# zkQML: Verifiable and Privacy-Preserving Inference for Quantum Machine Learning (Student Abstract)*

**Seung Kwon Lee[1], Seok Bin Son[1], Joongheon Kim[1], Hoh Peter In[1,2]**

[1] Korea University, Seoul, Republic of Korea
[2] DAO Solution Inc., Seoul, Republic of Korea
{peaceskl, lydiasb, joongheon, hoh_in}@korea.ac.kr

## Abstract

Quantum machine learning (QML) has attracted growing interest for their ability to achieve superior performance with significantly fewer parameters. However, the high cost and scarcity of current hardware push inference to cloud-hosted quantum devices, creating a tension between verifiability and confidentiality. This work proposes a novel framework that converts quantum neural network operations into classical arithmetic circuits that faithfully approximate genuine quantum computations. By encrypting these circuits with zero-knowledge proofs, it ensures computational validity while concealing internal parameters. Experimental results show that our classical circuits achieve fidelity above 0.9996 and total variation distance below 1% compared to actual quantum computations, verifying the practicality of trustworthy and privacy-preserving quantum inference.

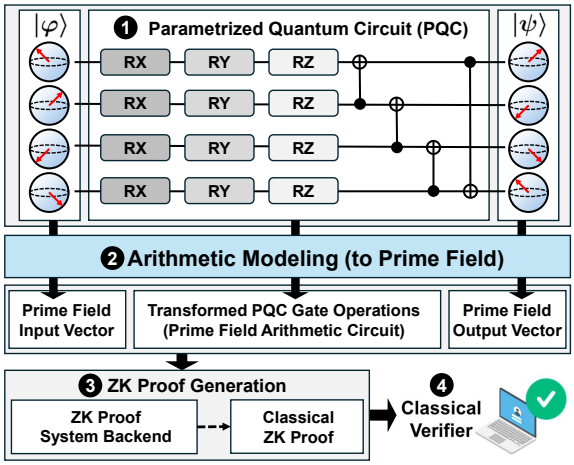


Figure 1: Overview of zkQML.

## Introduction

With the rapid progress of artificial intelligence (AI), quantum computing has emerged as a promising paradigm for extending AI capabilities (Roh et al. 2025). Quantum machine learning (QML) models are widely studied for their supremacy to retain expressive power while using fewer parameters than classical neural networks (Son et al. 2025). However, today's costly and scarce hardware forces inference to be outsourced to cloud-hosted quantum devices, creating tension between ensuring the **trustworthiness** of inference and protecting the **confidentiality** of model parameters. Existing proof techniques often demand quantum-capable verifiers (Fitzsimons et al. 2017) or interactive protocols (Mahadev 2018), limiting scalability. Meanwhile, the emerging field of ZKML[1] (Chen et al. 2024) provides non-interactive proofs for classical models but struggles to scale to large ones (Sun et al. 2024).

To address these challenges, this paper proposes **zkQML**, a framework that bridges ZKML with QML by targeting parametrized quantum circuits (PQCs). PQCs serve as the inference layer in QML-based neural network architectures. Therefore, the proposed framework focuses on PQC to enable privacy-preserving and classical proofs of quantum inference. This integration marks a crucial step toward trustworthy AI, where verifiable reasoning enhanced with zero-knowledge proofs ensures the reliability of quantum inference without compromising privacy. While QML provides superior expressive power over classical ML, its reliance on remote quantum hardware prevents classical users from verifying the correctness of results. zkQML closes this gap by combining quantum expressivity with zero-knowledge verification, enabling secure, non-interactive verification of QML inference in untrusted environments. Consequently, even verifiers without quantum capability can trust the correctness of quantum computations, while secret internal parameters remain protected and computation integrity is guaranteed.

Our contributions are threefold: (1) The first framework for non-interactive, classical-verifier proofs of QML inference, enabling privacy-preserving validation without exposing model parameters; (2) Novel arithmetic modeling techniques that encode PQCs as proof constraints, balancing consistency and efficiency; (3) A prototype system showing the practicality of privacy-preserving quantum inference.

*Hoh Peter In's postal address: Wujung Info & Telecom Bldg #411, 145 Anam-ro, Seoul 02841, Korea (Phone: 82-2-3290-3206, E-mail: hoh_in@korea.ac.kr) (Corresponding Authors: Hoh Peter In, Joongheon Kim)

[1]Zero-Knowledge Proofs for Machine Learning

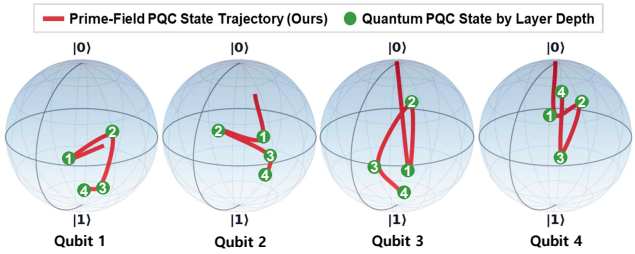


Figure 2: Bloch-sphere visualization of state evolution in the prime-field PQC, compared with the quantum PQC outputs.

## zkQML Framework

zkQML provides a bridging architecture that converts quantum operations into classical arithmetic circuits and generates zero-knowledge proofs, enabling classical verification of quantum computations.

**Procedure.** As shown in Fig. 1, the framework comprises four stages: ① the quantum state passes through the trained PQC to produce an inference; ② through arithmetic modeling, the input, PQC gate operations, and output are transformed into prime-field vectors and an arithmetic circuit; ③ the proof backend compiles these vectors and the arithmetic circuit into constraints, generating a classical zero-knowledge inference proof that hides the internal parameters; and ④ the proof is transmitted to a classical verifier, which performs cryptographic validation of the PQC inference in a non-interactive manner. This process ensures trustworthiness and privacy while maintaining succinctness, hiding parameters, and enabling entirely classical verification.

**Arithmetic Modeling.** The core of bridging *QML* with *ZKML* is the arithmetic encoding of PQCs. Since zero-knowledge proofs operate over prime fields, PQC operations must be reformulated as field arithmetic while preserving the consistency of computation. Quantum states are represented as amplitude vectors with complex coefficients, decomposed into real and imaginary parts, and mapped into fixed-point form over a prime field. This encoding is preserved under PQC transformations, maintaining states and parameters as field elements within bounded precision. Parameterized gates such as $U3(\theta, \phi, \lambda)$ and $CU3$ are modeled as linear transformations, using low-degree polynomial approximations for trigonometric functions and control logic enforced via masking. These encodings inevitably introduce discretization and approximation error. To guarantee consistency, zkQML enforces agreement between the original and arithmetic-encoded PQC outputs within tolerance $\varepsilon$, capturing fixed-point errors. Agreement is evaluated by *fidelity* for global state similarity and *total variation (TV) distance* for comparing probability distributions. This ensures that the encoded PQC remains faithful to the original computation, with verifiable correctness.

## Experimental Results

**Setting.** zkQML is experimentally evaluated on an Intel i7-11700 CPU with 32 GB RAM utilizing quantum computing simulation libraries (e.g., TorchQuantum v0.2.0) and the Halo2 v0.3 ZK toolchain (Rust v1.89.0, Pasta curves). This experiment evaluates a four-qubit PQC-based QML inference model trained on MNIST. The PQC layer depth is varied from 1 to 4 to scale the quantum operations. The circuit employs the widely studied $U3CU3$ design, which captures both single-qubit rotations and entangling gates.

| Depth | 1 | 2 | 3 | 4 |
|---|---|---|---|---|
| **Fidelity** | 0.9997500 | 0.9996591 | 0.9996525 | 0.9996560 |
| **TV Distance** | 0.0089895 | 0.0090602 | 0.0062066 | 0.0079591 |

Table 1. Comparison of original and arithmetic-encoded 4-qubit PQC of *U3CU3* layers.

**Results.** This experiment measures how closely the modeled prime-field PQC approximates the actual quantum PQC operations, thereby verifying that the cryptographically proven PQC arithmetic circuits constitute a trustworthy proof of conventional QML inference. By increasing the layer depth, we further examine whether a stable error bound is maintained as the quantum operations scale. Fig. 2 illustrates the trajectory of the input state vector through the prime-field PQC up to depth four. The markers denote the corresponding states obtained from a quantum PQC simulator at each layer depth. The close alignment between the trajectory and the markers shows that the prime-field PQC modeled in zkQML faithfully tracks the evolution of actual quantum operations.

Table 1 reports the quantitative results. For depths up to four, the fidelity remains above 0.9996 and the TV distance below 0.01, with values near 1 and 0 corresponding to higher similarity and smaller discrepancy, respectively. These results confirm that the modeled prime-field PQC preserves inference consistency with negligible error as the depth increases. This further suggests that zkQML can provide practical and consistency-preserving inference proofs for classical verifiers.

## Conclusion

This paper presents **zkQML**, the first framework for non-interactive, classical-verifier proofs of QML inference. By bridging QML with cryptographic proof systems, zkQML combines the expressive power of quantum models with the verifiability and privacy guarantees of zero-knowledge proofs. Our experiments demonstrate that zkQML can faithfully reproduce the inference behavior of established quantum simulators, yielding fidelity above 0.9996 and total variation distance below 0.01 across multiple PQC depths. This result indicates that the arithmetic-encoded prime-field PQC in zkQML preserves the functional equivalence of genuine quantum inference, validating its correctness under noise-free simulation conditions. In essence, zkQML confirms that verifiable quantum inference can be represented and proven entirely through classical arithmetic operations without direct quantum execution. Looking ahead, extending zkQML to real NISQ hardware will require robust noise modeling and proof calibration to ensure reliable verification under practical conditions.

## Acknowledgments

This research was supported by the Korea University International Joint Research Support Program of 2025 (No. K2521071).

## References

Chen, B.-J.; et al. 2024. ZKML: An Optimizing System for ML Inference in Zero-Knowledge Proofs. In *Proc. European Conference on Computer Systems (EuroSys)*, 560–574. Athens, Greece.

Fitzsimons, J. F.; et al. 2017. Unconditionally Verifiable Blind Quantum Computation. *Physical Review A*, 96: 012303.

Mahadev, U. 2018. Classical Verification of Quantum Computations. In *Proc. IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, 259–267. Paris, France.

Roh, E. J.; et al. 2025. Hybrid Quantum-Classical Style Transfer. In *Proc. AAAI Conference on Artificial Intelligence*, 29480–29481. Philadelphia, PA, USA.

Son, S. B.; et al. 2025. Toward Uniform Quantum Federated Aggregation: Heterogeneity Exclusion Using Entropy and Fidelity. *IEEE Internet of Things Journal*, 12(5): 5732–5741.

Sun, H.; et al. 2024. zkLLM: Zero Knowledge Proofs for Large Language Models. In *Proc. ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 4405–4419. Salt Lake City, UT, USA.