STEVENS
INSTITUTE OF TECHNOLOGY
1870

Ph.D. DISSERTATION DEFENSE

Candidate:	Jiarui Li
Degree:	Doctor of Philosophy
School/Department:	Charles V. Schaefer, Jr. School of Engineering and Science / Department of Electrical and Computer Engineering
Date:	Monday, December 8th, 2025
Time:	10:30 a.m. to 12:30 p.m. (Eastern)
Location:	https://stevens.zoom.us/j/96511073355
Title:	Efficient Privacy-Preserving Data Outsourcing with Verifiable Computation
Chairperson:	Dr. Min Song, Department of Electrical and Computer Engineering
Committee Members:	Dr. Lei Wu, Department of Electrical and Computer Engineering Dr. Hao Wang, Department of Electrical and Computer Engineering Dr. Hui Wang, Department of Computer Science Dr. Shucheng Yu, Department of Computer Science and Engineering, Yeshiva University

ABSTRACT

The recent rapid development of AI-driven applications has propelled cloud computing into the spotlight due to its capacity to enable users to utilize the cloud infrastructure from service providers. However, conventional cloud services raise several privacy concerns. One primary concern is that these AI applications necessitate transmitting sensitive user data to the cloud servers for processing, creating risks to data privacy. Moreover, the rise of distributed training paradigms such as federated learning highlights the attacks from malicious training participants. In this dissertation, I address these challenges by integrating lightweight cryptographic schemes into the data pipeline. My studies can be summarized as following parts.

First, I develop a quantum-resistant Multi-Authority Attribute-Based Encryption (MA-ABE) scheme built upon a new ring-based linear secret-sharing scheme and an efficient ring-based lattice sampler. These components result in a decentralized MA-ABE that achieves fine-grained access control, post-quantum security, and significantly improved computational efficiency, highlighting its feasibility for real-world secure data-sharing scenarios in untrusted cloud environments.

Second, I present a secure deep neural network inference offloading framework that enables resource-constrained clients to outsource expensive linear operations to a cloud server through an interactive scalar product evaluation protocol to preserve data privacy while still achieving practical performance suitable for real-time inference.

Finally, I introduce a trusted-execution-environment (TEE)–assisted integrity-verification frameworks for federated learning that address the fundamental challenge of verifying honest local training on untrusted or potentially malicious participants. The first design verifies model update correctness by reconstructing the local training behavior within an enclave, while the second introduces an accumulator-based design that records the intermediate gradients, enabling lightweight integrity checks without reproducing the full training process. Together, these techniques defend against a broad class of integrity violations, including tampering with training code, manipulating local computation, and producing dishonest or poisoned model updates.

These contributions form a cohesive and practical foundation for end-to-end secure and privacy-preserving AI pipelines, offering cryptographic robustness, confidentiality guarantees, and verifiable trustworthiness across the lifecycle of data sharing, model inference, and distributed model training in untrusted environments.