Quantum Federated Learning for Noisy and Imbalanced State Discrimination

Rocco Ballester[1,2,3], Christian Blum[1,3], Jesus Cerquides[1,3] and Luis Artiles[2]

[1]Artificial Intelligence Research Institute (IIIA-CSIC), Barcelona 08193, Spain
[2]Strategic Platform, Barcelona 08018, Spain
[3]Autonomous University of Barcelona (UAB), Barcelona 08193, Spain

Abstract

Quantum Federated Learning (QFL) is an emerging framework that integrates Federated Learning (FL) with Quantum Computing (QC), enabling collaborative training of quantum models across distributed and privacy-preserving environments. While prior work has primarily focused on classical machine learning tasks using Variational Quantum Circuits (VQCs), the application of QFL to inherently quantum problems remains largely unexplored. In this work, we investigate QFL in the context of one of the most fundamental tasks in quantum information theory, namely, Quantum State Discrimination (QSD). We design and evaluate a set of toy yet realistic scenarios involving small, imbalanced, and noisy quantum datasets, reflecting practical constraints common in quantum sensing and metrology. Our findings demonstrate that QFL can successfully overcome local dataset biases and quantum noise, achieving near-optimal performance comparable to analytical results. All code and experiments are made publicly available at https://github.com/roccobb/QFL4QSD to support reproducibility and encourage further research.

Keywords

Quantum Federated Learning, Quantum State Discrimination, Artificial Intelligence, Distributed Quantum Computing

1. Introduction

The rapid advancement of both Machine Learning (ML) and Quantum Computing (QC) is transforming the landscape of modern computing and Artificial Intelligence (AI) [1, 2, 3]. ML has become a cornerstone of contemporary AI, enabling significant breakthroughs in data analysis, pattern recognition, decision-making, and many more [1]. However, the explosive growth of data has exposed key limitations of centralized ML models, including storage bottlenecks, high communication costs, and heightened risks of data leakage [4, 5]. These concerns, along with increasing emphasis on privacy and compliance with regulations such as the EU's General Data Protection Regulation (GDPR), have motivated a shift toward decentralized, privacy-preserving approaches such as Federated Learning (FL), a paradigm in which multiple participants collaboratively train a shared model without exchanging raw data [6, 7].

Parallelly, the field of QC has undergone significant advancements, progressing from theoretical foundations to the realization of Noisy Intermediate-Scale Quantum (NISQ) devices [2]. Unlike classical systems, quantum computers exploit phenomena such as superposition and entanglement to process information in fundamentally new ways. These principles have enabled quantum algorithms that offer computational advantages for specific tasks [8, 9], as well as the development of cryptographic protocols that guarantee information-theoretic security [10].

However, despite the promising capabilities enabled by QC, practical implementations remain constrained by the limitations of current hardware. Today's quantum devices are characterized by a small number of qubits, limited coherence times, high error rates, and the absence of full-scale quantum error correction [11]. These challenges place significant restrictions on the complexity and depth of quantum

AIQxQIA 2025: International Workshop on AI for Quantum and Quantum for AI | co-located with ECAI 2025, Bologna, Italy
rocco.ballester@iiia.csic.es (R. Ballester); christian.blum@iiia.csic.es (C. Blum); cerquide@iiia.csic.es (J. Cerquides); luis.artiles@strategicplatform.com (L. Artiles)
0000-0002-1736-3559 (C. Blum); 0000-0002-3752-644X (J. Cerquides); 0000-0001-6557-7677 (L. Artiles)

circuits that can be reliably executed, thereby limiting the scope of near-term quantum applications [12].

In addition, many practical quantum computing scenarios involve datasets that are inherently small, distributed, and sensitive [13, 14]. This is particularly evident in domains such as quantum sensing and quantum metrology, where data is often generated across multiple remote quantum devices or laboratories [15]. In such contexts, centralizing data is not only impractical due to hardware constraints but may also be undesirable due to privacy concerns or institutional boundaries [16].

These challenges, combined with the growing need for scalable, distributed quantum data processing [17], have motivated the integration of FL principles into the quantum domain—giving rise to the field known as Quantum Federated Learning (QFL).

Nevertheless, while QFL offers a promising avenue for collaborative quantum learning, the majority of existing works have primarily focused on applying it to classical tasks, particularly to standard classification problems using Variational Quantum Circuits (VQCs) and classical datasets [18]. In contrast, this work explores the use of QFL to address Quantum State Discrimination (QSD), a foundational problem in quantum information theory. We investigate this task under practical constraints, including noisy and imbalanced quantum datasets distributed across different clients. To the best of our knowledge, this is the first work to apply a federated learning framework to a genuinely quantum task—namely, QSD—highlighting the broader potential of QFL for advancing fundamental quantum information processing problems in realistic settings.

The main contributions of this work are as follows:

1. We introduce, for the first time to the best of our knowledge, the use of QFL to address a foundational quantum information problem, namely QSD.
2. We design and implement a set of toy but realistic experiments within the QFL framework, including:
   a) Small and imbalanced quantum datasets with varying degrees of class imbalance;
   b) Small, imbalanced quantum datasets subject to depolarizing noise, modeling realistic quantum noise in state preparation.
3. We demonstrate that QFL can effectively mitigate local bias and noise, achieving strong performance in QSD tasks. Our results are benchmarked against an analytical solution, validating the robustness and accuracy of the QFL approach.
4. To promote transparency and reproducibility, we release a public repository containing all code, experiments, and results available at https://github.com/roccobb/QFL4QSD.

The rest of this paper is organized as follows. Section 2 reviews related work in both QFL and QSD. Section 3 provides the necessary background on these two domains. In Section 4, we define the problem under study and detail the assumptions and challenges involved. Section 5 presents and discusses our experimental results. Finally, Section 6 concludes the paper and outlines potential directions for future work.

2. Related Work

QFL has quickly become a growing area of research at the intersection of quantum computing and distributed learning. Since its initial formulation in the early 2020s [19], QFL has attracted increasing interest for its potential to enable collaborative quantum model training without requiring centralized access to data, making it a promising candidate for real-world applications using current NISQ devices [20, 17, 18].

While some studies have explored the use of QFL to enhance communication security—leveraging quantum cryptographic techniques such as Quantum Key Distribution (QKD) to enable more private and robust federated protocols [21, 22, 23, 24, 25]—the predominant research focus has been on optimization, particularly in the context of classical classification tasks [18]. Thus, the vast majority of works in the field typically use quantum models, most often VQCs, to classify classical datasets such as MNIST or

CIFAR-10 within a federated framework (see [19, 26, 27, 28, 29, 30, 31, 32, 33, 34], to name a few). In this sense, the primary goal of most works is to obtain better convergence or improved efficiency within a federated framework, rather than to address a problem that is inherently quantum in nature.

Among the growing body of QFL literature, only a handful of works explicitly address the use of quantum data, in particular, [35] and [36]. However, these studies adopt relatively simple settings in which quantum data is artificially generated by applying single-qubit rotations, with labels assigned based on whether the resulting state exceeds a predefined excitation threshold. Although VQCs are employed within a federated learning framework, the classification tasks remain straightforward. In contrast, our work addresses a significantly more complex and foundational challenge, namely QSD, while incorporating practical constraints such as noisy states, class imbalance, and heterogeneous data distributions across clients.

On a different note, QSD stands out as a particularly important and impactful problem among the many challenges in quantum information theory [37, 38]. Its relevance spans a wide array of applications, including quantum communication, quantum cryptography, quantum metrology, and quantum sensing, where reliable state identification is essential to system performance and security [39].

A few recent works have explored the use of VQCs for QSD. Chen et al. [40] introduced a framework where parameterized quantum circuits are trained to approximate generalized quantum measurements, also known as Positive Operator-Valued Measure (POVM), enabling discrimination of non-orthogonal quantum states with strong generalization to unseen inputs. Similarly, Lee et al. [41] proposed a framework which learns the optimal POVM for minimum-error discrimination using a cost-function-based variational approach, showing performance close to semidefinite programming solutions. Lastly, [42] tested the performance of noisy quantum neural networks for QSD, showing that effective discrimination is possible even under realistic noise conditions.

These works establish the feasibility of using VQCs for QSD. Nevertheless, they typically assume access to centralized quantum data and rely on idealized settings, often involving large or even infinite datasets. In contrast, we address a more realistic and practical setting where multiple clients each hold small, finite quantum datasets that may be imbalanced and subject to noise. This decentralized approach better reflects the constraints faced in current quantum technologies and distributed quantum applications [14].

3. Fundamental Concepts

3.1. State Discrimination

QSD is the task of identifying an unknown quantum state, ρ, chosen from a known set of possible states $\rho \in \{\rho_1, \rho_2, ..., \rho_n\}$. The objective is to design an optimal measurement that best identifies ρ.

In this work, we focus on quantum states represented by density matrices acting on a 2-dimensional complex Hilbert space $\mathcal{H}$, that is, single-qubit states. In addition, we restrict our attention to the binary classification case, where the unknown state belongs to the set $\rho \in \{\rho_A, \rho_B\}$.

Thus, the goal is to construct a quantum measurement that, upon receiving an unknown qubit state drawn from $\{\rho_A, \rho_B\}$ identifies it as either class A or class B with the smallest possible probability of error.

Let the measurement, E_i, with $i \in \{0, 1\}$, be described by a 2-element POVM. A measurement outcome of $i = 0$ is interpreted as a guess that the system was prepared in state ρ_A, and similarly, a measurement outcome of $i = 1$ is interpreted as having received state ρ_B.

Assuming the state has been prepared in ρ_A or ρ_B with prior probabilities p_A and p_B, respectively, the probability of correctly identifying the state is:

$$P_{corr} = p_A Tr[\rho_A E_0] + p_B Tr[\rho_B E_1]. \tag{1}$$

If the states are orthogonal (i.e., $Tr[\rho_A \rho_B] = 0$), they can be perfectly distinguished. However, in a more realistic and general scenario, ρ_A and ρ_B are not orthogonal (i.e., $0 < Tr[\rho_A \rho_B] < 1$), and it is impossible to distinguish them with certainty.

The optimal strategy for this particular case was provided by Carl W. Helstrom (i.e., [43]).

Since $\sum_i E_i = \mathbb{I}$ for all POVMs, Eq. 1 can be written as

$$P_{corr} = p_A Tr[\rho_A E_0] + p_B Tr[\rho_B \mathbb{I} - \rho_B E_0] = p_B + Tr[(p_A \rho_A - p_B \rho_B) E_0]. \quad (2)$$

To maximize Eq. 2, the trace needs to be maximized. This is achieved when E_0 is a projector on the positive eigenspace of $p_A \rho_A - p_B \rho_B$.

Moreover, in this case, the solution is equivalent to finding a rotation in the Bloch sphere that aligns the optimal measurement direction with a given fixed measurement basis such that $E_i = U^\dagger |i\rangle \langle i| U$.

Finding this optimal rotation effectively reduces the problem of QSD to a variational optimization task. In this work, we employ a Variational Quantum Circuit (VQC) to learn the unitary transformation that implements the desired rotation. The quantum circuit is trained to minimize a suitable loss function that quantifies the discrimination error, allowing the model to approximate the optimal measurement strategy directly from data.

3.2. Quantum Federated Learning

FL is a decentralized machine learning paradigm that enables multiple clients to collaboratively train a shared model without exchanging their raw data. This approach addresses critical concerns around data privacy, communication costs, and scalability by performing local training on each client's dataset and aggregating model updates centrally [7].

Generally, a FL framework comprises a central server and a set of K parties, commonly referred to as clients, each holding a subset of the dataset. The goal is to collaboratively train a machine learning model without disclosing any local data. To accomplish this, each client performs local training and sends its model parameters to the central server, which aggregates the updates and returns the revised global parameters to all clients:

$$w_{t+1} \leftarrow \sum_{k=1}^{K} \frac{n_k}{n} w_{t+1}^k, \quad \text{where } w_{t+1}^k \leftarrow w_t^k - \eta \nabla f(w_t^k), \ \forall k, \quad (3)$$

where $\frac{n_k}{n}$ represents the fraction of dataset holded by client k, w_t^k are the model parameters of client k at step t, η is the learning rate, $f(\cdot)$ represents the machine learning model, and w_t are the global model parameters at step t.

This process is repeated iteratively until a convergence or stopping criterion is met.

Despite its advantages, FL also introduces several challenges, including statistical heterogeneity across clients, communication inefficiencies, and vulnerabilities to adversarial behavior or data poisoning attacks [4].

QFL emerges from the intersection of FL and Quantum Machine Learning (QML) and aims to address some of FL's mentioned challenges while leveraging the potential advantages of QC.

As discussed in Section 2, typical implementations of QFL combine QML models—most notably VQCs—at the client level to boost performance, while employing quantum communication protocols to address privacy concerns and mitigate communication bottlenecks. Nonetheless, we believe that QFL is particularly well-suited for scenarios where quantum data is inherently distributed across different devices or institutions and where data centralization is infeasible or undesirable.

4. Problem Definition

As stated in Section 3.1, the goal in binary QSD is to identify the true label $y \in \{0, 1\}$ of an unknown quantum state $\rho \in \mathcal{H}$, which is known to be prepared from one of two possible sources: ρ_A (class A) or ρ_B (class B). The optimal strategy for minimizing the probability of misclassification is given by Eq. 2, whose success probability depends on both:

- The actual quantum states ρ_A and ρ_B.
- Their respective prior probabilities $p_A = Pr(y = 0)$ and $p_B = Pr(y = 1)$, with $p_A + p_B = 1$.

In this work, we consider a more constrained and realistic version of this problem.
Let there be K clients $\{C_k\}_{k=1}^{K}$, each holding a private dataset of labeled quantum states:

$$D_k = \{(\rho_j^{(k)}, y_j^{(k)})\}_{j=1}^{M}, \text{ where } y_j^{(k)} \in \{0, 1\},$$

with M being the number of training samples per client (for simplicity, we assume that all clients have the same number of samples, i.e, $|D_k| = M, \ \forall k$).
Moreover, it is important to know that in our scenario, each client:

- Does not know the actual quantum states $\rho_j^{(k)}$, only the corresponding class labels $y_j^{(k)}$.
- Does not know the prior probabilities p_A, p_B of the true population from which the unknown test state is drawn.
- May have a class imbalance in their local dataset, such that the empirical class ratio

  $$r_k = \frac{1}{M} \sum_{j=1}^{M} \mathbb{1}[y_j^{(k)} = 0],$$

  differs from the true class prior p_A.
- May also have noisy quantum states in their local datasets due to imperfections in state preparation or transmission.

These constraints make the learning problem significantly more challenging. The presence of noise can degrade the training signal and introduce uncertainty, especially since clients lack access to state fidelity information and cannot correct for errors explicitly.
Similarly, although clients may suspect that their local datasets are imbalanced, and thus not representative of the true population, they cannot estimate or adjust for the true class priors due to the decentralized setup and privacy-preserving constraints. Collectively, these factors complicate the task of collaboratively learning an optimal global measurement strategy, such as the Helstrom measurement, making this problem both realistic and nontrivial.
We consider a VQC of a single trainable rotation gate $RY(\theta)$ and a final measurement in the computational basis in order to learn the optimal rotation for the aforementioned state discrimination task.
While this model could be extended to higher-dimensional state spaces, more complex circuits, or multi-class settings, we deliberately focus on a minimal configuration. This ensures that we can isolate and analyze the fundamental effects of federated optimization, data imbalance, and quantum noise, without conflating them with circuit complexity or optimization instability, while being able to compare our results with analytical solutions.

5. Experiments and Discussion

In this section, we evaluate the performance of our quantum model in both federated and non-federated training settings, and compare the results against the known analytical solution given by the Helstrom bound.
We consider two experimental configurations:

- **Data Imbalance Only.** Clients are affected only by varying local class ratios, with no quantum noise applied. This isolates the effect of statistical bias in local datasets.
- **Data Imbalance and Depolarizing Noise.** Clients experience both class imbalance and depolarizing noise during state preparation. This setting reflects a more realistic challenge, where both data heterogeneity and quantum noise are present.

In our experiments, we consider a binary QSD task between the states $\rho_A = |0\rangle\langle 0|$ and $\rho_B = |+\rangle\langle +|$, where $|+\rangle = \frac{1}{\sqrt{2}}|0\rangle + |1\rangle$, with prior probabilities $p_A = p_B = 0.5$. While this symmetric distribution simplifies the setup and provides a clear reference point, our methodology is applicable to any other class distribution, as will be discussed later.

In this particular setup, the optimal measurement corresponds to a rotation around the Y-axis by an angle of $\theta = arccos(1/\sqrt{2}) \approx 0.7854$ radians. This known analytical solution allows us to quantitatively assess how well our quantum variational model, trained under different conditions, approximates the optimal decision boundary.

Thus, for each configuration, we compare:

- **Local Training.** Each client trains its VQC independently using only local data.
- **Federated Training (QFL).** Clients participate in a federated optimization loop where only model parameters are shared with a central server. No raw data or quantum states are exchanged, preserving privacy and reducing communication complexity.

These experiments aim to assess whether QFL enables a collaborative and privacy-preserving learning process that converges to near-optimal discrimination performance, even under realistic constraints.

5.1. Experiment 1: Imbalanced Datasets

We begin by analyzing the impact of class imbalance across clients in the absence of quantum noise. In this setup, each client receives a local training dataset of size $M = 100$, consisting of 1-qubit states labeled according to class A or class B. During training, we use a single shot per circuit evaluation. However, since gradients are estimated using the parameter-shift rule, each input state is effectively used twice. Equivalently, we may regard the dataset as comprising 100 pairs of state inputs, with each pair corresponding to the two shifted evaluations required for gradient estimation.

The class distribution is not globally uniform but varies across clients, simulating a realistic FL scenario with non-identically distributed data. The class imbalance for client k is determined by a parameter, σ, that controls the variability of local class ratios. Specifically, each client's class-A proportion $r_k \in [0, 1]$ is sampled as:

$$r_k = \frac{1}{1 + exp(-z_k)}, \quad \text{with} \quad z_k \sim \mathcal{N}(0, \sigma^2).$$

Thus, the corresponding labels are drawn with class probabilities $(r_k, 1 - r_k)$. Larger values of σ yield greater diversity in class imbalance across clients. Note that while the true population priors are assumed to be $p_A = p_B = 0.5$, each client's empirical distribution may significantly deviate from this, especially when σ is large.

Nevertheless, the methodology applies to any class distribution, provided that the local class ratios are sampled from a Gaussian centered at the true population prior.

To better understand the impact of class imbalance on the local learning process, we conduct an experiment where a single client trains its VQC model on its 100 samples with varying proportions of class A states ratios. For each class ratio, the training is repeated 50 times with different random samples to account for statistical variability. The metric reported is the angular distance between the learned rotation and the analytically optimal solution.

As shown in Figure 1, high imbalances significantly degrade performance, while class ratios closer to the true prior $p_A = 0.5$ yield more accurate models with lower angular deviations, as expected.

Next, we investigate how well a single client and a federated setup with 10 clients learn the optimal rotation for the QSD task under a highly imbalanced and heterogeneous data distribution (using $\sigma = 4$).

Fig. 2a shows the training dynamics of a single client, which updates its model parameter after each individual state measurement and does not participate in any FL process. In contrast, Fig. 2b illustrates the evolution of the global model in a QFL setup. In this setting, each client performs local updates after every state measurement, but communicates its model parameter to a central server after every 5 measurements. The server aggregates these parameters by averaging and broadcasts the updated global

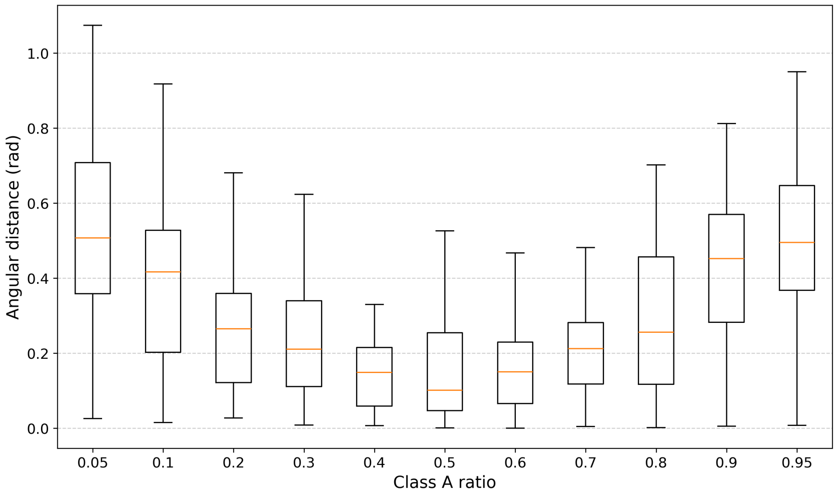

Figure 1: Effect of class imbalance on model accuracy for a single client. Each boxplot shows the distribution of angular distances (in radians) between the learned rotation and the optimal solution, over 50 runs. Performance degrades as the local class ratio deviates from the true prior $p_A = 0.5$.

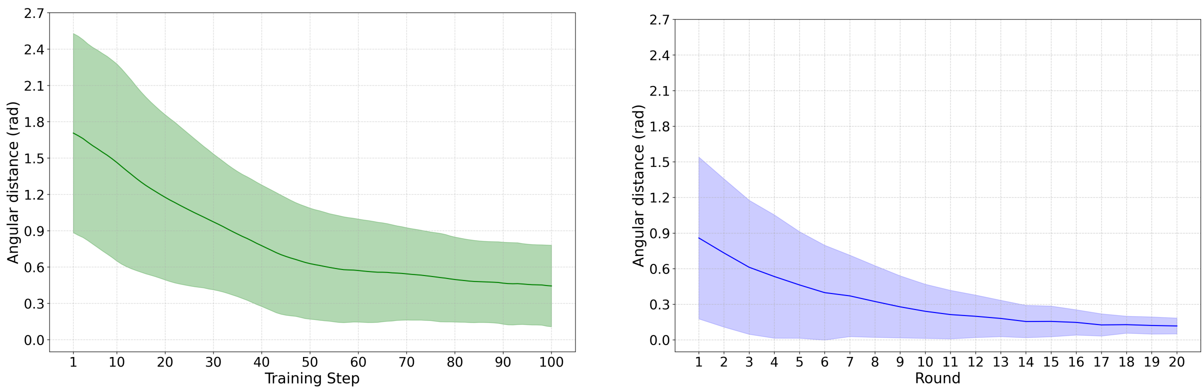

(a) Learning dynamics of a single client.

(b) Learning dynamics of a QFL setup with 10 clients and 20 communication rounds.

Figure 2: Evolution of the angular distance to the optimal Helstrom rotation under strong class imbalance ($\sigma = 4$). Shaded areas denote standard deviation across 50 runs.

model back to all clients. Thus, the QFL setup accounts for $100/5 = 20$ communication rounds. As before, both experiments are repeated 50 times to account for statistical variability.

As shown in Figure 2, the federated setup significantly outperforms the single-client case. Not only does the QFL model converge to a parameter much closer to the optimal rotation, but it also exhibits considerably lower variance across runs. This highlights the ability of QFL to mitigate the effects of local imbalance by leveraging information aggregated from multiple clients without the need to centralize or share local data.

5.2. Experiment 2: Imbalanced and Noisy Datasets

To simulate realistic hardware conditions, we also introduce quantum noise in the form of depolarizing noise. This noise is applied immediately after the state preparation stage, and acts as a quantum channel:

$$\mathcal{E}(\rho) = (1 - p)\rho + \frac{\rho}{3}(X\rho X + Y\rho Y + Z\rho Z),$$

where $p \in [0, 1]$ is the depolarization probability. This models a common and challenging form of noise present in many current quantum devices [15], and serves to test the robustness of QFL under noisy conditions.

Figure 3 shows the angular distance between the learned and optimal rotation as the noise probability increases for a single client setup. Fig. 3a corresponds to the case of balanced datasets ($\sigma = 0$) while Fig. 3b illustrates the performance under high imbalance ($\sigma = 4$). In both scenarios, we observe a consistent degradation in accuracy as the noise level increases. However, the degradation is significantly more pronounced when both imbalance and noise are present, highlighting their compounding effect.

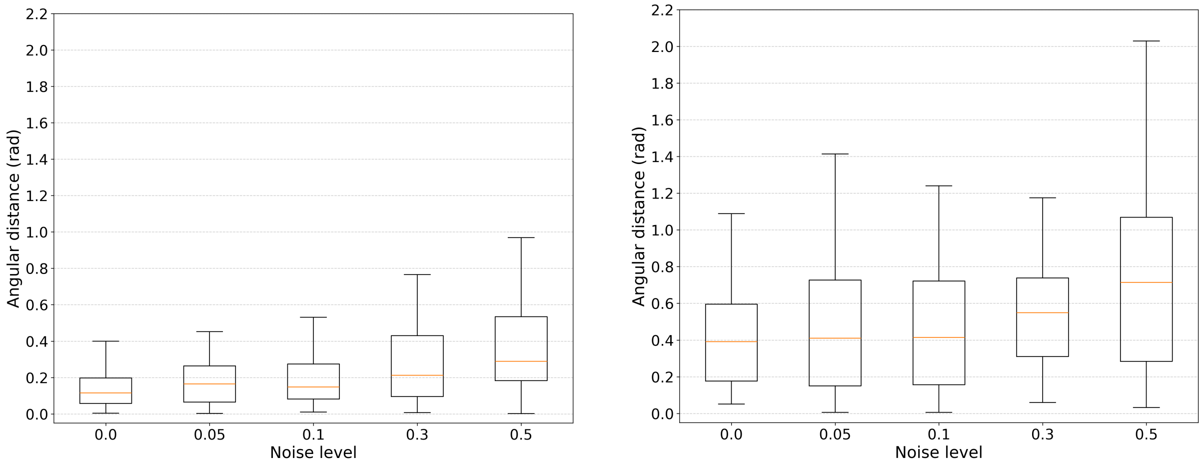

(a) Balanced datasets ($\sigma = 0$). (b) Imbalanced datasets ($\sigma = 4$).

Figure 3: Performance degradation of a single client under increasing depolarizing noise levels. All results are averaged over 50 independent trials to ensure statistical robustness.

As done in Section 5.1, we replicate here the training evolution analysis from Fig. 2 but under more challenging conditions. In particular, we consider a highly imbalanced setting ($\sigma = 4$) combined with strong depolarizing noise ($p = 0.5$).

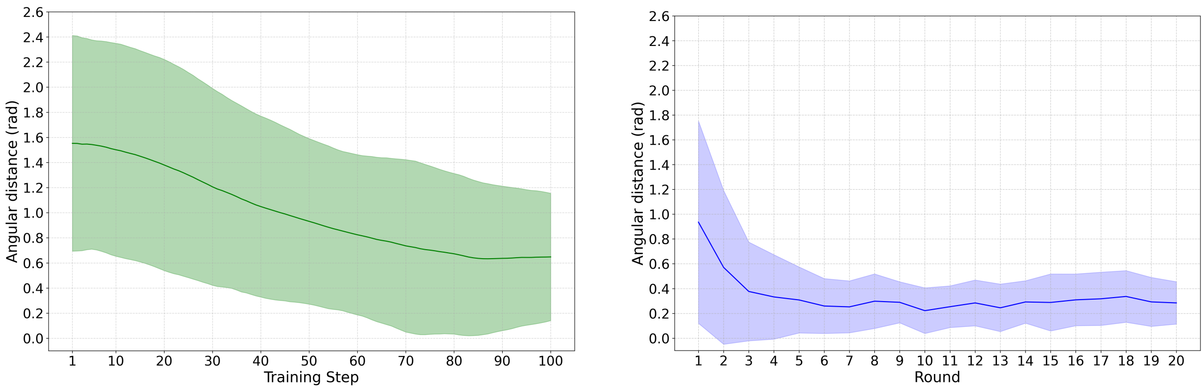

(a) Learning dynamics of a single client. (b) Learning dynamics of a QFL setup with 10 clients and 20 communication rounds.

Figure 4: Evolution of the angular distance to the optimal Helstrom rotation under strong class imbalance ($\sigma = 4$) and strong depolirizing noise ($p = 0.5$). Shaded areas denote standard deviation across 50 runs.

Fig. 4a shows the evolution of a single client's model parameter under a highly imbalanced and noisy scenario, while Fig. 4b shows the evolution of the QFL framework using 10 clients and 20 federated rounds. As before, each curve represents the average angular distance to the optimal rotation over 50 trials. We observe that the single-client setup struggles to find a good approximation to the optimal rotation, with large variance across trials. In contrast, the QFL setup converges more effectively and exhibits significantly reduced variance, demonstrating its robustness in the presence of both data imbalance and quantum noise.

These experiments demonstrate that QFL is a valuable framework for QSD under realistic constraints. By effectively mitigating the detrimental effects of data imbalance and depolarizing noise, QFL enables more accurate and stable training compared to isolated, local approaches. This highlights its potential for scalable and privacy-preserving quantum information processing in practical settings.

6. Conclusion and Future Work

In this work, we have demonstrated that QFL can be effectively employed to mitigate the impact of data imbalance and quantum noise in a decentralized and privacy-preserving setting. While most prior research on QFL has focused on classical machine learning tasks using quantum models, our results show that QFL can also benefit genuinely quantum problems—such as QSD—by enabling robust learning in scenarios where clients have limited, biased, and noisy access to quantum data. Given the constraints of current and near-term quantum hardware, QFL presents a promising framework for practical distributed quantum learning applications.

Looking ahead, several directions for future research emerge. First, a systematic study of key hyperparameters—such as the number of clients, the amount of data per client, the number and frequency of federated rounds, and the aggregation strategies used—could help identify configurations that optimize performance across different quantum learning scenarios. In our experiments, these choices were made heuristically; a more principled exploration could yield deeper insights into the dynamics of QFL and its limitations.

Another natural direction for future research is to increase the complexity of the QSD problem itself. This could involve extending the task to multi-qubit systems, introducing more than two classes, or simulating more realistic quantum environments by incorporating multiple, heterogeneous noise sources. These extensions would not only bring the problem closer to real-world quantum applications but also further test the scalability and robustness of the QFL framework in more demanding scenarios.

Along the same lines, future work could also examine the role of quantum model architecture in distributed learning. While this study focused on a simple VQC to allow for analytical comparison, more expressive ansätze—potentially involving entanglement or deeper layers—could enhance learning performance in more complex settings. Understanding how circuit design influences learning dynamics and generalization in QFL could provide useful guidelines for building more powerful and adaptable quantum models for distributed applications.

Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT in order to: Grammar and spelling check, Paraphrase and reword. After using this tool, the authors reviewed and edited the content as needed and takes full responsibility for the publication's content.

References

[1] S. Tufail, H. Riggs, M. Tariq, A. I. Sarwat, Advancements and Challenges in Machine Learning: A Comprehensive Review of Models, Libraries, Applications, and Algorithms, Electronics 12 (2023) 1789. URL: https://www.mdpi.com/2079-9292/12/8/1789. doi:10.3390/electronics12081789, number: 8 Publisher: Multidisciplinary Digital Publishing Institute.

[2] M. Schuld, F. Petruccione, Supervised Learning with Quantum Computers, Quantum Science and Technology, Springer International Publishing, Cham, 2018. URL: https://link.springer.com/10.1007/978-3-319-96424-9. doi:10.1007/978-3-319-96424-9.

[3] M. Schuld, I. Sinayskiy, F. Petruccione, An introduction to quantum machine learning, Contemporary Physics 56 (2015) 172–185. URL: https://doi.org/10.1080/00107514.2014.964942. doi:10.1080/00107514.2014.964942, publisher: Taylor & Francis _eprint: https://doi.org/10.1080/00107514.2014.964942.

[4] B. S. Guendouzi, S. Ouchani, H. El Assaad, M. El Zaher, A systematic review of federated learning: Challenges, aggregation methods, and development tools, Journal of Network and Computer Applications 220 (2023) 103714. URL: https://linkinghub.elsevier.com/retrieve/pii/S1084804523001339. doi:10.1016/j.jnca.2023.103714.
[5] A. Oussous, F.-Z. Benjelloun, A. Ait Lahcen, S. Belfkih, Big Data technologies: A survey, Journal of King Saud University - Computer and Information Sciences 30 (2018) 431–448. URL: https://www.sciencedirect.com/science/article/pii/S1319157817300034. doi:10.1016/j.jksuci.2017.06.001.
[6] O. Lynskey, Complete and Effective Data Protection, Current Legal Problems 76 (2023) 297–344. URL: https://doi.org/10.1093/clp/cuad009. doi:10.1093/clp/cuad009.
[7] B. McMahan, E. Moore, D. Ramage, S. Hampson, B. A. y. Arcas, Communication-Efficient Learning of Deep Networks from Decentralized Data, in: Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, PMLR, 2017, pp. 1273–1282. URL: https://proceedings.mlr.press/v54/mcmahan17a.html, iSSN: 2640-3498.
[8] Y. Liu, S. Arunachalam, K. Temme, A rigorous and robust quantum speed-up in supervised machine learning, Nature Physics 17 (2021) 1013–1017. URL: http://arxiv.org/abs/2010.02174. doi:10.1038/s41567-021-01287-z, arXiv:2010.02174 [quant-ph].
[9] P. W. Shor, Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer, SIAM Journal on Computing 26 (1997) 1484–1509. URL: https://epubs.siam.org/doi/10.1137/S0097539795293172. doi:10.1137/S0097539795293172, publisher: Society for Industrial and Applied Mathematics.
[10] C. H. Bennett, G. Brassard, Quantum cryptography: Public key distribution and coin tossing, Theoretical Computer Science 560 (2014) 7–11. URL: https://linkinghub.elsevier.com/retrieve/pii/S0304397514004241. doi:10.1016/j.tcs.2014.05.025.
[11] J. Preskill, Quantum Computing in the NISQ era and beyond, Quantum 2 (2018) 79. URL: http://arxiv.org/abs/1801.00862. doi:10.22331/q-2018-08-06-79, arXiv:1801.00862 [cond-mat, physics:quant-ph].
[12] N. P. de Leon, K. M. Itoh, D. Kim, K. K. Mehta, T. E. Northup, H. Paik, B. S. Palmer, N. Samarth, S. Sangtawesin, D. W. Steuerman, Materials challenges and opportunities for quantum computing hardware, Science 372 (2021) eabb2823. URL: https://www.science.org/doi/full/10.1126/science.abb2823. doi:10.1126/science.abb2823, publisher: American Association for the Advancement of Science.
[13] M. E. Rivero-Angeles, Quantum-based wireless sensor networks: A review and open questions, International Journal of Distributed Sensor Networks 17 (2021) 15501477211052210. URL: https://doi.org/10.1177/15501477211052210. doi:10.1177/15501477211052210, publisher: SAGE Publications.
[14] Z. Zhang, Q. Zhuang, Distributed quantum sensing, Quantum Science and Technology 6 (2021) 043001. URL: https://dx.doi.org/10.1088/2058-9565/abd4c3. doi:10.1088/2058-9565/abd4c3, publisher: IOP Publishing.
[15] V. Giovannetti, S. Lloyd, L. Maccone, Advances in quantum metrology, Nature Photonics 5 (2011) 222–229. URL: https://www.nature.com/articles/nphoton.2011.35. doi:10.1038/nphoton.2011.35, publisher: Nature Publishing Group.
[16] M. Chehimi, S. Y.-C. Chen, W. Saad, S. Yoo, Federated Quantum Long Short-term Memory (FedQLSTM), 2023. URL: http://arxiv.org/abs/2312.14309. doi:10.48550/arXiv.2312.14309, arXiv:2312.14309 [quant-ph].
[17] L. Pira, C. Ferrie, An invitation to distributed quantum neural networks, Quantum Machine Intelligence 5 (2023) 23. URL: https://doi.org/10.1007/s42484-023-00114-3. doi:10.1007/s42484-023-00114-3.
[18] R. Ballester, A Comprehensive Literature Review on Quantum Federated Learning, 2025. URL: https://zenodo.org/records/15105827. doi:10.5281/zenodo.15105827.
[19] S. Y.-C. Chen, S. Yoo, Federated Quantum Machine Learning, 2021. URL: http://arxiv.org/abs/2103.12010, arXiv:2103.12010 [quant-ph].

[20] J. Wu, T. Hu, Q. Li, Distributed Quantum Machine Learning: Federated and Model-Parallel Approaches, IEEE Internet Computing 28 (2024) 65–72. URL: https://ieeexplore.ieee.org/document/10508212/?arnumber=10508212. doi:10.1109/MIC.2024.3361288, conference Name: IEEE Internet Computing.
[21] W. Li, S. Lu, D.-L. Deng, Quantum federated learning through blind quantum computing, Science China Physics, Mechanics & Astronomy 64 (2021) 100312. URL: https://doi.org/10.1007/s11433-021-1753-3. doi:10.1007/s11433-021-1753-3.
[22] B. Polacchi, D. Leichtle, L. Limongi, G. Carvacho, G. Milani, N. Spagnolo, M. Kaplan, F. Sciarrino, E. Kashefi, Multi-client distributed blind quantum computation with the Qline architecture, Nature Communications 14 (2023) 7743. URL: https://doi.org/10.1038/s41467-023-43617-0. doi:10.1038/s41467-023-43617-0.
[23] Q. Xia, Z. Tao, Q. Li, Defending Against Byzantine Attacks in Quantum Federated Learning, in: 2021 17th International Conference on Mobility, Sensing and Networking (MSN), 2021, pp. 145–152. URL: https://ieeexplore.ieee.org/document/9751508/?arnumber=9751508. doi:10.1109/MSN53354.2021.00035.
[24] W. Yamany, N. Moustafa, B. Turnbull, OQFL: An Optimized Quantum-Based Federated Learning Framework for Defending Against Adversarial Attacks in Intelligent Transportation Systems, IEEE Transactions on Intelligent Transportation Systems 24 (2023) 893–903. URL: https://ieeexplore.ieee.org/document/9641742/. doi:10.1109/TITS.2021.3130906, conference Name: IEEE Transactions on Intelligent Transportation Systems.
[25] Y. Zhang, C. Zhang, C. Zhang, L. Fan, B. Zeng, Q. Yang, Federated Learning with Quantum Secure Aggregation, 2023. URL: http://arxiv.org/abs/2207.07444. doi:10.48550/arXiv.2207.07444, arXiv:2207.07444 [quant-ph].
[26] L. Sünkel, P. Altmann, M. Kölle, T. Gabor, Quantum Federated Learning for Image Classification:, in: Proceedings of the 16th International Conference on Agents and Artificial Intelligence, SCITEPRESS - Science and Technology Publications, Rome, Italy, 2024, pp. 936–942. URL: https://www.scitepress.org/DigitalLibrary/Link.aspx?doi=10.5220/0012421200003636. doi:10.5220/0012421200003636.
[27] H. Lee, S. Park, Performance Analysis of Quantum Federated Learning in Data Classification, The Journal of Korean Institute of Communications and Information Sciences 49 (2024) 264–269. URL: https://www.dbpia.co.kr. doi:10.7840/kics.2024.49.2.264.
[28] J. Shi, T. Chen, S. Zhang, X. Li, Personalized Quantum Federated Learning for Privacy Image Classification, 2024. URL: http://arxiv.org/abs/2410.02547. doi:10.48550/arXiv.2410.02547, arXiv:2410.02547.
[29] D. Gurung, S. R. Pokhrel, A Personalized Quantum Federated Learning, in: Proceedings of the 8th Asia-Pacific Workshop on Networking, ACM, Sydney Australia, 2024, pp. 175–176. URL: https://dl.acm.org/doi/10.1145/3663408.3665806. doi:10.1145/3663408.3665806.
[30] T. Wang, H.-H. Tseng, S. Yoo, Quantum Federated Learning with Quantum Networks, in: ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2024, pp. 13401–13405. URL: https://ieeexplore.ieee.org/abstract/document/10447516. doi:10.1109/ICASSP48485.2024.10447516, iSSN: 2379-190X.
[31] D. De, M. Nandy Pal, D. Hazra, FQPDR: federated quantum neural network for privacy-preserving early detection of diabetic retinopathy, Evolutionary Intelligence (2024). URL: https://doi.org/10.1007/s12065-024-00971-2. doi:10.1007/s12065-024-00971-2.
[32] Z. Qu, L. Zhang, P. Tiwari, Quantum Fuzzy Federated Learning for Privacy Protection in Intelligent Information Processing, IEEE Transactions on Fuzzy Systems (2024) 1–12. URL: https://ieeexplore.ieee.org/document/10572363/?arnumber=10572363. doi:10.1109/TFUZZ.2024.3419559, conference Name: IEEE Transactions on Fuzzy Systems.
[33] Y. Song, Y. Wu, S. Wu, D. Li, Q. Wen, S. Qin, F. Gao, A quantum federated learning framework for classical clients, Science China Physics, Mechanics & Astronomy 67 (2024) 250311. URL: https://doi.org/10.1007/s11433-023-2337-2. doi:10.1007/s11433-023-2337-2.
[34] W. J. Yun, J. P. Kim, S. Jung, J. Park, M. Bennis, J. Kim, Slimmable quantum federated learning,

arXiv preprint arXiv:2207.10221 (2022).
[35] M. Chehimi, W. Saad, Quantum Federated Learning with Quantum Data, in: ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2022, pp. 8617–8621. doi:10.1109/ICASSP43922.2022.9746622, iSSN: 2379-190X.
[36] S. Rayhan, An Approach to Work with Quantum Data in Federated Learning, 2023.
[37] A. Chefles, Quantum State Discrimination, Contemporary Physics 41 (2000) 401–424. URL: http://arxiv.org/abs/quant-ph/0010114. doi:10.1080/00107510010002599, arXiv:quant-ph/0010114.
[38] J. A. Bergou, Discrimination of quantum states, Journal of Modern Optics 57 (2010) 160–180. URL: https://doi.org/10.1080/09500340903477756. doi:10.1080/09500340903477756, publisher: Taylor & Francis _eprint: https://doi.org/10.1080/09500340903477756.
[39] J. Bae, L.-C. Kwek, Quantum state discrimination and its applications, 2017. URL: https://arxiv.org/abs/1707.02571v2. doi:10.1088/1751-8113/48/8/083001.
[40] H. Chen, L. Wossnig, S. Severini, H. Neven, M. Mohseni, Universal discriminative quantum neural networks, Quantum Machine Intelligence 3 (2021) 1. URL: http://arxiv.org/abs/1805.08654. doi:10.1007/s42484-020-00025-7, arXiv:1805.08654 [quant-ph].
[41] D. Lee, K. Baek, J. Huh, D. K. Park, Variational quantum state discriminator for supervised machine learning, 2023. URL: https://arxiv.org/abs/2303.03588v1. doi:10.1088/2058-9565/ad0a05.
[42] A. Patterson, H. Chen, L. Wossnig, S. Severini, D. Browne, I. Rungger, Quantum state discrimination using noisy quantum neural networks, Physical Review Research 3 (2021) 013063. URL: https://link.aps.org/doi/10.1103/PhysRevResearch.3.013063. doi:10.1103/PhysRevResearch.3.013063.
[43] C. W. Helstrom, Quantum detection and estimation theory, Journal of Statistical Physics 1 (1969) 231–252. URL: https://doi.org/10.1007/BF01007479. doi:10.1007/BF01007479.