Quantum Federated Learning Experiments in the Cloud with Data Encoding

Shiva Raj Pokhrel, Naman Yash, Jonathan Kua, Gang Li and Lei Pan

ABSTRACT

Quantum Federated Learning (QFL) is an emerging concept that aims to unfold federated learning (FL) over quantum networks, enabling collaborative quantum model training along with local data privacy. We explore the challenges of deploying QFL on cloud platforms, emphasizing quantum intricacies and platform limitations. The proposed data-encoding-driven QFL, with a proof of concept (GitHub Open Source) using genomic data sets on quantum simulators, shows promising results.[1]

1 INTRODUCTION

Quantum computing has unlocked unprecedented computational capabilities, offering solutions to problems beyond classical computers' reach [1]. This breakthrough holds particular promise in machine learning, where quantum algorithms can vastly accelerate data processing, impacting sectors like healthcare, finance, and cybersecurity [2, 3]. By harnessing quantum computing, we can tackle previously insurmountable challenges, marking a significant milestone in computational science and its practical applications.

Our focus lies on Quantum Federated Learning (QFL), a frontier merging federated learning (FL) principles with quantum machine learning (QML) over quantum networks [4]. FL, designed to train models on decentralized devices while preserving data privacy, offers a potent solution for data analysis. By marrying FL with QML, our goal is to enhance computational efficiency and model performance within the quantum realm, all while safeguarding data privacy. Our key contribution lies in crafting and implementing a QFL algorithm, leveraging the potential of cloud-based quantum computing platforms. The analysis of various quantum cloud service providers is crucial to grasp the unique challenges and opportunities in this burgeoning field. Understanding existing quantum cloud infrastructures' capabilities, limitations, and progression roadmaps is paramount for translating FL principles into practical quantum implementations.

In classic FL, each client trains its model locally, sending model parameters to a global server for aggregation without
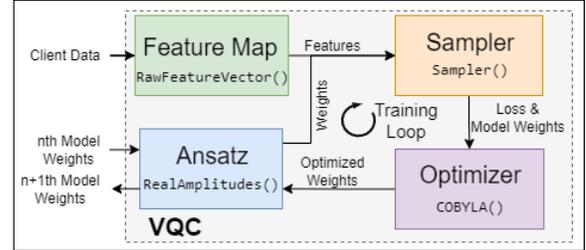
Figure 1: A high level view of local learning in the proposed QFL Process consisting of several key components. The *Feature Map* ingests input data and encodes it into a quantum state. Following this, the *Ansatz* comes into play as a parameterized quantum circuit, with its parameters being iteratively fed by the *Optimizer*–optimization objective function is driven by the outcomes from the *Sampler*.

sharing raw data [5]. Our paper demonstrates the feasibility of QFL using IBM's Qiskit quantum computing library, distinguishing itself by integrating FL into QML, specifically in cloud-based platforms. Our contributions encompass an in-depth analysis of current cloud-based quantum resources for QFL suitability and a proposal for a data-encoding-driven QFL implementation. By comparing state-of-the-art FL techniques, we aim to surmount challenges associated with deploying QFL algorithms.

The novelty of this research lies in its potential to expedite QFL adoption, paving the way for quantum-enhanced machine learning models over the cloud trained efficiently in a distributed setting while upholding local privacy. We design a novel process, as illustrated with Qiskit components in Fig. 1, that can be perceived as approximately transforming input data into a quantum state, exploring and exploiting it using a customizable parameterized quantum circuit, and iteratively optimizing the parameters to steer and achieve the desired outcome based on the global objective function. The background details are discussed later in Sec. 3 with implementation details deferred later in Sec. 3.2.

In the proposed QFL realization, clients transform their unique data into quantum states using a Feature Map, then process them with a parameterized quantum circuit (Ansatz) where local training is conducted using Qiskit and the updated weights are aggregated centrally, and global weights are returned to clients for local model updates.

---

[1]Authors are from IoT & SE Lab, School of IT, Deakin University, shiva.pokhrel@deakin.edu.au.

2 STATE OF QUANTUM COMPUTING

Quantum computing and QML promise dramatic speedups over classical machines for certain tasks. A prime example is the popular Shor's algorithm for factoring integers, where quantum computing provides an exponential advantage over the best-known classical methods [6]. Major tech firms are currently racing to build proprietary quantum hardware and offer cloud access. However, despite rapid progress quantum hardware development is still in nascent development stages [2].

Major tech companies like IBM, Amazon and Microsoft now offer cloud-based access to quantum computers. This allows users to experiment with the latest quantum algorithms without investing in specialized labs or hardware. IBM has emerged as an early front-runner, recently unveiling a 433-qubit quantum processor that currently reigns as the most powerful quantum system.[2] In contrast, Amazon and Microsoft provide access to a diverse range of third-party quantum processors, like IONQ's Harmony and Rigetti's Ankaa-1. By using AWS Braket and Azure Quantum, researchers can test algorithms on systems like IONQ's 32-qubit Harmony or Rigetti's 80-qubit Ankaa processor. As underlying quantum technologies continue rapidly improving, we should see more powerful processors integrated into the cloud ecosystems [7, 8].

QML shows promise but remains in early developmental stages. While programming frameworks like IBM's Qiskit are emerging as robust tools, there is a lack of standardization of different software libraries and limited developer forums and up-to-date resources which pose challenges to newcomers in the field. Cloud-based access aims to increase experimentation by removing on-premises infrastructure barriers. However, resource contention on shared quantum hardware can result in prolonged queuing delays before algorithms execute. This severely impacts techniques like QFL (a full-fledge federation of QML) which require repeated access to quantum circuits for the model training.

While QFL holds immense potential, the current landscape faces challenges in hardware development, software standardization, and efficient cloud-based access for QFL. Looking at specific QML roadmaps, IBM aims to unveil Quantum System Two in 2024, a modular platform paving the way for a scalable quantum-centric ecosystem. By 2025, IBM targets Quantum Condor, a processor with over 1,000 qubits. IONQ focuses on practical quantum devices, projecting increased algorithmic qubits in their Forte and Tempo systems. Google's focus on increasing qubits, achieving a logical qubit prototype, and forming industry partnerships showcases its commitment to practical quantum applications. Meanwhile, Rigetti Computing's strides include launching Ankaa and Lyra processors and collaborations to enhance quantum cloud services, with plans to deploy testbeds at Fermilab by 2025.

Developed by IBM, Qiskit[3] can play an important role in driving the substantial research and development progress of QFL. Its capabilities in quantum simulation and QML are key factors in overcoming QFL operationalization challenges and pushing the boundaries of this emerging field.

3 OPERATIONALIZING QFL IN QISKIT

In the realm of quantum computing, Qiskit's quantum simulators [9] emerge as indispensable tools, empowering developers to test quantum algorithms on classical computers. These simulators serve as flexible platforms that facilitate robust algorithm development and testing processes. Their accessibility benefits users of all levels of proficiency, reducing the iterative journey of algorithm refinement [10]. By providing a means to test algorithms on classical machines, they eliminate hardware access barriers while faithfully emulating the expected behavior of actual quantum systems. This capability accelerates rapid prototyping, empowering developers to iterate and refine without incurring costly overhead. Embracing Qiskit's quantum simulators is not just a choice, it is a strategic imperative for advancing distributed quantum algorithms development efficiently and effectively.

3.1 Modeling QML Libraries

Advances in the Qiskit's QML libraries will be cutting-edge contributions to advancing QFL. The existing library includes pre-built implementations of classifiers like the Variational Quantum Classifier (VQC) and the Quantum Support Vector Regressor (QSVR). Extending these classifiers are vital in constructing QFL models, which form the foundation for success of the framework proposed in this paper. By offering ready-to-use modules, Qiskit simplifies the development cycle, allowing researchers and developers to focus on tailoring QFL models to specific application requirements.

Fig. 1 illustrates the quantum machine learning process enabled by extending the VQC. In our QFL implementation, the VQC uses `RawFeatureVector()` as a feature map to encode classical data into quantum states. This feature map transforms classical input vectors into quantum states through a series of quantum gates, effectively preparing the quantum data for analysis. The `Sampler()` component executes the quantum circuits on various backends, which can be either simulators or real quantum hardware, thereby providing the flexibility to test and run quantum machine learning models in different environments. As shown in Fig. 1, the ansatz is a

---

[2]IBM unveils 400 qubit-plus quantum processor and next-generation IBM Quantum System Two, link

[3]https://qiskit-community.github.io/qiskit-machine-learning/

parameterized quantum circuit that is repeatedly run, where each iteration or 'rep' applies a sequence of gates whose parameters are adjusted in each step of the optimization process. The `RealAmplitudes()` in ansatz is used within variational circuits to construct quantum states that are optimized for our QFL operationalization. The optimization is performed by employing the `COBYLA()` algorithm, a gradient-free optimizer suitable for the noisy conditions of current quantum technology. `COBYLA` iteratively adjusts the parameters of the ansatz in an attempt to minimize a loss function, guiding the model towards the most accurate classifications possible given the input data. This synergy of components represents the architecture of a VQC, which transforms classical data for quantum processing, leverages quantum computation for machine learning tasks, and optimizes its performance over multiple iterations.[4]

It is worth noting that Qiskit's quantum simulators and libraries collectively empower researchers and developers to develop and evaluate new QML and QFL algorithms easily. Furthermore, the inclusion of extended libraries within Qiskit shows a broader commitment to the evolution of distributed quantum machine learning algorithms.

3.2 Implementation in Qiskit

In our pursuit of new strategies to operationalize QFL, we embark on an evolutionary journey from traditional FL. Witnessing the decentralization of model training in FL, where clients actively contribute while preserving data privacy, marks a pivotal phase. However, the transition to QFL introduces practical challenges demanding a fresh perspective.

Currently, achieving QFL with quantum devices on each client poses significant hurdles, primarily due to the absence of small-scale quantum devices on client-side platforms. The quantum computing landscape is dominated by technologies like superconducting qubits and ion trap-based quantum computers, which lack mobility.

In our current QFL implementation, client and global model training occurs primarily on the server. Clients provide data utilized for training on the server, which then returns client models. This orchestration, occurring after each epoch, represents a pragmatic compromise balancing quantum computing capabilities with the absence of client-side quantum devices and edge computing limitations.

To pioneer the operationalization of QFL, we leverage IBM's quantum cloud platform, tapping into its array of quantum processors and simulators. Utilizing Qiskit's simulation capabilities forms the foundation for implementing the QFL framework, opting for established frameworks such as QFedAvg. Our proposed design and implementation of QFL is shown in Figure 2 and is detailed in the following.

---

[4]https://qiskit-community.github.io/qiskit-machine-learning/
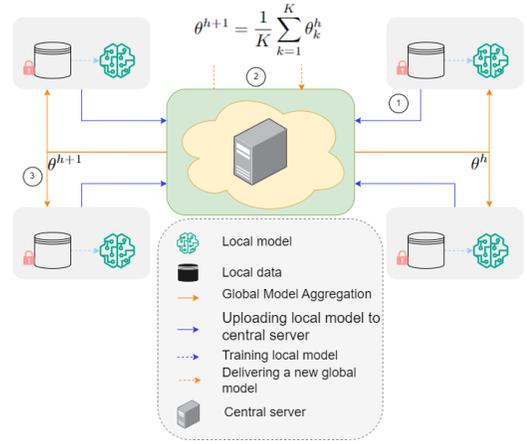
Figure 2: *QFL Process.* Each client in the setup trains models locally and only shares its parameters with a central server (step 1). The server aggregates these to enhance the global model(step 2), then circulates the updated parameters back to clients(step 3), preserving data privacy as client data stays local.

We consider a server, as shown in Figure 2, coexists with a set K of quantum computing clients, collectively training a distributed quantum neural network model (QNN) [11]. Each client possesses a local dataset comprising quantum states $|\psi_m\rangle$ and corresponding labels y_m for $m = 1, ..., M_k$, which serve as the foundation for training an individual QCNN with learnable parameters θ_k.

The model parameter vector, denoted as θ_k, encapsulates all learnable parameters essential for both quantum circuit operations and classical components. Our principal objective lies in optimizing the QNN model's performance, enabling accurate predictions based on the quantum states within each client's dataset. This innovative approach sets the stage for exploring the intricate intersection of quantum computing and federated learning, opening avenues for transformative advancements in QFL. For simplicity and tractability, we utilize the mean squared error (MSE) loss function:

$$J(\theta_k) = \frac{1}{2M_k} \sum_{m=1}^{M_k} (y_m - f_{\theta_k}(|\psi_m\rangle))^2 \qquad (1)$$

and adopt locally via the VQC [12]. In the proposed approach, each client device has access to its own dataset and trains its model locally. The primary goal is along the lines of FL design principles as to collectively train a global model without sharing raw data among clients, thereby preserving data privacy [13].

Figure 3 shows the federation process of the proposed QFL, which involves the following stages:

Local Training: Each client uses its quantum circuit to process its part of the dataset. The local quantum model is then trained using this data.

Parameter Sharing: Upon completion of training in each epoch, clients send their quantum circuit parameters (model weights) to a central server.

Aggregation: The server aggregates these local updates using techniques like simple averaging, weighted averaging, or best pick [13]. This aggregation results in an updated global model.

Global Model Update: The updated weights of the global model are sent back to the client devices.

Local Model Update: Each client updates its local model using global weights. The weight updation formula evolves as

$$\theta^{h+1} = \alpha \times \theta^h + (1 - \alpha) \times \text{global}(\theta^{h+1}) \qquad (2)$$

where α is a weighting factor determined by the current epoch and specific to each client.
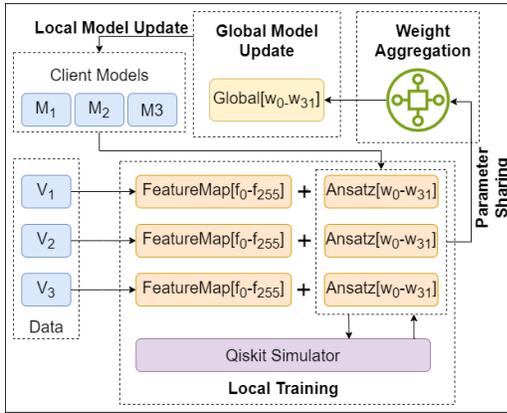
Figure 3: *Details of the QFL framework.* Each client possesses unique data (V1, V2, V3), which undergoes transformation into quantum states using a Feature Map of size 255. These quantum states are then processed by a parameterized quantum circuit (Ansatz) with a depth of 3 and 32 weights. Local training using these quantum components is performed using the Qiskit Simulator. After training, the updated weights from the Ansatz are sent to the server for aggregation, where parameters from all clients are combined, and finally, the aggregated global weights are returned to the clients for local model updates.

3.3 Analyzing Aggregation Schemes

Three weight aggregation schemes are employed to refine the global model by managing client updates effectively. Firstly, Simple Averaging calculates the average of all local model updates. Mathematically, the new global model parameters

θ^{h+1} are determined as the mean of parameters from individual client models, represented by θ_k^h for the k^{th} client's model in the h^{th} iteration.

$$\theta^{h+1} = \frac{1}{K} \sum_{k=1}^{K} \theta_k^h \qquad (3)$$

This method assumes equal contribution from all clients, regardless of the size or diversity of their local datasets.

Weighted Averaging improves the aggregation process by considering the relative importance of each client's update. It assigns weights based on the significance of each client's contribution, resulting in a weighted sum of individual client model updates. This approach allows for a more tailored integration of contributions, accounting for variations in client performance. In this technique, each client's update is weighted by a factor w_k reflecting its significance in the learning process as

$$\theta^{h+1} = \sum_{k=1}^{K} w_k \theta_k^h, \qquad (4)$$

where θ^{h+1} represents the updated parameters of the model after the h-th iteration. The parameters are updated as a weighted sum of the parameters from K client models, denoted by θ_k^h, where each model's contribution is scaled by a weight w_k.

Finally, the *Best Pick* scheme adopts a selective approach, incorporating updates only from clients meeting a predefined performance criterion. The resulting global model is a weighted sum of selected client updates, prioritizing accuracy and reliability. By setting a performance threshold, only updates surpassing this threshold contribute to the global model. The global model update can be represented as:

$$\theta^{h+1} = \sum_{k \in \mathcal{S}} w_k \theta_k^h \qquad (5)$$

where summation is taken over the set $\mathcal{S}$, which represents the subset of clients whose updates satisfy the criterion. The weighting factor w_k assigned to each client model's parameters, θ_k^h, is determined by the performance of the client, emphasizing updates from clients deemed more accurate or reliable, potentially enhancing the overall performance and efficiency of the model.

3.3.1 Datasets. In our experiments, we use a genomic dataset to train a decentralized QFL model, crucial to understanding genomic sequences, consisting of labeled data points (X, Y) representing different genomic characteristics. We chose this dataset due to its sufficient sample size for algorithm testing and the promising potential of genomics for future QFL applications. Our algorithm, demonstrated on the IBM Cloud platform, utilizes multiple QML models and

---

Algorithm 1 Proof of Concept QFL over IBM Cloud

Input: Data (X, Y), Max iterations I, Initial client weights $w_{i\,i=1}^n$, Initial global weights w_g, Number of clients n

1: **Procedure Initialize**
2: Split genomic data into user-specific datasets:
3: $(x_i, y_i) \sim (X, Y)$ for $i = 1, \ldots, n$
   —————————————-
1: **Procedure QFL Learning with Training**
2: **for** $epoch = 1, \ldots, I$ **do**
3:     Select a subset of participating clients, say m clients
4:     **for** each client i in the selected subset **do**
5:         Initialize client-specific quantum circuit Q_i
6:         Encode client data x_i into quantum states X_{q_i}
7:         Optimize client's model to minimize loss and update weights: w_i
8:         $w_i = \text{minimize loss}(f_{Q_i}(X_{q_i}), y_i)$
9:         Upload local updates w_i to the global server
10:    **end for**
11:    Aggregate local updates at the global server:
12:    Broadcast updated global weights w_g to all clients
13:    **for** each client i **do**
14:        Update local model weights with global weights:
15:    **end for**
16: **end for**

---

follows the essential steps outlined in Algorithm 1, encompassing data preparation, model training, and aggregation of updates from multiple clients. Each client, in our proposed QFL implementation, utilizes a quantum circuit to process genomic data and update model weights based on its local dataset, subsequently aggregated by the server to refine the global model, ensuring data privacy while leveraging collective learning.

3.4 Encoding driven QFL

In our implementation of QFL, data encoding stands out as a critical step, transforming classical data into a format compatible with quantum processing. Our QFL implementation using a genomic dataset showcases the potential of quantum facilities to manage complex, high-dimensional data. By employing amplitude encoding, we simplify model training with our 200-feature dataset, leveraging only a fraction of the qubits typically required. This approach proves particularly advantageous in genomic data analysis, where traditional computing methods may struggle with data scale and complexity.

We consider three primary encoding techniques: basis encoding, amplitude encoding, and angle encoding, as follows. Basis encoding, while straightforward, suffers from inefficiencies in qubit usage, particularly for larger datasets.
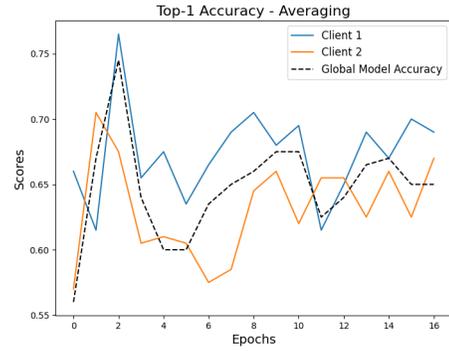
Figure 4: Comparision of the evolution of Top-1 Accuracy over epochs for the global model and clients models using the averaging technique.
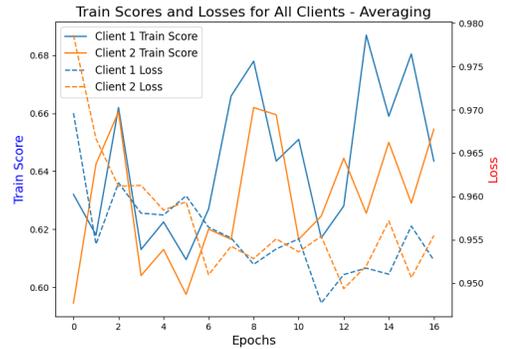
Figure 5: Evolution of Top-1 Accuracy and Loss over epochs for clients using the averaging technique.

Amplitude encoding emerges as a more efficient alternative, especially for handling substantial data volumes. This technique capitalizes on the amplitudes of a quantum state to represent data points, offering a compact representation that can encode vast amounts of information with minimal qubits. For instance, our 200-feature genomic dataset can be effectively encoded using just 8 qubits, making amplitude encoding a practical choice within current quantum computing limitations. This method proves invaluable in quantum machine learning, offering a streamlined approach to represent high-dimensional data.

In contrast, angle encoding, akin to basis encoding in its reliance on one qubit per feature, falls short in scalability for larger datasets due to its demanding qubit requirements. Given these considerations, we opt for amplitude encoding, which efficiently utilizes qubits and aligns with the 25-qubit limit of Qiskit's `aer_simulator` utilized in our research.

As shown in Step 6, our Algorithm 1 efficiently uses quantum circuits for data encoding and processing, presenting a novel methodology for genomic data analysis within a
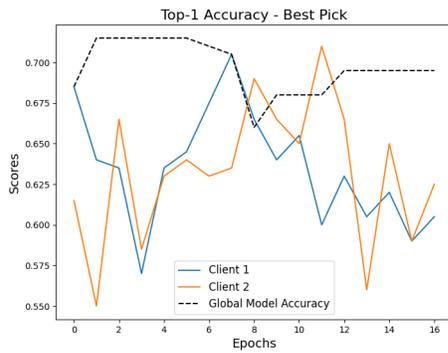
Figure 6: Evolution of Top-1 Accuracy for the global model and individual clients using the Best Pick scheme.

Figure 7: Temporal evolution of the training scores and losses for all clients using the Best Pick.

federated learning framework. Utilizing the IBM Cloud platform for quantum computing, our implementation offers scalability and accessibility, potentially paving the way for transformative advancements in genomic research.

3.5 Performance Evaluation

We conducted a series of experiments to train the developed Quantum Federated Learning (QFL) framework using three distinct global aggregation methods: *Averaging, Best Pick, and Weighted Averaging*. The results of these experiments are depicted in Figures 4 to 9 and are detailed below.

In Figure 4 averaging method exhibited varying Top-1 accuracy between clients, with the global model's accuracy typically positioned between the clients' accuracies. Both the high and low performances of the individual client models directly influenced the global model, placing its accuracy within the range of the client's accuracy. The accuracy and training loss patterns for all clients demonstrated in Figure 5 exhibit a consistent decrease in loss and a tractable improvement in accuracy over epochs.

Figure 8: Temporal evolution of the Top-1 Accuracy for the global model and individual clients using the Weighted Averaging.

In contrast, the best pick method demonstrated a more selective behavior in Figure 6 of our experiments, often resulting in the global model's accuracy surpassing individual clients' accuracies. This observation suggests that the global model performance was driven by the highest-performing clients, with minimal or no contributions from poorly performing ones. The decreasing loss and increasing accuracy observed in Figure 7 over epochs indicated progressive improvements in the training of clients.

With weighted averaging, see the results in Figures 8 and 9. The QFL performance closely matched the highest performing clients. This approach minimally affected the global model's performance by poor-performing clients while significantly benefiting from higher-performing ones. This method ensures that the global model takes advantage of the strengths of the higher performing clients while mitigating the impact of the lower performing ones. Training scores and losses exhibited a steady trend of reduced loss and improved accuracy.

Figure 9: Temporal evolution of training scores and losses for all clients using the Weighted Averaging.

4 CONCLUSION

This article highlights the efficacy of different aggregation methods in the operationalization of QFL over Qiskit, demonstrating how data encoding and weighted averaging enhance overall performance by leveraging the strengths of high-performing clients while mitigating the impact of low-performing ones. More research is needed to refine QFL techniques, Qiskit implementation, and transition from simulators to real quantum hardware for successful implementation on cloud computing infrastructure.

REFERENCES

[1] S. S. Gill, A. Kumar, H. Singh, M. Singh, K. Kaur, M. Usman, and R. Buyya, "Quantum computing: A taxonomy, systematic review and future directions," *Software: Practice and Experience*, vol. 52, no. 1, pp. 66–114, 2022.

[2] H. T. Larasati, M. Firdaus, and H. Kim, "Quantum federated learning: Remarks and challenges," in *2022 IEEE 9th International Conference on Cyber Security and Cloud Computing (CSCloud)/2022 IEEE 8th International Conference on Edge Computing and Scalable Cloud (EdgeCom)*, pp. 1–5, 2022.

[3] M. Chehimi and W. Saad, "Quantum federated learning with quantum data," in *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 8617–8621, 2022.

[4] M. C. Caro, H.-Y. Huang, M. Cerezo, K. Sharma, A. Sornborger, L. Cincio, and P. J. Coles, "Generalization in quantum machine learning from few training data," *Nature communications*, vol. 13, no. 1, p. 4919, 2022.

[5] S. R. Pokhrel and J. Choi, "Federated learning with blockchain for autonomous vehicles: Analysis and design challenges," *IEEE Transactions on Communications*, vol. 68, no. 8, pp. 4734–4746, 2020.

[6] P. W. Shor, "Algorithms for quantum computation: Discrete logarithms and factoring," in *Proceedings 35th Annual Symposium on Foundations of Computer Science*, pp. 124–134, 1994.

[7] S. L. et al., "Pricing plans - azure quantum." https://learn.microsoft.com/en-us/azure/quantum/pricing.

[8] Amazon, "Amazon braket pricing." https://aws.amazon.com/braket/pricing/.

[9] IonQ Staff, "The value of classical quantum simulators," March 2023.

[10] M. Kashif and S. Al-Kuwari, "Qiskit as a simulation platform for measurement-based quantum computation," in *2022 IEEE 19th International Conference on Software Architecture Companion (ICSA-C)*, pp. 152–159, 2022.

[11] I. Cong, S. Choi, and M. D. Lukin, "Quantum convolutional neural networks," *Nature Physics*, vol. 15, no. 12, pp. 1273–1278, 2019.

[12] H. Yano, Y. Suzuki, K. M. Itoh, R. Raymond, and N. Yamamoto, "Efficient discrete feature encoding for variational quantum classifier," *IEEE Transactions on Quantum Engineering*, vol. 2, pp. 1–14, 2021.

[13] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Artificial intelligence and statistics*, pp. 1273–1282, PMLR, 2017.