

ERGUTE BAO

MBZUAI, Abu Dhabi, UAE

baoergute8@gmail.com | <https://erguteb.github.io> | +971 058 573 2696 | +33 07 53 15 00 10

RESEARCH INTEREST

I am interested in establishing *rigorous practices for private and secure AI*: identifying risks in existing systems, formalizing these problems, studying them towards creating practical solutions with rigorous guarantees.

I am particularly interested in the following directions:

- Differential privacy (DP) for database (DB), federated learning (FL), and large language models (LLMs).
- Synthetic data generation with formal privacy and security guarantees
- Robustness of LLMs in adversarial environments.

EDUCATION

National University of Singapore (NUS) 2024

Ph.D. in Computer Science

Advisor: Xiaokui Xiao

The Chinese University of Hong Kong (CUHK) 2018

B.Sc. in Computer Science (First Class Honors) ELITE Stream, Minor in Mathematics.

EXPERIENCES

Postdoc Researcher @ **MBZUAI, Abu Dhabi, UAE** June 2025 - Present

Hosted by Ting Yu

Research Intern @ **Tongyi Lab, Alibaba Group, China** Spring 2025, Summer 2024

Project: Resolving the Tension between Privacy, Computation Resources, and Utility in LLM Fine-tuning.

Hosted by Fei Wei, Yaliang Li, and Bolin Ding.

Research Intern @ **Sea AI Lab, SEA, Singapore** Spring 2023

Project: A Generic Private and Secure Vertical Federated Learning Framework (patent filed).

Hosted by Tianyu Pang and Chao Du.

Research Intern @ **DAMO Academy, Alibaba Group, China** Summer 2022

Project: One-Round Federated Learning Protocol with Differential Privacy (patent filed).

Hosted by Yaliang Li and Bolin Ding.

Summer Student Intern @ **Engineering Faculty, CUHK, Hong Kong** Summer 2017, Summer 2016

Project: Dimensionality Reduction and Similarity Search for High-dimensional Data.

Advised by Jinfeng Li and James Cheng

RECENT WORKS

- **Auditing Apple's DifferentialPrivacy.framework: Implementation Bugs, Misconfigurations, and Practical Risks.** To appear in IEEE Symposium on Security and Privacy (S&P), 2026.

The first systematic analysis to reveal that the closed-sourced data collection framework implemented in Apple's MacOS 14.2 and 15.6 is not as private or secure as they have claimed.

- **Overcoming the Retrieval Barrier: Indirect Prompt Injection in the Wild for LLM Systems.** To appear in USENIX Security Symposium (USENIX Security), 2026.

A mathematical framework and a holistic approach for indirect prompt injection for retrieval-based LLMs, uncovering underestimated privacy and security risks/issues in the existing paradigms.

- **Unlocking the Power of Differentially Private Zeroth-order Optimization for Fine-tuning LLMs.** Appeared in USENIX Security Symposium (**USENIX Security**), 2025.

We present a practical and effective algorithm—compatible with accessible GPUs such as RTX4070—for fine-tuning LLMs while ensuring formal privacy guarantees and strong model utility.

SELECTED AWARDS

- **Third** place in the 2018 Differential Privacy Synthetic Data Challenge held by the National Institute of Standards and Technology of the USA.

In this first world-wide coding contest in applying differential privacy in practice, the goal is to design privacy-preserving data publishing algorithms for the US Census. Our team designed and implemented PrivBayes, a differentially private method for releasing data which circumvents the curse of dimensionality. Our solution was invited to the *Journal of Privacy and Confidentiality* (see also here and here).

- **First** place in the 2020 Privacy Temporal Map Challenge held by the National Institute of Standards and Technology of the USA.

We design a practical algorithm for generating real-world temporal map data while preserving differential privacy. Our solution is based on the Markov random field, which improves over PrivBayes in terms of data utility while resolving consistency issues and maintaining the computation efficiency (see also here and here).

FULL LIST OF PUBLICATIONS

1. Rishav Chourasia, **E. Bao**, Uzair Javaid, and Xiaokui Xiao.
Auditing Apple’s DifferentialPrivacy.framework: Implementation Bugs, Misconfigurations, Practical Risks. IEEE Symposium on Security and Privacy (**S&P**), 2026.
2. Hongyan Chang, **E. Bao**, Xinjian Luo, and Ting Yu.
Overcoming the Retrieval Barrier: Indirect Prompt Injection in the Wild for LLM Systems. USENIX Security Symposium (**USENIX Security**), 2026.
3. Jianxin Wei, **E. Bao**, Xiaokui Xiao, and Ting Yu.
SaGD: A Node-Level Private Graph Learning Framework with Sensitivity-Aware Gradient Descent. The Web Conference (**WWW**), 2026.
4. Yangfan Jiang, Fei Wei, **E. Bao**, Yaliang Li, Bolin Ding, Yin Yang, and Xiaokui Xiao.
Accurate Table Question Answering with Accessible LLMs. IEEE International Conference on Data Engineering (**ICDE**), 2026.
5. **E. Bao**, Yangfan Jiang, Fei Wei, Xiaokui Xiao, Zitao Li, Yaliang Li, and Bolin Ding.
Unlocking the Power of Differentially Private Zeroth-order Optimization for Fine-tuning LLMs. USENIX Security Symposium (**USENIX Security**), 2025.
6. **E. Bao**, Fei Wei, Xiaokui Xiao, Yin Yang, Tianyu Pang, and Chao Du.
Towards Learning on Vertically Partitioned Data with Distributed Differential Privacy. IEEE International Conference on Data Engineering (**ICDE**), 2025.
7. Jianxin Wei, Yizheng Zhu, Xiaokui Xiao, **E. Bao**, Yin Yang, Kuntai Cai, and Beng Chin Ooi.
GCON: Differentially Private Graph Convolutional Network via Objective Perturbation. IEEE International Conference on Data Engineering (**ICDE**), 2025.
8. Fei Wei, **E. Bao**, Xiaokui Xiao, Yin Yang, and Bolin Ding.
AAA: an Adaptive Mechanism for Locally Differential Private Mean Estimation. International Conference on Very Large Data Bases (**VLDB**), 2024.

9. **E. Bao**, Dawei Gao, Xiaokui Xiao, Yaliang Li.
Communication Efficient and Differentially Private Logistic Regression under the Distributed Setting.
ACM SIGKDD Conference on Knowledge Discovery and Data Mining (**KDD**), 2023.
10. Jianxin Wei, **E. Bao**, Xiaokui Xiao, Yin Yang.
DPIS: an Enhanced Mechanism for Differentially Private SGD with Importance Sampling.
ACM SIGSAC Conference on Computer and Communications Security (**CCS**), 2022.
11. **E. Bao**, Yizheng Zhu, Xiaokui Xiao, Yin Yang, Beng Chin Ooi, Benjamin H.M. Tan, Khin M.M. Aung.
Skellam Mixture Mechanism: a Novel Approach to Federated Learning with Differential Privacy.
International Conference on Very Large Data Bases (**VLDB**), 2022.
12. **E. Bao**, Yin Yang, Xiaokui Xiao, and Bolin Ding.
CGM: An Enhanced Mechanism for Streaming Data Collection with Local Differential Privacy
International Conference on Very Large Data Bases (**VLDB**), 2021.
13. **E. Bao**, Xiaokui Xiao, Jun Zhao, Dongping Zhang and Bolin Ding.
Synthetic Data Generation with Differential Privacy via Bayesian Networks
Journal of Privacy and Confidentiality (**JPC**), 2021, 11(3).
Invited paper, based on our solution for 2018 NIST DP challenge.

ACADEMIC SERVICES

Program committee member for the following conferences:

- ACM International Conference on Management of Data (SIGMOD): 2027
- International Conference on Very Large Data Bases (VLDB): 2026 2027
- International Conference on Data Engineering (ICDE): 2026
- International Conference on Database Systems for Advanced Applications (DASFAA): 2023-2024 (Demo) 2026 (Demo).

Reviewer for the following journals:

- IEEE Transactions on Knowledge and Data Engineering (TKDE)
- The International Journal on Very Large Data Bases (VLDBJ)
- ACM Transactions on Knowledge Discovery from Data (TKDD)
- IEEE Transactions on Big Data (TBD)
- IEEE Transactions on Information Forensics and Security (TIFS)
- IEEE Transactions on Dependable and Secure Computing (TDSC)
- IEEE Transactions on Privacy (ToP)
- Springer International Journal of Information Security (IJIS)

INVITED TALKS

The Theory and Practice of Differential Privacy in Distributed Settings. Tongyi Lab, Alibaba Group (2025).
Additive Discrete Noise Mechanisms for Differential Privacy. Privacy reading group, Michigan State University (2024).

The Skellam Mechanism for Distributed Differential Privacy. Privacy Innovation Lab, TikTok (2023).