

PristiQ: A Co-Design Framework for Preserving Data Security of Quantum Learning in the Cloud

Zhepeng Wang^{†,§}, Yi Sheng[†], Nirajan Koirala[‡], Kanad Basu^{*}, Taeho Jung[‡], Cheng-Chang Lu[¶], Weiwen Jiang^{†,§}

[†]Department of Electrical and Computer Engineering, George Mason University

[§]Quantum Science and Engineering Center, George Mason University

[‡]Department of Computer Science and Engineering, University of Notre Dame

^{*}Department of Electrical & Computer Engineering, University of Texas at Dallas

[¶]Qradle Inc.

{zwang48, wjiang8}@gmu.edu

Abstract—Benefiting from cloud computing, today’s early-stage quantum computers can be remotely accessed via the cloud services, known as Quantum-as-a-Service (QaaS). However, it poses a high risk of data leakage in quantum machine learning (QML). To run a QML model with QaaS, users need to locally compile their quantum circuits including the subcircuit of data encoding first and then send the compiled circuit to the QaaS provider for execution. If the QaaS provider is untrustworthy, the subcircuit to encode the raw data can be easily stolen. Therefore, we propose a co-design framework for Preserving the data security of QML with the QaaS paradigm, namely PristiQ. By introducing an encryption subcircuit with extra secure qubits associated with a user-defined security key, the security of data can be greatly enhanced. And an automatic search algorithm is proposed to optimize the model to maintain its performance on the encrypted quantum data. Experimental results on simulation and the actual IBM quantum computer both prove the ability of PristiQ to provide high security for the quantum data while maintaining the model performance in QML.

Index Terms—Cloud Quantum Computing, Quantum Machine Learning, Quantum Data Security

I. INTRODUCTION

Over decades of development in quantum science, we are now witnessing the breakthrough of technologies in designing quantum computers: the number of qubits in the actual quantum computers has increased rapidly, from 5 qubits in 2016 to 433 qubits in 2022 [15]. With access to these actual quantum computers via the cloud services, known as Quantum-as-a-Service (QaaS), various studies have been done to explore the quantum advantage in different problems [33], [45], [47], where quantum machine learning (QML), in particular, quantum neural network (QNN) has attracted great research efforts recently [12], [17], [20], [34], [39], [44].

Although QNN is promising in many real applications, the current QaaS paradigm for the execution of QNN poses a high risk of data leakage due to the raw data delivery to the cloud. Fig. 1 shows a typical threat model in the QaaS paradigm for QML applications. In this case, the local user and the attacker may train their own QNN models on the same data but target different tasks. For example, with the same medical images,

the user could build a QNN for classification while the attacker could build a different QNN for segmentation. The training process of a QNN is usually executed in a hybrid quantum-classical way. For instance, the classical data \mathbf{Y}_1 will first be encoded as a quantum circuit $D(\mathbf{Y}_1)$ with a data encoding method (e.g., amplitude encoding). $D(\mathbf{Y}_1)$ and the quantum circuit $QC(\theta)$ for the QNN will be provided to the local quantum compiler, which will generate a compiled quantum circuit $QC(\theta) \cdot D(\mathbf{Y}_1)$. The compiled circuit is then sent to the quantum cloud provider to run the forward pass. And the forward results will be sent back to the user to compute the gradients of learnable parameters θ through parameter shift [24]. For the quantum circuit $QC(\omega)$ of the attacker, it can be trained in the same way.

After the training process, the compiled circuit $QC(\theta) \cdot D(\mathbf{X})$ will be sent to the quantum cloud provider for inference to get an accurate classification of sensitive user data \mathbf{X} . However, if the quantum cloud provider is untrustworthy, the attacker can directly get access to $QC(\theta) \cdot D(\mathbf{X})$. The encoding quantum circuit $D(\mathbf{X})$ can be decoupled from $QC(\theta) \cdot D(\mathbf{X})$ since the data encoding algorithms have fixed circuit patterns that are easy to be detected. Moreover, with $QC(\omega)$, an accurate segmentation of $D(\mathbf{X})$ can be made, which means that the attacker can obtain more information from $D(\mathbf{X})$ with data analysis methods. Such kind of data leakage can not only cause privacy breaches to the users but also bring huge damage to the reputation of the institution that leverages the power of QNNs for their applications. Therefore, the preservation of data security is of great importance in QML applications.

There exist many works on blind quantum computing [4], [8], which focus on the security of quantum circuits. However, they assume that the user has a lightweight quantum terminal to preprocess the data in quantum states, which does not hold in the current QaaS paradigm for quantum applications. Traditional secure computation methods like homomorphic encryption [9], [26] cannot be directly applied to quantum computing since the involved complicated computing operations cannot be implemented by quantum operations. Therefore, a quantum computing-based methodology is necessary to protect the raw input data to the QNNs in the context of QaaS.

PristiQ is a prescription medicine to treat depression. We envision our framework to treat the depression by data leakage in future quantum machine learning research.

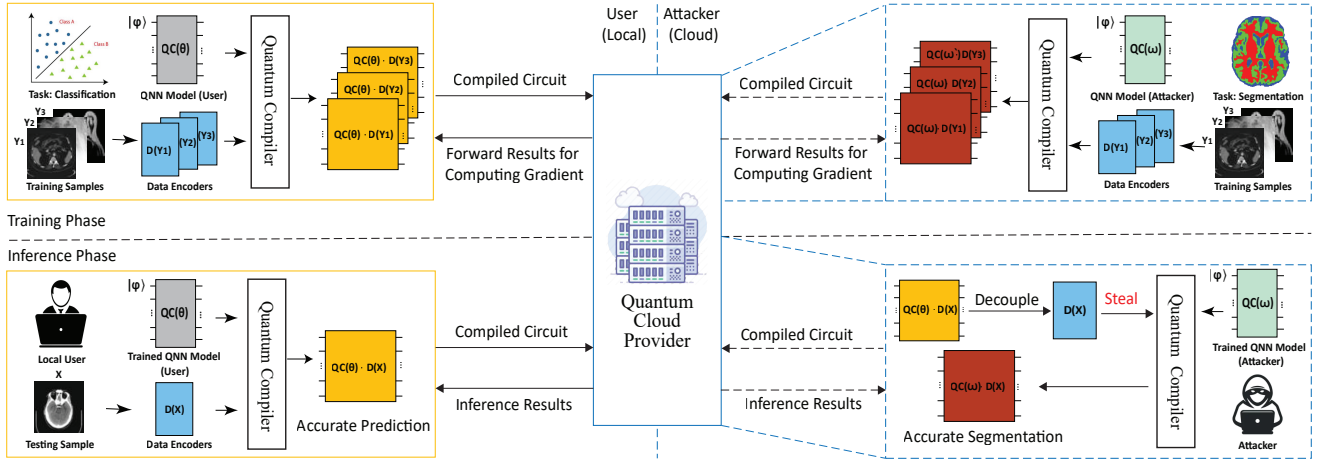


Fig. 1. Illustration of the typical threat model in QaaS paradigm for QML applications: The attacker in the untrustworthy cloud quantum provider can steal the user data for its own QML task.

In this paper, we make the very first attempt to preserve data security in QML applications by formulating a circuit-compiler-model cross-layer design framework, namely PristiQ. The proposed framework is composed of three components: *i) PriCircuit* incorporates an encryption subcircuit into the subcircuit of data encoding to keep data security; *ii) PriCompiler* obfuscates the encryption subcircuit; and *iii) PriModel* automatically searches for the optimal design of the QNN model in order to maintain the performance when directly conducting computations on the encrypted data. These components will work together to provide data security while maintaining high performance for the QML tasks through QaaS.

The main contributions of this paper are as follows.

- The first and the main contribution of this paper is the proposal of PristiQ, which is, to the best of our knowledge, the first framework to secure input data to QNNs under QaaS paradigm.
- PristiQ provides a fundamental understanding of the design of a secure quantum computing system: a cross-layer co-design is required to preserve data security in quantum computing. This design philosophy is not limited to QML applications but can also be applied to other types of quantum applications.
- PristiQ brings the concept of model adaptation to the design of QNN, which preserves high performance on the encrypted data.

II. RELATED WORK

Quantum Neural Network. Since the fundamental operation in deep learning on classical computing is neuron computation, [5] proposed the very first quantum neuron to mimic the behavior of the classical neuron. Based on this work, [3] proposed the quantum recurrent neural network (QRNN), while QuantumFlow [17] proposed a co-design framework to implement the quantum neuron onto a quantum processor with the demonstration of quantum advantage. There is another type of QNN, which has no classical counterparts, known as variational quantum circuits (VQCs) [19], [31], [34]. The QNNs built with VQCs include the quantum generative

adversarial networks (QGAN) [7] and the quantum convolutional neural networks (QCNN) [6]. These works provided the fundamentals and motivation for the research of QNNs. [41] proposed a hybrid quantum neuron architecture design by mixing these two types of quantum neuron designs to achieve better performance. The potential of QNNs has inspired many works to study it from different aspects. To make the QNNs available in the noisy intermediate-scale quantum (NISQ) era, [21] conducted research on the robustness of Quantumflow on NISQ devices. [11] proposed a novel QNN compression algorithm to improve the performance of QNNs under quantum noise. The high variance of quantum noise across devices and time brought new challenges to the development of QML. Therefore, [32] studied the reproducibility of QML while [13] proposed calibration [13] method. [29], [30] are works on the visualization of quantum states and quantum noise, providing tools to further analyze the behavior of QML models with the temporal and spatial variance on the quantum noise. Moreover, although most of the existing works mainly focus on deploying QNNs for inference on quantum devices, there exist recent works on enabling the training of QNNs on quantum devices [37], [38].

Secure Quantum Computing. Techniques in blind quantum computing [4], [8] are proposed to preserve the security of quantum circuits under the assumption that the client has a local quantum terminal to rotate and measure a single qubit while the communication between the client and server relies on a quantum communication channel. However, these assumptions do not hold in the QaaS paradigm for QML. [14] customizes a distributed method to protect the quantum circuit for prime factorization with classical local terminal only. But it cannot be directly extended to QML and it requires the involvement of multiple distributed quantum computers. [23] proposes an encryption algorithm to protect the data encoded by basis encoding, which is not an efficient and practical encoding method in QML. [40] proposed a distributed algorithm to preserve the security of the QML model in the cloud. But it focuses on the protection of the QML circuit instead of the quantum data circuit.

Privacy-Preserving Machine Learning without Encryption.

Various approaches have been proposed to protect data privacy in the scenarios where machine learning is performed over distributed datasets possessed by users. Differential privacy [1], [2] ensures individual data is hidden from an untrusted aggregator by the users who add accuracy-dependent perturbation to the shared data. Federated learning [16], [43] is another option which lets users share model parameters with an untrusted aggregator instead of their datasets. There also exist a hybrid approach that combines the differential privacy and federated learning [42], [46]. In this case, the computation itself is not secured, but the data is preprocessed such that it does not leak sensitive information. Such work is orthogonal to our work where the computation process is secured by encrypting the input data.

Secure Machine Learning with Secure Computation. Some approaches secure the machine learning process by performing computations on encrypted data. This way, outsourced computation/data is secure against the cloud. There exist approaches based on homomorphic encryption [9], [26], secure multiparty computation [22], [25], [27], [28], the combination of the previous two [18], and trusted execution environment technologies [10], [35]. These are the most similar to our work in terms of what is secured in the machine learning process. However, these focus on classical computing environments, and such encryption methodologies can hardly be applicable in quantum computing scenarios.

III. PRISTIQ FRAMEWORK

Fig. 2 shows an overview of PristiQ. Given the data encoding subcircuit $D(\mathbf{X})$ and QNN subcircuit $QC(\theta)$ as inputs, which is shown in Fig. 2 (a), PristiQ will sequentially go through the three components: *PriCircuit* in Fig. 2(b), *PriCompiler* in Fig. 2(c) and *PriModel* in Fig. 2(d).

Based on $D(\mathbf{X})$ and $QC(\theta)$, *PriCircuit* will reconstruct the original circuit $QC(\theta) \cdot D(\mathbf{X})$ by inserting an encryption subcircuit $E(\delta)$ between $D(\mathbf{X})$ and $QC(\theta)$, which introduces extra qubits, as shown in Fig. 2(b). To avoid the attacker detecting the boundary between $D(\mathbf{X})$ and $E(\delta)$ easily, *PriCompiler* is designed to obfuscate $D(\mathbf{X})$ and $E(\delta)$. More specifically, it utilizes the compiling optimization to form a merged circuit $E(\delta) \cdot D(\mathbf{X})$, making their boundary indistinguishable. In this way, the attacker could only extract the encrypted data encoded by $E(\delta) \cdot D(\mathbf{X})$ instead of the raw input data encoded by $D(\mathbf{X})$. However, directly applying the original QNN subcircuit $QC(\theta)$ to the encrypted data may incur significant performance degradation. Therefore, to maintain the performance on encrypted data, *PriModel* will revise $QC(\theta)$ by replacing it with subcircuit $QC'(\theta')$. The neural architecture of the QNN for $QC'(\theta')$ is designed by using an automatic search engine based on reinforcement learning, aiming at the best performance on the encrypted data. In the following content of this section, we will introduce the details of all three components in PristiQ.

A. PriCircuit

Design Principle: The main purpose of *PriCircuit* is to provide data security, i.e., the raw input data should be hidden; however, the encrypted data should also preserve the critical information contained by the raw input data, such that the encrypted data can still be learned effectively. Unlike conventional neural networks in classical computing, most quantum neural networks [5], [17] do not extract the spatial features. Therefore, the most important information in the raw input data is the relative relationship among its features. Based on this observation, we propose a two-stage encryption process for *PriCircuit*. And the details are as follows.

Design overview: Fig. 3 shows the detailed design of encryption subcircuit $E(\delta)$ generated by *PriCircuit* with an example. As shown in Fig. 3(a), there are N qubits used to encode raw input data \mathbf{X} with 2^N features, named data qubits. And M extra qubits are added to the N data qubits, called secure qubits. The first stage applies the subcircuit $S(\delta)$, which is a part of $E(\delta)$, to operate on the newly added M secure qubits with random parameter set δ to perform amortization on the amplitudes from data qubits to secure qubits; then, the second stage utilizes the remaining part of $E(\delta)$ (i.e., subcircuit P) to perform permutation on the amplitudes of the quantum state processed by $S(\delta)$ since the permutation matrix can be implemented by quantum circuits. Based on the design of *PriCircuit*, we define the security key in PristiQ as a tuple $\langle \delta, P \rangle$, where δ is the random parameter set used in the amortization with scaling while P is the design of random permutation circuit.

Stage 1: amortization w/ scaling. As shown in Fig. 3(a), stage 1 starts with initial quantum states S_1 , where $S_1 = |\mathbf{X}\rangle \otimes |0\rangle^{\otimes M} = |q_0 q_1 \dots q_{N-1} q_N \dots q_{N+M-1}\rangle$. Here $|\mathbf{X}\rangle$ is the encoded quantum state for the raw input data \mathbf{X} , using the data encoding subcircuit $D(\mathbf{X})$. Note that for S_1 , only quantum basis state $|q_0 q_1 \dots q_{N-1} 0 \dots 0\rangle$ has non-zero amplitudes since no quantum gates are applied to the M secure qubits. For example, in Fig. 3(b), the number of data qubits N is 2 while the number of secure qubits M is 2 as well. In this case, $S_1 = |q_0 q_1 q_2 q_3\rangle = |\mathbf{X}\rangle \otimes |00\rangle$ and only quantum state $|q_0 q_1 00\rangle$ has non-zero amplitudes, which is clearly shown in the state vector of S_1 in Fig. 3(c).

Stage 1 focuses on the transformation from state S_1 to state S_2 , where each secure qubit is operated by Ry gate with a parameter that is randomly generated. More specifically, the k^{th} secure qubit is denoted as Sec_k and its initial state is $|0\rangle$. After applying an Ry gate with a parameter δ_k to it, the state of Sec_k is converted to $[\cos \frac{\delta_k}{2}, \sin \frac{\delta_k}{2}]^T$, which is denoted by $|QS_k\rangle$. Therefore, the quantum state $|Q\rangle$, which is the state composed of all the secure qubits, can be computed as,

$$|Q\rangle = |QS_0\rangle \otimes |QS_1\rangle \otimes \dots \otimes |QS_k\rangle \otimes \dots \otimes |QS_{M-1}\rangle. \quad (1)$$

By combining the quantum states of data qubits $|\mathbf{X}\rangle$ and secure qubits $|Q\rangle$, we have $S_2 = |\mathbf{X}\rangle \otimes |Q\rangle$, where $|Q\rangle$ is served as the **scaling** vector with 2^M scaling factors for the **amortization** on the amplitudes from data qubits to secure qubits. For example, in Fig. 3(b), the amplitude on $|q_0 q_1 00\rangle$ is amortized to the other three states (i.e., $|q_0 q_1 01\rangle$, $|q_0 q_1 10\rangle$,




Fig. 2. Illustration of PristiQ framework: (a) the data encoding subcircuit $D(\mathbf{X})$ and QNN subcircuit $QC(\theta)$ as the inputs; (b) *PriCircuit*, adding the encryption subcircuit $E(\delta)$; (c) *PriCompiler*, obfuscating $D(\mathbf{X})$ and $E(\delta)$ by forming $E(\delta) \cdot D(\mathbf{X})$; (d) *PriModel*, revising QNN subcircuit $QC(\theta)$ to $QC'(\theta')$ for maintaining the performance on encrypted data.

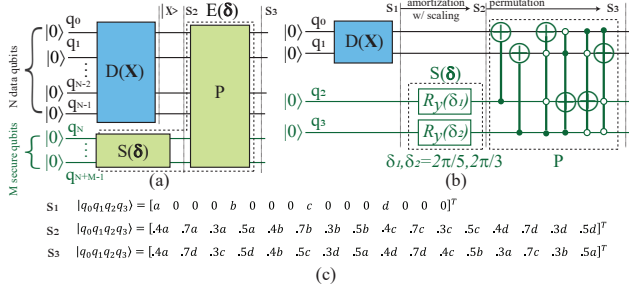


Fig. 3. *PriCircuit*: (a) encryption subcircuit $E(\delta)$ composed of $S(\delta)$ and P ; (b) an example of *PriCircuit* with 2 data qubits and 2 secure qubits; (c) transformations of quantum states: $S_1 \rightarrow S_2$ and $S_2 \rightarrow S_3$.

$|q_0 q_1 11\rangle$) with scaling factors in $|Q\rangle$ determined by the parameter set δ . For instance, in Fig. 3(b), $|\mathbf{X}\rangle = [a, b, c, d]^T$. Since $\delta = [\delta_0, \delta_1] = [\frac{\pi}{5}, \frac{\pi}{3}]$, we have $|QS_0\rangle = [\cos \frac{\pi}{5}, \sin \frac{\pi}{5}] = [0.81, 0.59]$ and $|QS_1\rangle = [\cos \frac{\pi}{3}, \sin \frac{\pi}{3}] = [0.5, 0.87]$. According to Equation (1), we have $|Q\rangle = [0.4, 0.7, 0.3, 0.5]$. We can then obtain $S_2 = [a \cdot |Q\rangle, b \cdot |Q\rangle, c \cdot |Q\rangle, d \cdot |Q\rangle]^T$. The specific value of S_2 is shown in Fig. 3(c). And there exist four groups of values in S_2 , where the values in each group are multiplied with a unique amplitude from $|\mathbf{X}\rangle$. Since all the groups share the same scaling vector $|Q\rangle$, the relative magnitude between arbitrary pair of groups is equal to that between the unique amplitudes they multiply with. In this way, the relative relationship between features within raw input data $|\mathbf{X}\rangle$ is maintained.

Stage 2: permutation. Although the random parameter set δ is unknown to the attacker, stage 1 is not sufficient for the encryption of data since the attacker could get δ and $|\mathbf{X}\rangle$ by solving a system of trigonometric equations if they know the processing of stage 1 is applied. Therefore, *PriCircuit* further generates quantum state S_3 by randomly permuting the amplitudes in S_2 with a permutation matrix P . Since permutation matrix P is unknown to the attacker, the system of trigonometric equations cannot be formulated and thus not be solved. Note that because the permutation matrix is a unitary matrix, it is always feasible to implement it with a corresponding quantum circuit, as shown in Fig. 3(a). Therefore, we have $S_3 = P \cdot S_2$. For the example in Fig. 3(b), it utilizes multiple CNOT gates to implement a permutation matrix and the specific value of S_3 after permutation is shown in Fig. 3(c). And it is obvious that there is no explicit group




Fig. 4. *PriCompiler*: (a) a typical subcircuit in $D(\mathbf{X})$; (b) the subcircuit with dummy CNOT gates; (c) the compiled subcircuit of (b) using *PriCompiler*.

of values that shares the same multiplier after the permutation.

B. PriCompiler

Potential Threat and Design Principle: The main purpose of *PriCompiler* is to obfuscate $E(\delta)$. This is motivated by the fact that the boundary between the secure qubits and data qubits might be detected based on the design of *PriCircuit* since there are only R_y gates on the secure qubits in $S(\delta)$ while the data qubits on $D(\mathbf{X})$ contains multiple two-qubit gates. If the boundary is identified, then $D(\mathbf{X})$ can be extracted.

Therefore, we propose to introduce two-qubit gates on the secure qubits in $S(\delta)$, making the boundary between secure qubits and data qubits indistinguishable. However, naively adding more two-qubit gates may introduce extra undesired operations on the input data and thus disrupt the features for inference. Therefore, *PriCompiler* should try to avoid actual interaction with $D(\mathbf{X})$, which brings new challenges for obfuscation. Fortunately, by reviewing the QaaS paradigm in Fig. 1, we observe that the quantum compiler can be leveraged to make the obfuscation.

Design Details: *PriCompiler* follows a two-stage design flow: (1) a dummy gate insertion and (2) unitary gate decomposition. Based on the fact that two consecutive CNOT gates can be canceled out with each other (i.e., $I = CNOT \cdot CNOT$), we propose a method to create dummy gates and introduce two-qubit gates on the secure qubits. Fig. 4 shows an example of the whole design flow. The circuit in Fig. 4(a) is a typical subcircuit in $D(\mathbf{X})$. In this subcircuit, the three data qubits are operated with two-qubit gates while the single secure qubit only has an R_y gate with a random parameter δ . In Fig. 4(b), an R_y gate with two dummy CNOT gates is added to the

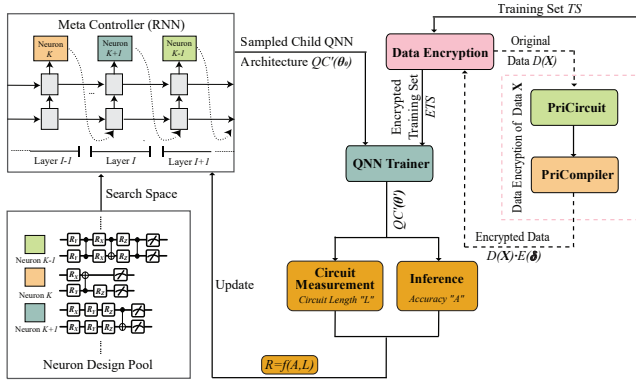


Fig. 5. The workflow of PriModel

secure qubit for obfuscation. More specifically, we first add an Ry gate with a randomly generated parameter α . We then adjust the parameter of the original Ry gate to $\delta - \alpha$ to ensure the functional correctness of the circuit. Between these two Ry gates, we add two consecutive CNOT gates. After this, the first stage is completed.

In the second stage, given the circuit with dummy CNOT gates from the first stage, the default quantum compiler will remove them. But in *PriCompiler*, it adds a barrier between the pair of dummy CNOT gates to split the circuit into two parts, i.e., p_1 and p_2 , as shown in Fig. 4(b). In p_1 and p_2 , each subcircuit will be assembled into a block if it includes exactly three qubits. In this example, p_1 and p_2 can be regarded as two blocks. Since all the gates in each block correspond to a unitary gate, the unitary matrix of each unitary gate can be calculated and further decomposed to a new circuit with different designs consisting of chosen basis gates. Following this rule, the resultant circuit compiled from the one in Fig. 4(b) is presented in Fig. 4(c). In Fig. 4(c), it is clearly shown that multiple two-qubit gates are introduced on the secure qubit q_3 . Due to the limited space, we only keep the CNOT gate in the circuit in Fig. 4(c). And the consecutive CNOT gates can thus not be canceled out due to the existence of Rz gates between them.

C. PriModel

After the process of *PriCircuit* and *PriCompiler*, the total number of input qubits is increased from N to $N + M$. Therefore, the dimensions of features to the QNN model for processing is expanded from 2^N to 2^{N+M} as well. However, the original QNN subcircuit $QC(\theta)$ only contains operations on the N data qubits without consideration of M secure qubits, as shown in Fig. 2(c). Moreover, the original QNN subcircuit $QC(\theta)$ is designed and optimized for the raw input data instead of the encrypted data. As a result, directly applying this QNN model to the encrypted data can degrade the performance of QNN model, which is shown in Section IV. Therefore, a reinforcement learning-based algorithm, namely *PriModel*, is proposed to automatically search for the best QNN architecture for the encrypted data.

Fig. 5 shows the overview of *PriModel*. It is composed of three components: (1) a meta controller to guide the whole

process which is implemented with a recurrent neural network (RNN), (2) a neuron design pool to serve as the search space for the architecture of QNN, which consists of multiple quantum neuron designs. (3) a security-aware QNN evaluator (i.e., the right part of Fig. 5) to measure each sampled solution from the meta controller.

More specifically, for the meta controller, the RNN cell in time step I corresponds to the design of the quantum neuron in layer I of the QNN, which generates a probability distribution for the sampling from the neuron design pool. The sampled neuron design for each layer will be connected to form a child QNN $QC'(\theta_0)$, which is initialized with θ_0 and sent to the evaluator. During the evaluation, $QC'(\theta_0)$ will first be trained on the encrypted training set ETS , which is obtained by encrypting the data in the original training set TS . For each original data $D(X)$ in TS , it will go through *PriCircuit* and *PriCompiler* to get the encrypted data $D(X) \cdot E(\delta)$. And the outcome will be saved in ETS . The trained QNN model $QC'(\theta')$ is then evaluated to calculate two metrics, i.e., circuit length L and accuracy A . These two metrics are used for calculating the reward R to update the meta controller. The reward function is defined as,

$$R = A - b - \lambda \frac{L}{L_{base}}, \quad (2)$$

where b is the exponential moving average of the accuracies of previous sampled QNNs. L_{base} is a preset baseline circuit length for normalization. According to the available quantum computing resources, the user can use λ to control the trade-off between the accuracy and the complexity of circuit implementation for the sampled QNN.

IV. EXPERIMENTAL RESULTS

This section presents the experimental results of PristiQ. We first evaluate the effectiveness of the encryption component of PristiQ (i.e., *PriCircuit*). Then, we show the importance of PristiQ to enable accurate inference on the encrypted data by the results on the noiseless quantum simulator. Finally, we evaluate PristiQ on the noisy quantum simulator and actual quantum computer (i.e., IBMQ Manila) to show the effect of quantum noise on PristiQ.

A. Experimental Setting

Dataset. We evaluate PristiQ on MNIST-2 (class), -3 (class); Fashion-2 (class), -3 (class). When evaluated on the noiseless simulator, the data are downsampled to a resolution of 4×4 from 28×28 , which needs 4 data qubits with amplitude encoding. When evaluated in noisy environments (i.e., noisy simulator of IBMQ Manila and IBMQ Manila), the data is downsampled to a resolution of 4×2 with 3 data qubits. In this case, due to the limited quantum resources, we use 100 samples and 150 samples from MNIST-2 and MNIST-3 for evaluation, respectively.

Security Keys. To generate the security keys, the range of random sampling of rotation angles is $[\frac{1\pi}{8}, \frac{7\pi}{8})$ and the permutation matrix P is implemented with random generation of CNOT gates. When evaluated on the noiseless simulator, for each encryption setting (i.e., number of secure qubits, dataset),

TABLE I
EVALUATION OF *PriCircuit*

Dataset	# Secure Qubits	Model Source	Accuracy (%)	PSNR(dB)
MNIST-2	0	User	99.06	34.28
		Attacker	99.06	
	1	Attacker	54.10 \pm 16.06	19.35 \pm 2.21
MNIST-3	0	User	92.56	34.82
		Attacker	92.29	
	1	Attacker	39.05 \pm 6.12	20.00 \pm 2.18
Fashion-2	0	User	87.19	26.06
		Attacker	86.54	
	1	Attacker	58.97 \pm 17.33	18.15 \pm 2.06
Fashion-3	0	User	77.41	25.32
		Attacker	77.38	
	1	Attacker	44.26 \pm 10.86	17.81 \pm 1.97
	2	Attacker	35.82 \pm 10.48	17.02 \pm 1.47

we randomly generated 8 security keys. When evaluated in noisy environments, for each encryption setting, we randomly generated 3 security keys.

Metrics. We use the peak signal-to-noise ratio (PSNR) to serve as a quantitative metric to evaluate the difference between two images. The smaller the PSNR is, the larger the difference is, indicating a better quality of encryption and security. Since amplitude encoding requires normalization of the features within the original data, it introduces differences between the quantum data and the original data. Therefore, we computed the PSNR between the original data and the quantum data (i.e., the PSNR when the number of secure qubits is 0) to be the baseline for the comparison of PSNR.

QNNs. All the evaluated child QNNs are implemented with TorchQuantum [36]. They are trained with batch size 64 for 20 epochs, using the Adam optimizer with a learning rate of 0.05. They are then compiled and executed by Qiskit. Besides, for *PriModel*, there are 6 options for quantum neural designs at each layer, including 5 different quantum neuron designs from [34] in the neuron design pool and one identity operation.

B. *PriCircuit* Effectively Protects Data

As Fig. 1 shows, we have two types of models for evaluation in this part, i.e., the model for user and the model for attacker. In general cases, the two models should target different types of tasks with the same dataset. But in this experiment, we trained two different models which both target the classification task for simplicity.

Table I reports the results for the evaluation of *PriCircuit*. It clearly shows that without encryption, the data stolen from the user can be utilized by the attacker with its own model effectively. More specifically, the performance of the model for attacker is close to that of the model for user and thus verifies the threat model in Fig. 1. For example, on MNIST-3, the accuracy of the model for attacker is 92.29 % while that of the model for user is 92.56 %.

When the data is encrypted by *PriCircuit*, the security of data can be preserved. It means that the model for the attacker

TABLE II
EVALUATION OF *PristiQ*

Dataset	# Secure Qubits	Model Type	# Param	Circuit Length	Accuracy (%)
MNIST-2	0	User (Vanilla)	56	82	99.06
		User (Vanilla)	56	82	53.65 \pm 14.98
	1	User (PristiQ)	59.5	96.88	99.15 \pm 0.05
MNIST-3	0	User (Vanilla)	56	82	49.81 \pm 4.79
		User (PristiQ)	65	135	99.10 \pm 0.04
	2	User (PristiQ)	65	135	99.10 \pm 0.04
Fashion-2	0	User (Vanilla)	44	52	92.56
		User (Vanilla)	44	52	40.07 \pm 9.52
	1	User (PristiQ)	48.5	72.25	95.08 \pm 0.27
Fashion-3	0	User (Vanilla)	44	52	40.22 \pm 9.96
		User (PristiQ)	58.5	115.25	95.03 \pm 0.24
	2	User (PristiQ)	58.5	115.25	95.03 \pm 0.24
Fashion-2	0	User (Vanilla)	80	103	87.19
		User (Vanilla)	80	103	58.97 \pm 17.35
	1	User (PristiQ)	41.5	59.38	88.30 \pm 0.19
Fashion-3	0	User (Vanilla)	80	103	54.39 \pm 13.54
		User (PristiQ)	45.5	79.5	88.23 \pm 0.16
	2	User (PristiQ)	45.5	79.5	88.23 \pm 0.16
Fashion-3	0	User (Vanilla)	48	46	77.41
		User (Vanilla)	48	46	36.97 \pm 9.26
	1	User (PristiQ)	57.5	90.38	79.13 \pm 0.09
Fashion-3	0	User (Vanilla)	48	46	36.80 \pm 13.42
	2	User (PristiQ)	67	135.88	79.07 \pm 0.10

performs badly on the encrypted data while the difference between the encrypted data and the original data is increased significantly. For example, with only 1 secure qubit, *PriCircuit* can reduce the accuracy and PSNR to 39.05 % and 20.0 dB on average on MNIST-3, respectively.

Moreover, increasing the number of secure qubits can further enhance the data security. For instance, on MNIST-3, the accuracy drops from 39.05 % to 36.29 % and the PSNR reduces from 20.0 dB to 18.69 dB respectively when increasing the secure qubits from 1 to 2.

C. *PristiQ* Can Enable Accurate Inference While Protecting Data

Table II shows the results for the evaluation of *PristiQ* on the noiseless simulator, where the model “User (Vanilla)” refers to the model for user optimized on the original data. Without encryption (i.e., the number of secure qubits is 0), it is clearly shown that only using *PriCircuit* to protect data will lead to the performance collapse of the model for user. For example, on MNIST-2, the accuracy drops from 99.06 % to 49.81 % on average with 2 secure qubits.

By applying *PristiQ* to building the model, where *PriCircuit* is included, the performance of the model (i.e., “User (PristiQ)” in Table II) can be recovered on all the evaluated datasets. For example, on MNIST-2, the accuracy is recovered to 99.10% on average when encrypted with 2 secure qubits, which is even improved by 0.04 % compared with the model optimized on the original data.

D. Evaluation of *PristiQ* in Noisy Environments

Fig. 6 shows the effect of quantum noise on the performance of *PristiQ* on MNIST-2 and MNIST-3. In Fig. 6, the orange bar (Vanilla) denotes the model for user optimized on the original data while the blue bar (PristiQ) denotes the model optimized on the encrypted data with *PristiQ*. Fig. 6 (a) and (c) show the




Fig. 6. Results of PristiQ for MNIST-2 on (a) simulator of IBMQ Manila and on (b) IBMQ Manila, and for MNIST-3 on (c) simulator of IBMQ Manila and on (d) IBMQ Manila.

results on the noisy simulator of IBMQ Manila, while Fig. 6 (b) and (d) show the results on the actual quantum computer (i.e., IBMQ Manila).

For the vanilla model, we can conclude that its performance degrades significantly on the encrypted data generated by *PriCircuit* in noisy environments. For example, with 2 secure qubits, the accuracy decreases from 98 % to 49 % on the noisy simulator and drops from 95 % to 64 % on the actual quantum computer on MNIST-2.

By applying PristiQ to optimizing the model, the performance can still be recovered even with the quantum noise. For instance, with 2 secure qubits, the accuracy can be recovered from 49 % to 97 % on the noisy simulator and increased from 64 % to 95 % on the actual quantum computer on MNIST-2.

V. CONCLUSION

In this paper, we made the very first exploration of the problem of data security for QML in the cloud. We proposed PristiQ, a framework to preserve data security in QML. By creating an encryption subcircuit with a user-defined security key, the important information in the original data is protected. Besides, PristiQ utilizes an automatic model optimizer to achieve high performance on the encrypted data. Extensive experiments in the noiseless quantum simulator and noisy quantum environments are conducted to show the effectiveness of PristiQ. Moreover, the design philosophy of PristiQ, a cross-layer co-design from the circuit level and compiler level to the application level, could be applied to guide the design of a secure quantum computing system in applications beyond QML.

ACKNOWLEDGMENT

This work is partly supported by the National Science Foundation (NSF) OAC-2311949 and OAC-2320957. The research used IBM Quantum resources via the Oak Ridge Leadership Computing Facility at the Oak Ridge National Lab, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. This project was also supported by resources provided by the Office of Research Computing at George Mason University (URL: <https://orc.gmu.edu>) and funded in part by grants from the National Science Foundation (Award Number 2018631). And we also thank Mason's QSEC and C-TASC centers for their support.

REFERENCES

- [1] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, "Deep learning with differential privacy," in *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, 2016, pp. 308–318.
- [2] P. C. M. Arachchige, P. Bertok, I. Khalil, D. Liu, S. Camtepe, and M. Atiquzzaman, "Local differential privacy for deep learning," *IEEE Internet of Things Journal*, vol. 7, no. 7, pp. 5827–5842, 2019.
- [3] J. Bausch, "Recurrent quantum neural networks," *arXiv preprint arXiv:2006.14619*, 2020.
- [4] A. Broadbent, J. Fitzsimons, and E. Kashefi, "Universal blind quantum computation," in *2009 50th Annual IEEE Symposium on Foundations of Computer Science*. IEEE, 2009, pp. 517–526.
- [5] Y. Cao, G. G. Guerreschi, and A. Aspuru-Guzik, "Quantum neuron: an elementary building block for machine learning on quantum computers," *arXiv preprint arXiv:1711.11240*, 2017.
- [6] I. Cong, S. Choi, and M. D. Lukin, "Quantum convolutional neural networks," *Nature Physics*, vol. 15, no. 12, pp. 1273–1278, 2019.
- [7] P.-L. Dallaire-Demers and N. Killoran, "Quantum generative adversarial networks," *Physical Review A*, vol. 98, no. 1, p. 012324, 2018.
- [8] J. F. Fitzsimons, "Private quantum computation: an introduction to blind quantum computing and related protocols," *npj Quantum Information*, vol. 3, no. 1, pp. 1–11, 2017.
- [9] R. Gilad-Bachrach, N. Dowlin, K. Laine, K. Lauter, M. Naehrig, and J. Wernsing, "Cryptonets: Applying neural networks to encrypted data with high throughput and accuracy," in *International conference on machine learning*. PMLR, 2016, pp. 201–210.
- [10] H. Hashemi, Y. Wang, and M. Annavaram, "Darknight: A data privacy scheme for training and inference of deep neural networks," *arXiv preprint arXiv:2006.01300*, 2020.
- [11] Z. Hu, P. Dong, Z. Wang, Y. Lin, Y. Wang, and W. Jiang, "Quantum neural network compression," in *Proceedings of the 41st IEEE/ACM International Conference on Computer-Aided Design*, 2022, pp. 1–9.
- [12] Z. Hu, J. Li, Z. Pan, S. Zhou, L. Yang, C. Ding, O. Khan, T. Geng, and W. Jiang, "On the design of quantum graph convolutional neural network in the niscq-era and beyond," in *2022 IEEE 40th International Conference on Computer Design (ICCD)*. IEEE, 2022, pp. 290–297.
- [13] Z. Hu, R. Wölle, M. Tian, Q. Guan, T. Humble, and W. Jiang, "Toward consistent high-fidelity quantum learning on unstable devices via efficient in-situ calibration," in *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, vol. 1. IEEE, 2023, pp. 848–858.
- [14] H.-L. Huang, Q. Zhao, X. Ma, C. Liu, Z.-E. Su, X.-L. Wang, L. Li, N.-L. Liu, B. C. Sanders, C.-Y. Lu *et al.*, "Experimental blind quantum computing for a classical client," *Physical review letters*, vol. 119, no. 5, p. 050503, 2017.
- [15] IBM, "Ibm unveils world's largest quantum computer at 433 qubits," <https://www.newscientist.com/article/2346074-ibm-unveils-worlds-largest-quantum-computer-at-433-qubits>, 2022, accessed: November, 2022.
- [16] M. Jiang, T. Jung, R. Karl, and T. Zhao, "Federated dynamic graph neural networks with secure aggregation for video-based distributed surveillance," *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2022.

- [17] W. Jiang, J. Xiong, and Y. Shi, "A co-design framework of neural networks and quantum circuits towards quantum advantage," *Nature communications*, vol. 12, no. 1, pp. 1–13, 2021.
- [18] C. Juvekar, V. Vaikuntanathan, and A. Chandrakasan, "{GAZELLE}: A low latency framework for secure neural network inference," in *27th USENIX Security Symposium (USENIX Security 18)*, 2018, pp. 1651–1669.
- [19] J. Li, A. Li, and W. Jiang, "Quapprox: A framework for benchmarking the approximability of variational quantum circuit," *arXiv preprint arXiv:2402.08261*, 2024.
- [20] J. Li, Z. Wang, Z. Hu, P. Date, A. Li, and W. Jiang, "A novel spatial-temporal variational quantum circuit to enable deep learning on nisc devices," in *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, vol. 1. IEEE, 2023, pp. 272–282.
- [21] Z. Liang, Z. Wang, J. Yang, L. Yang, Y. Shi, and W. Jiang, "Can noise on qubits be learned in quantum neural network? a case study on quantumflow," in *2021 IEEE/ACM International Conference On Computer Aided Design (ICCAD)*. IEEE, 2021, pp. 1–7.
- [22] Y. Luo, N. Xu, H. Peng, C. Wang, S. Duan, K. Mahmood, W. Wen, C. Ding, and X. Xu, "Aq2pnn: Enabling two-party privacy-preserving deep neural network inference with adaptive quantization," in *Proceedings of the 56th Annual IEEE/ACM International Symposium on Microarchitecture*, 2023, pp. 628–640.
- [23] U. Mahadev, "Classical homomorphic encryption for quantum circuits," *SIAM Journal on Computing*, no. 0, pp. FOCS18–189, 2020.
- [24] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, "Quantum circuit learning," *Physical Review A*, vol. 98, no. 3, p. 032309, 2018.
- [25] H. Peng, S. Huang, T. Zhou, Y. Luo, C. Wang, Z. Wang, J. Zhao, X. Xie, A. Li, T. Geng *et al.*, "Autorep: Automatic relu replacement for fast private network inference," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 5178–5188.
- [26] H. Peng, R. Ran, Y. Luo, J. Zhao, S. Huang, K. Thorat, T. Geng, C. Wang, X. Xu, W. Wen *et al.*, "Lingcn: Structural linearized graph convolutional network for homomorphically encrypted inference," *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [27] H. Peng, S. Zhou, Y. Luo, N. Xu, S. Duan, R. Ran, J. Zhao, S. Huang, X. Xie, C. Wang *et al.*, "Rrnet: Towards relu-reduced neural network for two-party computation based private inference," *arXiv preprint arXiv:2302.02292*, 2023.
- [28] M. S. Riazi, C. Weinert, O. Tkachenko, E. M. Songhori, T. Schneider, and F. Koushanfar, "Chameleon: A hybrid secure computation framework for machine learning applications," in *Proceedings of the 2018 on Asia conference on computer and communications security*, 2018, pp. 707–721.
- [29] S. Ruan, Z. Liang, Q. Guan, P. Griffin, X. Wen, Y. Lin, and Y. Wang, "Violet: Visual analytics for explainable quantum neural networks," *arXiv preprint arXiv:2312.15276*, 2023.
- [30] S. Ruan, R. Yuan, Q. Guan, Y. Lin, Y. Mao, W. Jiang, Z. Wang, W. Xu, and Y. Wang, "Venus: A geometrical representation for quantum state visualization," in *Computer Graphics Forum*, vol. 42, no. 3. Wiley Online Library, 2023, pp. 247–258.
- [31] M. Schuld, R. Sweke, and J. J. Meyer, "Effect of data encoding on the expressive power of variational quantum-machine-learning models," *Physical Review A*, vol. 103, no. 3, p. 032430, 2021.
- [32] P. Senapati, Z. Wang, W. Jiang, T. S. Humble, B. Fang, S. Xu, and Q. Guan, "Towards redefining the reproducibility in quantum computing: A data analysis approach on nisc devices," in *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, vol. 1. IEEE, 2023, pp. 468–474.
- [33] P. W. Shor, "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer," *SIAM review*, vol. 41, no. 2, pp. 303–332, 1999.
- [34] S. Sim, P. D. Johnson, and A. Aspuru-Guzik, "Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms," *Advanced Quantum Technologies*, vol. 2, no. 12, p. 1900070, 2019.
- [35] F. Tramer and D. Boneh, "Slalom: Fast, verifiable and private execution of neural networks in trusted hardware," *arXiv preprint arXiv:1806.03287*, 2018.
- [36] H. Wang, Y. Ding, J. Gu, Z. Li, Y. Lin, D. Z. Pan, F. T. Chong, and S. Han, "Quantumnas: Noise-adaptive search for robust quantum circuits," in *The 28th IEEE International Symposium on High-Performance Computer Architecture (HPCA-28)*, 2022.
- [37] H. Wang, J. Gu, Y. Ding, Z. Li, F. T. Chong, D. Z. Pan, and S. Han, "Quantumnat: quantum noise-aware training with noise injection, quantization and normalization," in *Proceedings of the 59th ACM/IEEE design automation conference*, 2022, pp. 1–6.
- [38] H. Wang, Z. Li, J. Gu, Y. Ding, D. Z. Pan, and S. Han, "Qoc: quantum on-chip training with parameter shift and gradient pruning," in *Proceedings of the 59th ACM/IEEE Design Automation Conference*, 2022, pp. 655–660.
- [39] H. Wang, P. Liu, J. Cheng, Z. Liang, J. Gu, Z. Li, Y. Ding, W. Jiang, Y. Shi, X. Qian *et al.*, "Quest: Graph transformer for quantum circuit reliability estimation," *arXiv preprint arXiv:2210.16724*, 2022.
- [40] Z. Wang, J. Li, Z. Hu, B. Gage, E. Iwasawa, and W. Jiang, "Qumos: A framework for preserving security of quantum machine learning model," in *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, vol. 1. IEEE, 2023, pp. 1089–1097.
- [41] Z. Wang, Z. Liang, S. Zhou, C. Ding, Y. Shi, and W. Jiang, "Exploration of quantum neural architecture by mixing quantum neuron designs," in *2021 IEEE/ACM International Conference On Computer Aided Design (ICCAD)*. IEEE, 2021, pp. 1–7.
- [42] K. Wei, J. Li, M. Ding, C. Ma, H. H. Yang, F. Farokhi, S. Jin, T. Q. Quek, and H. V. Poor, "Federated learning with differential privacy: Algorithms and performance analysis," *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 3454–3469, 2020.
- [43] Q. Yang, Y. Liu, Y. Cheng, Y. Kang, T. Chen, and H. Yu, "Federated learning," *Synthesis Lectures on Artificial Intelligence and Machine Learning*, vol. 13, no. 3, pp. 1–207, 2019.
- [44] C. Zhan and H. Gupta, "Quantum sensor network algorithms for transmitter localization," in *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, vol. 1. IEEE, 2023, pp. 659–669.
- [45] C. Zhan, H. Gupta, and M. Hillery, "Optimizing initial state of detector sensors in quantum sensor networks," *ACM Transactions on Quantum Computing*, Mar. 2024. [Online]. Available: <http://dx.doi.org/10.1145/3655028>
- [46] Y. Zhao, J. Zhao, M. Yang, T. Wang, N. Wang, L. Lyu, D. Niyato, and K.-Y. Lam, "Local differential privacy-based federated learning for internet of things," *IEEE Internet of Things Journal*, vol. 8, no. 11, pp. 8836–8853, 2020.
- [47] H.-S. Zhong, H. Wang, Y.-H. Deng, M.-C. Chen, L.-C. Peng, Y.-H. Luo, J. Qin, D. Wu, X. Ding, Y. Hu *et al.*, "Quantum computational advantage using photons," *Science*, 2020.