

# ZHANG Xinyi

DLB 625G, 34 Renfrew Rd, HKBU, Hong Kong  
☎ +852-56027517 · ✉ csxyzhang@comp.hkbu.edu.hk

Final year PhD candidate in the Department of Computer Science at Hong Kong Baptist University with experience in designing and implementing various high-performance, secure database and distribution-aware data structures.

**EDUCATION**      **Hong Kong Baptist University**      Hong Kong, China  
Ph.D. in Computer Science      *Sep 2021 – Present*  
Supervisor: Prof. Xu Jianliang

**Huazhong University of Science and Technology**      Wuhan, China  
Bachelor of Engineering in Electronics & Information Engineering      *Sep 2017 – Jun 2021*  
GPA: 3.7/4.0 (Rank: 15%)

**PROFESSIONAL EXPERIENCE**      **Nanyang Technological University**      Singapore  
**Visiting Scholar**      *March 2026 – Present*

- Supervisor: Prof. Wang Qichen
- Focus on "Learned Indexes for Query Optimization," specifically aiming to integrate learned index structures into structure-guided execution plans for acyclic queries.

**Hong Kong Baptist University**      Hong Kong, China  
**Ph.D. Candidate**      *Sep 2021 – Present*

- Supervisor: Prof. Xu Jianliang
- Designed novel algorithms and indexes for relational and vector databases to support efficient and secure query processing in a wide range of systems.
- Designed novel distribution-aware data structures to handle complex database workloads while ensuring robust performance.
- Research resulted in publications at top-tier conferences.

**ByteDance**      Beijing, China  
**Intern**      *Apr 2020 – Oct 2020*

- Optimized the inter-cluster communication for internal database clusters of ByteDance to enhance throughput and reduce latency.

**Huazhong University of Science and Technology**      Wuhan, China  
**Researcher**      *May 2019 – Feb 2020*

- Design a fast consensus algorithm of blockchain combining SDN, dPOS and PBFT consensus algorithm.
- Build a practical blockchain system for ASTRI as a integrity-assured record for their public distributed database.

**RESEARCH INTERESTS**

- Oblivious query processing for data federation and outsourcing cloud computing.
- Scalable secure query processing for vector database based on Trusted Execution Environments.
- Robust updatable learned index under real-world workload.
- Distributed vector database design

**PUBLICATIONS**

1. **X. Zhang**, L. Liang, A. Anastasia, and J. Xu, "HIRE: A Hybrid Learned Index for Robust and Efficient Performance under Mixed Workloads," *Proceedings of the ACM on Management of Data (SIGMOD '26)*, vol. 4, no. 1, Feb. 2026, Full Paper.
2. R. Tang, X. Zhu, **X. Zhang**, L. Chen, X. Li, M. Yuan, and J. Xu, "EDGE: DBMS-Empowered Boolean Decomposition for GIG Synthesis," in *ACM/IEEE Design Automation Conference (DAC '25)*, IEEE, 2025, pp. 1–7, Full Paper.
3. **X. Zhang**, Q. Wang, C. Xu, Y. Peng, and J. Xu, "FedKNN: Secure Federated k-Nearest Neighbor Search," *Proceedings of the ACM on Management of Data (SIGMOD '24)*, vol. 2, no. 1, Mar. 2024, Full Paper.

RESEARCH PROJECTS	<p><b>FedKNN: Secure Federated k-Nearest Neighbor Search</b> <span style="float: right;">SIGMOD '24</span></p> <ul style="list-style-type: none"> <li>Proposed FedKNN, a system for secure and privacy-preserving federated kNN search, achieving up to 4.8× efficiency improvement over SOTA by optimizing local computations with a Distribution-Aware algorithm.</li> </ul> <p><b>HIRE: A Hybrid Learned Index for Robust and Efficient Performance</b> <span style="float: right;">SIGMOD '26</span></p> <ul style="list-style-type: none"> <li>Designed HIRE, a hybrid in-memory index combining learned predictions with traditional robustness, which reduces tail latency by 98% and improves throughput by 41.7× compared to SOTA approaches.</li> </ul> <p><b>Privacy-Preserving Approximate Nearest Neighbor Search</b> <span style="float: right;">Under Review</span></p> <ul style="list-style-type: none"> <li>Proposed a secure ANN index structure within trusted execution environments, mitigating side-channel attacks by enforcing oblivious memory access patterns on proximity graph indexes. In this paper, we achieve plaintext-level throughput with provable privacy guarantees by devising a distribution-aware oblivious access algorithm, significantly outperforming SOTA solutions.</li> </ul> <p><b>Authenticated Approximate Nearest Neighbor Search</b> <span style="float: right;">Under Review</span></p> <ul style="list-style-type: none"> <li>Designed VIDA, a novel scheme for authenticated graph-based ANN search. VIDA leverages vector commitments and Inner Product Arguments to verify distance computations, thereby drastically reducing the overall data transmission size. To further optimize performance, we introduce MeST, a proof aggregation mechanism that compresses the verification of multiple vectors into a single cryptographic proof. In extensive experiment, our solution achieves up to a 5.2x speedup in end-to-end latency and reduces data transfer time by 99.3%.</li> </ul>
TALKS	<ol style="list-style-type: none"> <li>FedKNN: Secure Federated k-Nearest Neighbor Search, <i>2024 ACM SIGMOD International Conference on Management of Data</i>, Santiago, Chile, Jun. 2024.</li> <li>Secure Federated kNN Search, <i>Huawei</i>, Online, Nov. 2024.</li> <li>HIRE: A Hybrid Learned Index for Robust and Efficient Performance under Mixed Workloads, <i>202 HKBU-BNBU Joint Postgraduate Research Symposium</i>, Zhuhai, China, Jun. 2025.</li> </ol>
SERVICES	<ul style="list-style-type: none"> <li>Reviewer, IEEE Transactions on Parallel and Distributed Systems (TPDS) <span style="float: right;">2025</span></li> <li>Reviewer, IEEE Transactions on Knowledge and Data Engineering (TKDE) <span style="float: right;">2024-2026</span></li> <li>Reviewer, Springer Data Science and Engineering (DSE) <span style="float: right;">2024</span></li> <li>External Reviewer, ACM International Conference on Management of Data (SIGMOD) <span style="float: right;">2025</span></li> <li>External Reviewer, International Conference on Very Large Data Bases (VLDB) <span style="float: right;">2023-2025</span></li> <li>External Reviewer, IEEE International Conference on Data Engineering (ICDE) <span style="float: right;">2022-2025</span></li> </ul>
SKILLS	<p><b>Programming</b> Rust, C/C++, Golang, Python, L<sup>A</sup>T<sub>E</sub>X, Bash</p> <p><b>Tools</b> Docker, Vim, Git, Linux</p> <p><b>Languages</b> English, Chinese</p>
AWARDS	<ul style="list-style-type: none"> <li>Academic Excellence Scholarship, Huazhong University of Science and Technology <span style="float: right;">2018, 2021</span></li> <li>Excellent Teaching Assistant Performance Award, Hong Kong Baptist University <span style="float: right;">2022, 2023</span></li> <li>RPg Research Performance Award, Hong Kong Baptist University <span style="float: right;">2024, 2025</span></li> <li>ACM SIGMOD Student Travel Award, ACM <span style="float: right;">2024</span></li> <li>VLDB Student Travel Award, VLDB Endowment <span style="float: right;">2024</span></li> </ul>
REFEREES	<ul style="list-style-type: none"> <li>Prof. Xu Jianliang, Head &amp; Chair Professor, Hong Kong Baptist University</li> <li>Dr. Wang Qichen, Assistant Professor, Nanyang Technological University</li> <li>Prof. Byron Choi, Associate Head &amp; Professor, Hong Kong Baptist University</li> </ul>