

Naili Xing

Researcher in Databases & ML Systems | Apache Committer (SINGA) | Singapore PR
(+65) 9807 8455 | xingnaili14@gmail.com | Singapore
[LinkedIn](#) | [Google Scholar](#) | [GitHub](#)



EDUCATION

National University of Singapore — Ph.D. in Computer Science Aug 2021 – Sep 2026 (expected)
• Research focus: in-database structured data analytics; AI×Database. NUS Database Research Lab.

National University of Singapore — M.Sc. in Computer Engineering
University of Electronic Science and Technology of China — B.Eng. in Electrical and Electronics Engineering

AWARDS

Dean's Graduate Research Excellence Award — NUS School of Computing AY 2023/2024, Semester 2
Research Achievement Award — NUS School of Computing AY 2021/2022, Semester 2

PUBLICATIONS

First Author

- [1] **pTNAS: Progressive Neural Architecture Search for Tabular Data**
Proceedings of the 43rd International Conference on Machine Learning (ICML), 2026. [\[Paper\]](#)
- [2] **Database Native Model Selection: Harnessing Deep Neural Networks in Database Systems**
Proceedings of the VLDB Endowment (PVLDB), vol. 17, 2024. [\[Paper\]](#)
- [3] **Singa-Easy: An Easy-to-Use Framework for Multimodal Analysis**
Proceedings of the 29th ACM International Conference on Multimedia (MM), 2021. [\[Paper\]](#)
- [4] **Towards Effective Orchestration of AI×DB Workloads**
Under submission, 2026. [\[Paper\]](#)
- [5] **NQO: Query Optimization as a Learnable Function**
Under submission, 2026.
- [6] **MoQOE: Mixture of Learned Query Optimization Experts for End-to-End Efficiency**
Under submission, 2026.

Corresponding Author

- [7] **NeurBench: A Benchmark Suite for Learned Database Components with Drift Modeling**
Proceedings of the ACM on Management of Data (SIGMOD), 2026. [\[Paper\]](#)
- [8] **NeurIDA: Dynamic Modeling for Effective In-Database Analytics**
Proceedings of the VLDB Endowment (PVLDB), 2026. [\[Paper\]](#)

Second Author

- [9] **Powering In-Database Dynamic Model Slicing for Structured Data Analytics**
Proceedings of the VLDB Endowment (PVLDB), 2024. [\[Paper\]](#)
- [10] **Falcon: A Privacy-Preserving and Interpretable Vertical Federated Learning System**
Proceedings of the VLDB Endowment (PVLDB), 2023. [\[Paper\]](#)
- [11] **LogicTree-RAG: Logic Tree-guided Retrieval-Augmented Generation for Long-form Patent Drafting**
Under submission, 2026.

Co-author

- [12] **NeurDB: On the Design and Implementation of an AI-powered Autonomous Database**
Conference on Innovative Data Systems Research (CIDR), 2025. [\[Paper\]](#)
- [13] **NeurDB: An AI-powered Autonomous Data System**
Science China Information Sciences, 2024. [\[Paper\]](#)

EXPERIENCE

NUS Database Research Lab — Ph.D. Researcher Apr 2021 – Present
• Published at SIGMOD, VLDB, ICML, CIDR, and ACM Multimedia (see Publications).

Mindtranz Pte. Ltd. [Website] — Chief Executive / Technology Officer

Apr 2023 – Present

- Led a 5-person team to build IP8Value, an IP management platform for tech innovation and commercialization.
- Built LLM agent for automated patent drafting with logic-tree-guided RAG and structured generation.
- Architected end-to-end platform spanning business backend, ML/algorithm backend, and cloud infrastructure.

NUS Database Research Lab — Research Assistant

Jan 2020 – Apr 2021

- Designed and deployed scalable AutoML infrastructure on Kubernetes, enabling distributed training for ML services (e.g., Mask R-CNN object detection) and ELK/Kafka/Spark Streaming log monitoring.

Igloo (formerly Axinan) — Software Engineer

Apr 2019 – Dec 2019

- Built resume recommendation system using NLP (word2vec, NER) and deep learning classifiers (CNN, LSTM).

Self-employed — Trader

Jan 2021 – Present

- Trading with own capital; maintained approximately 80% win rate over the past four years.

OPEN-SOURCE CONTRIBUTIONS

NeurDB — AI-Powered Autonomous Data System

- **Core contributor:** Designed and built AI engine; integrated in-database ML operators (PREDICT) into PostgreSQL planner/executor and cost-based scheduling; TabPFN runtime and distributed inference. [\[Code\]](#) (C/Python)

Apache SINGA — Distributed Deep Learning Platform

- **Contributor:** Extended model support and integrated SINGA inference into PostgreSQL. [\[Code\]](#) (C++/Python)

Falcon — Privacy-Preserving Vertical Federated Learning System

- **Core contributor:** Built MapReduce-based coordinator for task scheduling; distributed VFL training/inference; secure algorithms (PHE/secret sharing) for linear/logistic regression and MLP. [\[Code\]](#) (Go/C++)

COOL — Cohort OnLine Analytical Processing System

- **Core contributor:** OLAP query optimization and distributed query processing (ZooKeeper, HDFS). [\[Code\]](#) (Java/Python)

MIT 6.824 — Distributed Systems

- **Personal project:** MapReduce and full Raft (leader election, log replication, persistence) in Go. [\[Code\]](#) (Go)