Liu**, Wenyu Gao, Xingwei Wang, Junbin Liang, “Participant Incentive Mechanism towards Quality-Oriented Sensing: Understanding and Application”, ACM Transactions on Sensor Networks (**TOSN**), vol. 15, no. 2, pp. 21:1-21:25, 2019.
- [3] **Rui Liu**, Junbin Liang, Jiannong Cao, Kehuan Zhang, Wenyu Gao, Lei Yang, Ruiyun Yu, “Understanding Mobile Users’ Privacy Expectations: A Recommendation-based Method through Crowdsourcing”, IEEE Transactions on Services Computing (**TSC**), vol. 12, no. 2, pp. 304-318, 2019.
- [4] Tien Tuan Anh Dinh, **Rui Liu**, Meihui Zhang, Gang Chen, Beng Chin Ooi, Ji Wang, “Untangling Blockchain: A Data Processing View of Blockchain Systems”, IEEE Transactions on Knowledge and Data Engineering (**TKDE**), vol. 30, no. 7, pp. 1366–1385, 2018.
- [5] **Rui Liu**, Junbin Liang, Jiannong Cao, Weiping Zhu, Sebastian Vansyckel, Christian Becker, “UIO-based Testbed Augmentation for Simulating Cyber-Physical Systems”, IEEE Intelligent Systems, vol. 33, no. 6, pp. 69–86, 2018.
- [6] **Rui Liu**, Jiannong Cao, Kehuan Zhang, Wenyu Gao, Lei Yang, Junbin Liang, “When Privacy Meets Usability: Unobtrusive Privacy Permission Recommendation System for Mobile Apps based on Crowdsourcing”, IEEE Transactions on Services Computing (**TSC**), vol. 11, no. 5, pp. 864–878, 2018.
- [7] Wenrui Diao, **Rui Liu**, Xiangyu Liu, Zhe Zhou, Zhou Li, Kehuan Zhang, “Accessing Mobile User’s Privacy based on IME Personalization: Understanding and Practical Attacks”, Journal of Computer Security (**JCS**), vol. 26, no. 3, pp. 283–309, 2018.
- [8] Junbin Liang, Jiannong Cao, **Rui Liu**, Tao Li, “Distributed Intelligent MEMS: A Survey and a Real-time Programming Framework”, ACM Computing Surveys (**CSUR**), vol. 49, no. 1, pp. 20:1–20:29, 2016.
- [9] Ruiyun Yu, **Rui Liu**, Xingwei Wang, Jiannong Cao, “Improving Data Quality with an Accumulated Reputation Model in Participatory Sensing Systems”, **Sensors**, vol. 14, no. 3, pp. 5573-5594, 2014.

### Demonstrations / Posters

- [1] Heng Zhang, Lingda Li, Donglin Zhuang, **Rui Liu**, Shuang Song, Dingwen Tao, Yanjun Wu, Shuaiwen Leon Song, “An Efficient Uncertain Graph Processing Framework for Heterogeneous Architectures”, ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (**PPoPP**), Virtual, 2021. [Poster]
- [2] Shan Jiang, Junbin Liang, Jiannong Cao, **Rui Liu**, “An Ensemble-level Programming Model with Real-Time Support for Multi-Robot Systems”, IEEE International Conference on Pervasive Computing and Communications (**PerCom**), Sydney, Australia. 2016. [Demo]

### Patents

- [1] Jesus Camacho Rodriguez, Kwanghyun Park, **Rui Liu**, Fotis Psallidas, Xiaoyong Zhu, Jinghui Mo, Rathijit Sen, Matteo Interlandi, Konstantinos Karanasos, Yuanyuan Tian, “Feature Store Data Preparation Optimization”, filed for **US Patent**, 2023.

### Miscellaneous

- [1] Yao Lu, Song Bian, Lequn Chen, Yongjun He, Yulong Hui, Matthew Lentz, Beibin Li, Fei Liu, Jialin Li, Qi Liu, **Rui Liu**, Xiaoxuan Liu, Lin Ma, Kexin Rong, Jianguo Wang, Yingjun Wu, Yongji Wu, Huanchen Zhang, Minjia Zhang, Qizhen Zhang, Tianyi Zhou, Danyang Zhuo, “Computing in the Era of Large Generative Models: From Cloud-Native to AI-Native”, arXiv Preprint, URL: <https://arxiv.org/pdf/2401.12230.pdf>.

## RECOGNITION

---

- ⚡ Travel Award of ICDE, 2023
- ⚡ CERES Graduate Scholar Fellowship, UChicago, 2017
- ⚡ Second Runner-up, Oracle Java Programming Contest in China, 2015
- ⚡ Honorable Mention, Research Competition of COMP@HK PolyU, 2013
- ⚡ Outstanding Graduate of Liaoning Province of China (top 1%), 2012
- ⚡ Google Excellence Scholarship, 2011
- ⚡ First Class Scholarship, Northeastern University (China), 2011
- ⚡ First Prize, China National Undergraduate Mathematical Contest in Modeling, 2010
- ⚡ Second Runner-up, IBM Mainframe Technology Application Contest in China, 2010
- ⚡ Best Mainframe Application, Citi Finance and Information Technology Contest in China, 2009

## TEACHING

---

### TAed Courses

- 2019, 2020, 2021, 2022, 2023: Introduction to Databases, UChicago
- 2018: Introduction to Data Science, UChicago
- 2014: Mobile Computing, HK PolyU
- 2013: Introduction to Internet Computing, HK PolyU
- 2013: Information Technology, HK PolyU

### Invited Lectures

- 2022.02: Introduction to Big Data Platforms, HK PolyU

## MEDIA COVERAGE

---

- 🔊 Disseminate: The Computer Science Research Podcast, ep. 49.
- 🔊 The Morning Paper for Blockbench.

## REFEREES

---

**Aaron J. Elmore** [aelmore@cs.uchicago.edu]  
Associate Professor  
University of Chicago

**Sanjay Krishnan** [skr@cs.uchicago.edu]  
Assistant Professor  
University of Chicago

**Michael J. Franklin** [mjfranklin@uchicago.edu]  
Morton D. Hull Distinguished Service Professor  
University of Chicago

*Last Updated: July 09, 2024*