Distributed PAC Learning from Quantum Data with Efficient Communication

Youming Tao[1], Yaowu Yang[2], Shuzhen Chen[3], Huiqun Li[4], Congwei Zhang[4], and Falko Dressler[1]

[1] TU Berlin, Berlin, Germany
{tao, dressler}@ccs-labs.org
[2] Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China
yangyaowu25s@ict.ac.cn
[3] Ocean University of China, Qingdao, China
szchen@ouc.edu.cn
[4] Shandong University, Qingdao, China
{lhqlee, xczhw}@mail.sdu.edu.cn

Abstract. We study distributed PAC learning from quantum data under classical communication constraints, where multiple clients collaboratively learn classical patterns from labeled quantum states. The key question is what collaboration can fundamentally gain in this setting. We show that collaboration can reduce the per-client quantum sample complexity by a linear factor of $1/K$ for achieving the same target error as centralized QPAC learning. To establish this result, we first propose `DQERM-AVG`, a distributed quantum empirical risk minimization algorithm that achieves this per-client reduction. We then characterize the communication required to obtain such a gain by proving a lower bound of $\Omega(|\mathcal{C}| + K)$ within our measurement-partitioning framework. Motivated by this result, we further develop `DQERM-RR`, which preserves the same per-client sample complexity while reducing the communication cost from $O(|\mathcal{C}|K)$ to the optimal $O(|\mathcal{C}| + K)$. Experiments on a quantum state classification task further support the theoretical findings.

Keywords: Distributed Learning · Quantum Data · PAC Learning.

1 Introduction

Recent advances in quantum sensor networks [10] and quantum edge computing networks [20] are enabling quantum data to be generated and processed at the network edge. As a result, learning from quantum data has become increasingly important. However, due to the fragile nature of quantum data as well as data privacy concerns, collecting all quantum data for centralized learning is highly impractical. On one side, transmitting qubits is inherently vulnerable to decoherence, leading to qubit degradation caused by environmental interactions, which significantly shortens their lifespan [23]. On the other side, data sharing causes privacy leakage so data owners can be reluctant to do so [25]. In addition, each individual device typically possesses only a limited amount of quantum data [26], due to constraints such as short coherence time, low sampling rate, or hardware limitations. These challenges motivate distributed learning methods that keep quantum data local while enabling collaboration across multiple devices.

While distributed quantum learning has attracted growing attention, existing efforts have focused primarily on the design and implementation of learning protocols [5, 19, 25, 27], rather than on their statistical foundations. In particular, the sample complexity required to identify a predictor that generalizes well remains open. This issue is especially critical in the quantum setting, where training samples are costly to generate, fragile under measurement, and therefore inherently limited. A natural framework for addressing this issue is quantum PAC (QPAC) learning [12–14, 24]. In QPAC, the learner is given classically-labeled quantum states and seeks a quantum measurement from a concept class $\mathcal{C}$ whose loss is close to that of the best candidate in $\mathcal{C}$. Given a target error ϵ and failure probability δ, the quantum sample complexity $n_{\mathcal{C}}(\epsilon, \delta)$ is the minimum number of training samples required to identify such a predictor with probability at least $1 - \delta$. Existing QPAC results, however, are largely confined to the centralized setting. In particular, for finite concept classes, compatibility-based measurement partitioning yields the sample complexity bound [12] $n_{\mathcal{C}}(\epsilon, \delta) \leq O(\frac{m}{\epsilon^2} \log \frac{|\mathcal{C}|}{\delta})$, where m denotes the minimum number of compatibility partition bins.

We consider a distributed quantum learning system consisting of a central coordinator and K clients. The coordinator stores the concept class $\mathcal{C}$, while each client holds n classically-labeled quantum states. Our goal is to understand the fundamental benefit of collaboration in this setting. In classical collaborative learning, collaboration among K learners can reduce the sample requirement of each learner by a linear factor of $1/K$ for achieving the same target accuracy, which is the optimal gain brought by collaboration [7]. This naturally raises the question of whether an analogous benefit can be achieved in distributed QPAC learning. In particular, we ask whether, by collaborating with other clients, each client can attain the same target error using only a $1/K$ fraction of the quantum samples required when learning alone under the centralized QPAC guarantee. This motivates our first question:

> ***Q1.*** *Can collaboration among K clients reduce the per-client sample complexity in distributed QPAC learning by a linear factor of $1/K$ compared with centralized QPAC learning?*

Beyond sample efficiency, communication cost is another fundamental concern in distributed learning. In this work, we focus on classical communication, since high-fidelity quantum communication across distributed processors remains technologically challenging and is currently limited to proof-of-principle demonstrations [8, 30]. This leads to our second question:

> ***Q2.*** *What classical communication cost is necessary to achieve such a linear per-client sample reduction, and can it be attained optimally?*

By addressing **Q1** and **Q2**, this work initiates the study of distributed QPAC learning. Our main contributions are summarized as follows:

Table 1. Summary of our contributions and comparison with existing works on PAC learning. We hide all constants under big-O notation. $\mathcal{C}$ is the concept class of all candidate predictors, m is the (minimum) number of bins generated by compatible partition, ϵ is the target error, and δ is the target failure probability.

Algorithm	Sample Complexity	Communication Overheads	Quantum?	Distributed?
Classical methods [3,4,6,22]	$O(\frac{1}{\epsilon^2}\log\frac{\lvert\mathcal{C}\rvert}{\delta})$	NA	✗	✗
`QERM` [12,24]	$O(\frac{m}{\epsilon^2}\log\frac{\lvert\mathcal{C}\rvert}{\delta})$	NA	✓	✗
`DQERM-AVG` (Our Algo. 1)	$O(\frac{m}{K\epsilon^2}\log\frac{\lvert\mathcal{C}\rvert}{\delta})$ (Thm. 1)	$O(\lvert\mathcal{C}\rvert K)$ (Thm. 2)	✓	✓
`DQERM-RR` (Our Algo. 2)	$O(\frac{1}{\epsilon^2}\lceil\frac{m}{K}\rceil\log\frac{\lvert\mathcal{C}\rvert}{\delta})$ (Thm. 4)	$\Theta(\lvert\mathcal{C}\rvert+K)$ (Thm. 3 and Thm. 5)	✓	✓

– We answer **Q1** affirmatively by designing `DQERM-AVG`, a distributed QPAC learning algorithm based on local quantum empirical risk minimization (QERM) and global averaging. We prove that `DQERM-AVG` reduces the per-client sample complexity by a linear factor of $1/K$ compared with centralized QPAC learning with communication cost $O(|\mathcal{C}|K)$, where K is # client.
– To address **Q2**, we first establish a communication lower bound of $\Omega(|\mathcal{C}|+K)$ for achieving the $1/K$ per-client sample reduction within the compatibility-based partition framework. This reveals a communication gap. We then develop `DQERM-RR`, which distributes subclasses of compatible predictors among clients using a round-robin allocation strategy. We prove that `DQERM-RR` preserves the same sample benefit as `DQERM-AVG` while attaining the optimal communication cost. Our theoretical results are summarized in Table 1.
– We further conduct numerical experiments on a quantum state classification task. The results corroborate our theoretical findings.

Due to space limitations, all proofs are deferred to the supplementary appendix.

2 Problem Formulation

System Model. We consider a distributed quantum learning system with one central coordinator and K clients. Each client communicates only with the coordinator, and there is no direct communication among clients. We account for both uplink and downlink communication. Following the standard communication model in distributed learning [11, 17], we measure communication cost in *words*, where each transmitted numerical value incurs one word.

Quantum Data Model. For any positive integer z, let $[z] \triangleq \{1, \ldots, z\}$. Let $\mathcal{X}$ denote the feature space of quantum states. For each $x \in \mathcal{X}$, let ρ_x be the density operator describing the corresponding quantum state on a Hilbert space H_X. Density operators are linear, self-adjoint, unit-trace, and positive semi-definite on H_X. Each quantum state is associated with a classical label. Let $\mathcal{Y}$ denote the finite classical label space. For compactness, we introduce an auxiliary quantum register to store the classical labels, i.e., $|y\rangle$ for each $y \in \mathcal{Y}$. Let $H_Y = \text{span}\{|y\rangle : y \in \mathcal{Y}\}$ denote the Hilbert space of the labels. Then any feature-label pair $(x, y) \in \mathcal{X} \times \mathcal{Y}$ is represented by the density operator $\rho_{xy} = \rho_x \otimes |y\rangle\langle y|$ on the joint Hilbert space $H_X \otimes H_Y$.

Distributed Learning Setup. Each client $k \in [K]$ holds a local training dataset $D^{(k)}$, denoted as $D^{(k)} = \{\rho_i^{(k)} \otimes |y_i^{(k)}\rangle\langle y_i^{(k)}|\}_{i=1}^n$. Let $D = \bigcup_{k\in[K]} D^{(k)}$

denote the full distributed training set. All samples are drawn i.i.d. from an unknown distribution $\mathcal{D}$ over $\mathcal{X} \times \mathcal{Y}$. The average density operator of samples drawn from distribution $\mathcal{D}$ is defined as $\rho_{XY} \triangleq \sum_{x\in\mathcal{X},\, y\in\mathcal{Y}} \mathcal{D}(x,y)\, \rho_x \otimes |y\rangle\langle y|$. Note that the quantum states in the training set can be either *pure* or *mixed* states. The coordinator stores a concept class $\mathcal{C}$ of candidate predictors, indexed by the set $\mathcal{J}$. Each predictor is a quantum measurement modeled by a positive operator-valued measure (POVM). Let $\hat{y} \in \mathcal{Y}$ be the prediction output by the quantum measurement. A POVM is denoted by $\mathcal{M} = \{M_{\hat{y}} : \hat{y} \in \mathcal{Y}\}$, where $M_{\hat{y}} = M_{\hat{y}}^\dagger \succeq 0$ for all $\hat{y} \in \mathcal{Y}$ and $\sum_{\hat{y}\in\mathcal{Y}} M_{\hat{y}} = I$. For a quantum state ρ_x, the probability of predicting label $\hat{y}$ under $\mathcal{M}$ is given by the Born rule, i.e., $\Pr(\hat{y}|x,\mathcal{M}) = \mathrm{tr}(\rho_x M_{\hat{y}})$. The learning algorithm can be modeled as a quantum measurement on the joint space of all training the samples, i.e., $H_{XY}^{\otimes Kn}$. The outcome of this measurement is a classical number in $\mathcal{J}$ indicating the index of the selected predictor in $\mathcal{C}$. We evaluate a predictor using a normalized loss function $\ell : \mathcal{Y}\times\mathcal{Y} \to [0,1]$. The population risk of a measurement $\mathcal{M}$ with respect to the underlying distribution $\mathcal{D}$ is $L_{\mathcal{D}}(\mathcal{M}) \triangleq \mathbb{E}[\ell(y,\hat{y})]$.

QPAC Learning Goal. Our goal is to learn a predictor from $\mathcal{C}$ that minimizes the population risk $L_{\mathcal{D}}$. However, due to the *no free lunch theorem* [28], one cannot in general identify the optimal predictor in $\mathcal{C}$ exactly from finite samples. Accordingly, we aim to output a measurement whose population risk is within ϵ of the best predictor in $\mathcal{C}$ with probability at least $1-\delta$.

Definition 1 ((ϵ,δ,K)-Distributed QPAC Learning). *A distributed quantum learning algorithm $\mathcal{A}$ agnostically (ϵ,δ)-QPAC learns a concept class $\mathcal{C}$ if there exists a function $n_{\mathcal{C}} : (0,1)\times(0,1)\times\mathbb{N} \to \mathbb{N}$ such that for every $\epsilon,\delta \in (0,1)$, whenever each of K clients is given at least $n_{\mathcal{C}}(\epsilon,\delta,K)$ i.i.d. training samples from $\mathcal{D}$, algorithm $\mathcal{A}$ outputs, with probability at least $1-\delta$, a measurement $\mathcal{M} \in \mathcal{C}$ satisfying $L_{\mathcal{D}}(\mathcal{M}) \le \inf_{\mathcal{M}'\in\mathcal{C}} L_{\mathcal{D}}(\mathcal{M}') + \epsilon$. The smallest such quantity $n_{\mathcal{C}}(\epsilon,\delta,K)$ is called the per-client quantum sample complexity of $\mathcal{C}$.*

3 The `DQERM-AVG` Algorithm

Design of `DQERM-AVG`. The central task in QPAC learning is to estimate the population risk $L_{\mathcal{D}}(\mathcal{M})$ for each predictor $\mathcal{M} \in \mathcal{C}$. We follow the framework of [12] and use empirical risk as its estimator. Let $\mathcal{Z}$ denote the image of the loss function ℓ, i.e., $\mathcal{Z} \triangleq \{\ell(y,\hat{y}) : y,\hat{y} \in \mathcal{Y}\}$, which is finite since $\mathcal{Y}$ is finite. For any predictor $\mathcal{M} = \{M_{\hat{y}} : \hat{y} \in \mathcal{Y}\}$, define the associated loss observable $\mathcal{L}_{\mathcal{M}} \triangleq \{L_z^{\mathcal{M}} : z \in \mathcal{Z}\}$, where $L_z^{\mathcal{M}} \triangleq \sum_{y,\hat{y}\in\mathcal{Y}:\, \ell(y,\hat{y})=z} M_{\hat{y}} \otimes |y\rangle\langle y|, \quad \forall z \in \mathcal{Z}$. Applying $\mathcal{L}_{\mathcal{M}}$ to a labeled sample $\rho_x \otimes |y\rangle\langle y|$ produces a random loss value $Z = \ell(y,\hat{Y}) \in \mathcal{Z}$, where $\hat{Y}$ is the label output by $\mathcal{M}$. For a realization of this measurement outcome, we write $z^{\mathcal{M}}(x,y)$. Then, for any dataset $\hat{D}$, the empirical risk of $\mathcal{M}$ is $\bar{L}_{\hat{D}}(\mathcal{M}) \triangleq \frac{1}{|\hat{D}|}\sum_{(x,y)\in\hat{D}} z^{\mathcal{M}}(x,y)$. The corresponding population risk can be written as $L_{\mathcal{D}}(\mathcal{M}) = \sum_{z\in\mathcal{Z}} z\, \mathrm{tr}\big(L_z^{\mathcal{M}}\rho_{XY}\big)$. Unlike classical ERM [6], where the same sample can be evaluated against all candidate predictors, quantum training samples cannot be reused after measurement due to state collapse, nor can they be copied because of the no-cloning theorem [29]. To evaluate multiple predictors on the same quantum sample, we exploit measurement compatibility.

Definition 2 (Compatibility [15]). *A collection of sharp*[5] *measurements* $\{\mathcal{M}_j \triangleq \{M_y^j : y \in \mathcal{Y}\}\}_{j\in\mathcal{J}}$ *is compatible if their operators mutually commute, i.e.,* $M_y^j M_{y'}^{j'} = M_{y'}^{j'} M_y^j$*, for all* $y, y' \in \mathcal{Y}$ *and* $j, j' \in \mathcal{J}$*.*

We partition $\mathcal{C}$ into subclasses whose predictors are mutually compatible.

Definition 3 (Compatibility Partition [12]). *Given a concept class* $\mathcal{C}$*, a compatibility partition is a family of disjoint subsets* $C_1, \ldots, C_m$ *such that* $\mathcal{C} = \biguplus_{r=1}^m C_r$*, and all predictors within each* C_r *are mutually compatible.*

Such a partition always exists, e.g., by taking singleton subsets. In this paper, we focus on a partition with the minimum number of subclasses, denoted by m. Computing an exact minimum partition is generally hard, but since compatibility depends only on $\mathcal{C}$, the partition can be computed offline by the coordinator. For each compatible subclass C_r, the corresponding loss observables $\{\mathcal{L}_{\mathcal{M}} : \mathcal{M} \in C_r\}$ are jointly measurable. Hence, we can construct a joint POVM $\mathcal{L}_{\text{QERM}}^{C_r} \triangleq \{L_{\mathbf{z}}^{C_r} = \prod_{j\in\mathcal{J}_{C_r}} L_{z_j}^{\mathcal{M}_j} : \mathbf{z} \in \mathcal{Z}^{|C_r|}\}$, where $\mathcal{J}_{C_r}$ is the index set of predictors in C_r, $\mathbf{z} = (z_1, \ldots, z_{|C_r|})$, and $\{L_z^{\mathcal{M}_j} : z \in \mathcal{Z}\}$ are the elements of $\mathcal{L}_{\mathcal{M}_j}$. This construction allows a single training sample to evaluate the empirical losses of all predictors in C_r simultaneously. In the distributed setting, the remaining challenge is how to leverage quantum samples stored across different clients to perform QERM for predictors in every compatible subclass $C_r \subset \mathcal{C}$. To address this, we adopt distributed averaging, a standard idea in distributed learning [9, 21], and combine it with the QERM procedure. This leads to the proposed `DQERM-AVG` in Algorithm 1. In `DQERM-AVG`, the coordinator distributes the QERM POVMs to the clients, rather than the predictor subclasses obtained from the compatibility partition. This design reduces both computation and communication costs. If the subclasses were sent instead, each client would need to construct the corresponding QERM POVMs locally, which is computationally inefficient, especially since the coordinator is typically more powerful than the clients. In addition, transmitting subclasses would require sending at least $|\mathcal{C}|$ predictor operators, whereas transmitting the QERM POVMs requires sending only $m \leq |\mathcal{C}|$ measurement operators.

Sample Complexity of `DQERM-AVG`. Our analysis relies on a quantum analogue of the Chernoff–Hoeffding inequality (Lemma 1), which controls the deviation between the empirical mean of measurement outcomes and their expectation. Based on that, we control the error empirical risk estimation in Lemma 2, which then immediately yields the sample complexity guarantee in Theorem 1

Lemma 1 (Quantum Chernoff–Hoeffding inequality [1, Theorem 19]). *Let* $\rho_1, \rho_2, \cdots, \rho_n$ *be i.i.d. random density operators on a finite-dimensional Hilbert space* H*. Let* $\bar{\rho} = \mathbb{E}[\rho_i]$ *be their average density operator. Let* $\mathcal{M}$ *be a discrete observable on* H *with outcomes in* $[a, b] \subset \mathbb{R}$*. If* Z_i *denotes the outcome of measuring* ρ_i *by* $\mathcal{M}$*, then for any* $t \geq 0$*,* $\mathbb{P}\left(\left|\frac{1}{n}\sum_{i=1}^n Z_i - \langle\mathcal{M}\rangle_{\bar{\rho}}\right| \geq t\right) \leq 2\exp(-\frac{nt^2}{2(b-a)})$*.*

[5] We follow the standard formulation of sharp quantum measurements.

Algorithm 1: `DQERM-AVG`

Input: Concept class $\mathcal{C}$, K clients each with n training samples.
Output: Index of the selected predictor in $\mathcal{C}$.

1 **Coordinator** partitions $\mathcal{C}$ into a family of compatible subclasses such that $\mathcal{C} = C_1 \uplus C_2 \uplus \cdots \uplus C_m$;
2 **Coordinator** constructs the POVM $\mathcal{L}_{\text{QERM}}^{C_r}$ for every subclass C_r as $\mathcal{L}_{\text{QERM}}^{C_r} \triangleq \{L_{\mathbf{z}}^{C_r} = \prod_{j \in \mathcal{J}_{C_r}} L_{z_j}^{\mathcal{M}_j} : \mathbf{z} \in \mathcal{Z}^{|C_r|}\}$;
3 **Coordinator** broadcasts $\{\mathcal{L}_{\text{QERM}}^{C_r}\}_{r=1}^{m}$ to all clients ;
4 **for every** ***client* k** **in parallel do**
5 Partition local training dataset $D^{(k)}$ into m parts evenly such that $D^{(k)} = D_1^{(k)} \uplus D_2^{(k)} \uplus \cdots \uplus D_m^{(k)}$ with each $|D_r^{(k)}| = n/m$;
6 Apply each $\mathcal{L}_{\text{QERM}}^{C_r}$ to the corresponding $D_r^{(k)}$ and obtain empirical loss values $\{\mathbf{z}_r^{(k)}(x,y)\}_{(x,y)\in D_r^{(k)}}$;
7 Compute the local empirical risk $\bar{L}_{\text{local}}^{(k)}(\mathcal{M}_j)$ of every predictor $\mathcal{M}_j$ in each subclass C_r as $\bar{L}_{\text{local}}^{(k)}(\mathcal{M}_j) = \frac{m}{n}\sum_{(x,y)\in D_r^{(k)}} z_{r,j}^{(k)}(x,y)$;
8 Upload every $\bar{L}_{\text{local}}^{(k)}(\mathcal{M}_j), j \in [m]$ to the coordinator ;
9 **Coordinator** obtains the global empirical risk $\bar{L}_{\text{global}}(\mathcal{M}_j)$ for each predictor $\mathcal{M}_j, j \in [m]$ by aggregating all corresponding local empirical risks $\bar{L}_{\text{local}}^{(k)}(\mathcal{M}_j)$ via averaging as $\bar{L}_{\text{global}}(\mathcal{M}_j) = \frac{1}{K}\sum_{k=1}^{K} \bar{L}_{\text{local}}^{(k)}(\mathcal{M}_j)$;
10 **Coordinator** selects predictor $\mathcal{M}_{j^*}$ with its index $j^* = \arg\min_{j\in[m]} \bar{L}_{\text{global}}(\mathcal{M}_j)$;
11 **Coordinator** returns the index j^* of the selected predictor ;

Lemma 2 (Generalization Error of `DQERM-AVG`). *For any predictor $\mathcal{M}_j \in \mathcal{C}$, with probability at least $1-\zeta$, $\left|\bar{L}_{\text{global}}(\mathcal{M}_j) - L_{\mathcal{D}}(\mathcal{M}_j)\right| \le \sqrt{\frac{2m}{nK}\log\frac{2}{\zeta}}$.*

Theorem 1. *For any $\epsilon, \delta \in (0,1)$, any finite concept class $\mathcal{C}$ is (ϵ,δ)-distributed QPAC learnable via `DQERM-AVG` with per-client quantum sample complexity upper bounded by $n_{\mathcal{C}}(\epsilon,\delta,K) = \frac{8m}{K\epsilon^2}\log\frac{2|\mathcal{C}|}{\delta}$.*

Theorem 1 shows that `DQERM-AVG` achieves quantum sample complexity of $O(\frac{m}{K\epsilon^2}\log\frac{|\mathcal{C}|}{\delta})$. Compared with the centralized QPAC bound of $O(\frac{m}{\epsilon^2}\log\frac{|\mathcal{C}|}{\delta})$, this establishes an explicit linear $1/K$ reduction in per-client sample complexity.

Communication Overhead of `DQERM-AVG`. The total cost consists of both downlink and uplink communication. Since the Hilbert space dimension is fixed and standard POVM parameterizations use a bounded number of coefficients per operator [2, 18], the size of each $\mathcal{L}_{\text{QERM}}^{C_r}$ scales as $O(|C_r|)$ words.

Theorem 2. *The downlink communication cost of `DQERM-AVG` (Step 2) is $O(|\mathcal{C}|)$ words, and the uplink communication cost (Step 8) is $O(|\mathcal{C}|K)$ words. Therefore, the total communication cost of `DQERM-AVG` is $O(|\mathcal{C}|K)$ words.*

4 Communication-Optimal Distributed QPAC via the `DQERM-RR` Algorithm

Communication Lower Bound. We establish a lower bound on the communication for any compatibility-partition-based algorithms with $1/K$ per-client sample reduction. It shows that `DQERM-AVG` is not communication-optimal.

Theorem 3. *Any compatibility-partition-based distributed QPAC learning algorithm that achieves a $1/K$ per-client sample reduction incurs $\Omega(|\mathcal{C}|+K)$ words of communication.*

Design of `DQERM-RR`. The communication gap comes from the distributed averaging step in `DQERM-AVG`. There, each predictor is evaluated on local samples at every client, and the corresponding local empirical risks must then be uploaded and averaged at the coordinator. As a result, every client needs to report empirical risk values for all assigned predictors, leading to the extra communication cost. To remove this overhead, `DQERM-RR` assigns each compatible subclass C_r, $r \in [m]$, to a single client, so that the predictors in C_r are evaluated only once and no cross-client averaging is needed. We use a round-robin assignment rule: subclass C_r is assigned to client $(r-1) \bmod K+1$. Each client then selects the predictor with the smallest empirical risk among those assigned to it, and reports only its index and empirical risk value to the coordinator. In this way, the uplink communication is reduced from sending all predictor-wise empirical risks to sending only $O(1)$ information per client. The algorithm is detailed in Algorithm 2.

Sample Complexity of `DQERM-RR`. Although `DQERM-RR` differs from `DQERM-AVG` in how QERM evaluations are distributed across clients, both algorithms are based on the same empirical risk minimization principle. The key difference is only in task allocation, not in the statistical estimation itself. Consequently, `DQERM-RR` preserves the same order of per-client sample complexity as `DQERM-AVG`. We first bound the generalization error of `DQERM-RR` in Lemma 3. Then, by applying the same argument as in the proof of Theorem 1, we immediately obtain the sample complexity bound in Theorem 4.

Lemma 3 (Generalization Error of `DQERM-RR`). *For any predictor $\mathcal{M}_j \in \mathcal{C}$, with probability at least $1-\zeta$, $|\bar{L}(\mathcal{M}_j) - L_{\mathcal{D}}(\mathcal{M}_j)| \leq \sqrt{\frac{2}{n}\lceil \frac{m}{K} \rceil \log \frac{2}{\zeta}}$.*

Theorem 4. *For any $\epsilon, \delta \in (0,1)$, any finite concept class $\mathcal{C}$ is (ϵ,δ)-distributed QPAC learnable via `DQERM-RR` with per-client quantum sample complexity upper bounded by $n_{\mathcal{C}}(\epsilon,\delta,K) = \frac{8}{\epsilon^2}\left\lceil \frac{m}{K} \right\rceil \log \frac{2|\mathcal{C}|}{\delta}$.*

Compared with `DQERM-AVG`, the factor m/K is replaced by $\lceil m/K \rceil$, which does not change the order of the sample complexity. Thus, `DQERM-RR` retains the same per-client sample efficiency as `DQERM-AVG`.

Communication Overhead of `DQERM-RR`. We characterize the communication cost of `DQERM-RR` in Theorem 5

Theorem 5. *The downlink communication cost of `DQERM-RR` (Step 3) is $O(|\mathcal{C}|)$ words, and the uplink communication cost (Step 9) is $O(K)$ words. Therefore, the total communication cost of `DQERM-RR` is $O(|\mathcal{C}|+K)$ words.*

Algorithm 2: `DQERM-RR`

Input: Concept class $\mathcal{C}$, K clients each with n training samples.
Output: Index of the selected predictor in $\mathcal{C}$.

Coordinator partitions $\mathcal{C}$ into a family of compatible subclasses such that $\mathcal{C} = C_1 \uplus C_2 \uplus \cdots \uplus C_m$;
for every $r \in [m]$ **do**
  Coordinator constructs the POVM $\mathcal{L}_{\text{QERM}}^{C_r}$ for every subclass C_r as $\mathcal{L}_{\text{QERM}}^{C_r} \triangleq \{L_{\mathbf{z}}^{C_r} = \prod_{j \in \mathcal{J}_{C_r}} L_{z_j}^{\mathcal{M}_j} : \mathbf{z} \in \mathcal{Z}^{|C_r|}\}$;
  Coordinator sends $\mathcal{L}_{\text{QERM}}^{C_r}$ to the client indexed by $(r-1) \bmod K + 1$;
for every *client* **k in parallel do**
  Partition local training dataset $D^{(k)}$ into $\lceil \frac{m}{K} \rceil$ parts evenly such that $D^{(k)} = D_1^{(k)} \uplus D_2^{(k)} \uplus \cdots \uplus D_{\lceil m/K \rceil}^{(k)}$ with each $|D_r^{(k)}| = n / \lceil \frac{m}{K} \rceil$;
  Apply each $\mathcal{L}_{\text{QERM}}^{C_r}$ assigned to the corresponding $D_r^{(k)}$ and obtain empirical loss values $\{\mathbf{z}_r^{(k)}(x,y)\}_{(x,y) \in D_r^{(k)}}$;
  Compute the empirical risk $\bar{L}(\mathcal{M}_j)$ of every predictor $\mathcal{M}_j$ in each subclass C_r as $\bar{L}(\mathcal{M}_j) = \frac{m}{n} \sum_{(x,y) \in D_r^{(k)}} z_{r,j}^{(k)}(x,y)$;
  Select predictor $\mathcal{M}_{j^{(k)*}}$ with its index $j^{(k)*} = \arg\min_{j \in [m]} \bar{L}(\mathcal{M}_j)$;
  Upload $j^{(k)*}$ and $\bar{L}_{\text{local}}^{(k)}(\mathcal{M}_{j^{(k)*}})$ to the coordinator
Coordinator selects predictor $\mathcal{M}_{j^*}$ with its index $j^* = \arg\min_{j \in \{j^{(1)*}, j^{(2)*}, \cdots, j^{(K)*}\}} \bar{L}(\mathcal{M}_j)$;
Coordinator returns the index j^* of the selected predictor ;

Combining Theorem 3 and Theorem 5, we conclude that `DQERM-RR` is communication-optimal within the compatibility-partition framework.

5 Numerical Simulation

Simulation Setup. We use the open-source Qiskit SDK [16] to simulate the quantum learning environment. Our task is a binary quantum state classification problem based on electron spin orientations. Specifically, each sample corresponds to a spin state whose axis is indexed by $\mathcal{X} = \left\{(\theta_i, \phi_j) = \left(\frac{i\pi}{11}, \frac{2j\pi}{11}\right) : 0 \le i, j \le 10\right\}$, which forms a finite grid on the Bloch sphere. The label space is $\mathcal{Y} = \{\text{blue}, \text{red}\}$. A sample is labeled *blue* if its spin axis lies in the orthant defined by $0 \le \theta \le \frac{\pi}{2}$, $0 \le \phi < \frac{\pi}{2}$, and *red* otherwise. To generate quantum states, we first sample random pure states on the two-dimensional Hilbert space H_X. Specifically, we draw a Haar-random unitary matrix U and take its first column as a state vector ψ, from which we construct the pure-state density matrix $\rho = \psi\psi^\dagger$. To model imperfect state preparation, we then perturb ρ by adding a small random complex matrix $\Delta \in \mathbb{C}^{2\times 2}$, whose real and imaginary parts are drawn i.i.d. from the uniform distribution $\mathcal{U}(-\iota/2, \iota/2)$, where $\iota = 0.01$ controls the noise magnitude. The perturbed state is then projected back to a valid density operator through Hermitian symmetrization, removal of negative eigenvalues, and trace normalization. In this way, we obtain i.i.d. mixed states that model noisy quantum samples. The concept class consists of binary POVMs. Each predictor $\mathcal{M}_j$ is

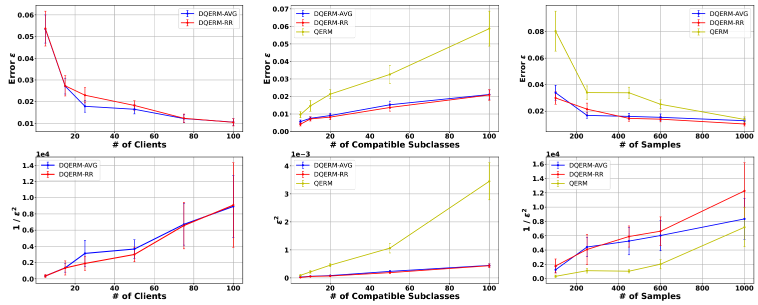

(a) (b) (c)

Fig. 1. Impact of # clients K, # partitioned subclasses m and # local sample n on error $\epsilon = \bar{L}(\mathcal{M}_{j^*}) - L_{\mathcal{D}}(\mathcal{M}_{j^*})$. **(a)**: ϵ vs. K. **(b)**: ϵ vs. $|\mathcal{C}|$. **(c)**: ϵ vs. n.

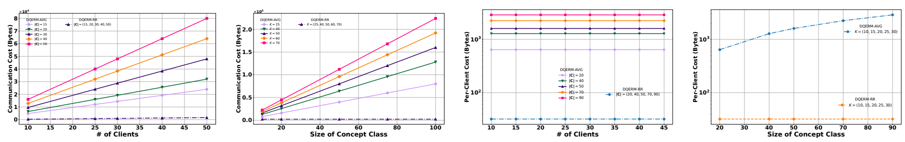

(a) (b) (c) (d)

Fig. 2. Communication cost comparison of `DQERM-AVG` and `DQERM-RR`. **(a):** Total cost vs. # clients K. **(b):** Total cost vs. concept class size $|\mathcal{C}|$. **(c):** Per-client cost vs. K. **(d):** Per-client cost vs. $|\mathcal{C}|$.

constructed as $\mathcal{M}_j = \{M_{j,0}, M_{j,1}\}$, where $M_{j,0}$ is generated from a random Hermitian matrix with eigenvalues clipped to $[0,1]$, and $M_{j,1} = I - M_{j,0}$. For loss function, we use the 0-1 loss.

Configuration and Baseline Methods. As this is the first work on distributed QPAC learning, we compare `DQERM-AVG` and `DQERM-RR` with the centralized QPAC baseline `QERM` from [12]. Unless otherwise specified, we use the default configuration $(K, m, n) = (25, 20, 200)$. Since the underlying sample distribution is not available in closed form, we approximate the population risk of each predictor by its empirical risk computed from 12,000 samples, which is much larger than the training size. For each data point, we report the average over 90 independent runs, with standard deviations shown as error bars. To study the impact of different parameters, we will vary one parameter at a time while keeping the others fixed.

Learning Error Evaluation. We first evaluate the learning performance of `DQERM-AVG`, `DQERM-RR`, and `QERM` using the learning error $\epsilon \triangleq \bar{L}(\mathcal{M}_{j^*}) - L_{\mathcal{D}}(\mathcal{M}_{j^*})$. We study the effect of the number of clients K for `DQERM-AVG` and `DQERM-RR`, and the effect of the number of compatible subclasses m and the number of local samples n for all three algorithms. The results are shown in Fig. 1.

Impact of # Clients. Fig. 1(a) shows that the learning error decreases as the number of clients increases for both `DQERM-AVG` and `DQERM-RR`. This is consistent with the fact that collaboration enables each client to benefit from a larger effective sample pool. Fig. 1(a) shows that $1/\epsilon^2$ grows approximately linearly with K, matching the theoretical scaling $\epsilon \propto 1/\sqrt{K}$ predicted by Lemma 2 and 3.

Impact of # Compatible Subclasses. Fig. 1(b) shows that the learning error increases with the number of compatible subclasses. Intuitively, a larger m means that the concept class is partitioned into more subclasses, so fewer samples are effectively used to evaluate each subclass, resulting in less accurate empirical risk estimation. Fig. 1(b) shows that ϵ^2 grows approximately linearly with m, which is consistent with the theoretical dependence $\epsilon \propto \sqrt{m}$. In addition, `DQERM-AVG` and `DQERM-RR` achieve lower learning error than the centralized baseline `QERM` across all tested configurations, reflecting the benefit of distributed collaboration.

Impact of # Local Sample. Fig. 1(c) shows that the learning error decreases as the local sample size increases. This is expected, since a larger training set yields a more accurate approximation of the population risk. Fig. 1(c) shows that $1/\epsilon^2$ grows approximately linearly with n, in agreement with the theoretical scaling $\epsilon \propto 1/\sqrt{n}$ from Lemma 2 and 3. As in the previous experiment, both `DQERM-AVG` and `DQERM-RR` consistently achieve lower learning error than `QERM`, again showing the benefit of collaboration.

Communication Cost Evaluation. Figure 2 compares the communication cost of `DQERM-AVG` and `DQERM-RR`. Since both algorithms have the same downlink scaling, we focus on uplink communication, which is also the more practically critical component. We examine its dependence on the number of clients K and the concept class size $|\mathcal{C}|$.

Communication Cost vs. # Clients. Fig. 2(a) shows that the total communication cost increases with the number of clients for both algorithms, but much more rapidly for `DQERM-AVG`. This is because `DQERM-AVG` requires every client to upload empirical risk values for all predictors, whereas `DQERM-RR` requires each client to report only the best predictor and its empirical risk among the assigned subclasses. As a result, `DQERM-RR` maintains substantially lower total communication cost. Fig. 2(c) shows that the per-client communication cost remains essentially constant w.r.t. K for both algorithms, but at a much smaller level for `DQERM-RR`.

Communication Cost vs. Concept Class Size. Fig. 2(b) shows that the total communication cost of `DQERM-AVG` increases linearly with $|\mathcal{C}|$, and the growth becomes more pronounced as the number of clients increases. In contrast, the total communication cost of `DQERM-RR` remains nearly unchanged as $|\mathcal{C}|$ varies, since each client communicates only a constant amount of information. Fig. 2(d) shows the same trend at the per-client level: the communication cost of `DQERM-AVG` scales linearly with $|\mathcal{C}|$, whereas that of `DQERM-RR` remains essentially constant. These results corroborate the theoretical communication advantage of `DQERM-RR`.

6 Conclusion

This work establishes a theoretical foundation for distributed learning from quantum data under the QPAC framework. We show that distributed quantum learning can achieve the same linear sample-saving factor as classical collaborative learning, despite the distinctive challenges imposed by quantum measurements, state collapse, and the no-cloning theorem. Beyond feasibility, we characterize the communication requirements of distributed QPAC learning. Our analysis shows that achieving optimal sample efficiency requires communication proportional to the concept class size and the number of clients. The DQERM-RR algorithm attains this lower bound while preserving the statistical guarantees of distributed empirical risk minimization. More broadly, this study illustrates how statistical learning theory can inform the design of scalable algorithms for emerging quantum networks. As an initial step, it also opens several directions for future research, including extensions to infinite concept classes, adaptive measurement strategies, and distributed learning under realistic quantum noise and hardware constraints.

Acknowledgments. This work was supported by by the Natural Science Foundation of Shandong Province (ZR2025QC1537) and the China Postdoctoral Science Foundation (CPSF) Postdoctoral Fellowship Program (GZC20251040).

References

1. Ahlswede, R., Winter, A.: Strong converse for identification via quantum channels. IEEE Transactions on Information Theory **48**(3), 569–579 (2002)
2. Barnett, S., Radmore, P.M.: Methods in theoretical quantum optics. OUP, Oxford, United Kingdom (2002)
3. Ben-David, S., Cesa-Bianchi, N., Long, P.M.: Characterizations of learnability for classes of {0,..., n}-valued functions. In: 5th Conference on Computational Learning Theory (COLT 1992). pp. 333–340. ACM, Pittsburgh, PA (1992)
4. Brukhim, N., Carmon, D., Dinur, I., Moran, S., Yehudayoff, A.: A characterization of multiclass learnability. In: 63rd IEEE Symposium on Foundations of Computer Science (FOCS 2022). pp. 943–955. IEEE, Denver, CO (2022)
5. Chehimi, M., Saad, W.: Quantum federated learning with quantum data. In: 47th IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2022). pp. 8617–8621. IEEE, Singapore, Singapore (2022)
6. Daniely, A., Sabato, S., Ben-David, S., Shalev-Shwartz, S.: Multiclass learnability and the ERM principle. The Journal of Machine Learning Research **16**, 2377–2404 (2015)
7. Duchi, J.C., Jordan, M.I., Wainwright, M.J., Zhang, Y.: Optimality guarantees for distributed statistical estimation. cs.it, arXiv (2014)
8. Dyakonov, M.: When will useful quantum computers be constructed? Not in the foreseeable future, this physicist argues. Here's why: The case against: Quantum computing. IEEE Spectrum **56**(3), 24–29 (2019)
9. Grimberg, F., Hartley, M.A., Karimireddy, S.P., Jaggi, M.: Optimal model averaging: Towards personalized collaborative learning. cs.lg, arXiv (2021)
10. Guo, X., Breum, C.R., Borregaard, J., Izumi, S., Larsen, M.V., Gehring, T., Christandl, M., Neergaard-Nielsen, J.S., Andersen, U.L.: Distributed quantum sensing in a continuous-variable entangled network. Nature Physics **16**, 281–284 (2020)
11. Hanna, O.A., Yang, L., Fragouli, C.: Solving multi-arm bandit using a few bits of communication. In: 25th International Conference on Artificial Intelligence and Statistics (AISTATS 2022). pp. 11215–11236. Valencia, Spain (2022)
12. Heidari, M., Padakandla, A., Szpankowski, W.: A theoretical framework for learning from quantum data. In: 54th IEEE International Symposium on Information Theory (ISIT 2021). pp. 1469–1474. IEEE, Melbourne, Australia (2021)

13. Heidari, M., Szpankowski, W.: Learning k-qubit Quantum Operators via Pauli Decomposition. In: 26th International Conference on Artificial Intelligence and Statistics (AISTATS 2023). pp. 490–504. PMLR, Valencia, Spain (2023)
14. Heidari, M., Szpankowski, W.: New bounds on quantum sample complexity of measurement classes. In: IEEE International Symposium on Information Theory (ISIT 2024). pp. 1515–1520. IEEE, Athens, Greece (2024)
15. Holevo, A.S.: Quantum systems, channels, information: a mathematical introduction. De Gruyter Studies in Mathematical Physics, De Gruyter (2019)
16. Javadi-Abhari, A., Treinish, M., Krsulich, K., Wood, C.J., Lishman, J., Gacon, J., Martiel, S., Nation, P.D., Bishop, L.S., Cross, A.W., Johnson, B.R., Gambetta, J.M.: Quantum computing with Qiskit. quant-ph, arXiv (2024), `https://arxiv.org/abs/2405.08810`
17. Kota, S.R., Karthik, P., Tan, V.Y.: Almost cost-free communication in federated best arm identification. In: 37th AAAI Conference on Artificial Intelligence (AAAI 2023). vol. 37, pp. 8378–8385. AAAI, Washington, D.C. (2023)
18. Leonhardt, U.: Measuring the quantum state of light. Cambridge Studies in Modern Optics, Cambridge University Press (1997)
19. Liu, Z.P., Cao, X.Y., Liu, H.W., Sun, X.R., Bao, Y., Lu, Y.S., Yin, H.L., Chen, Z.B.: Practical quantum federated learning and its experimental demonstration. quant-ph 2501.12709, arXiv (2025)
20. Ma, L., Ding, L.: Hybrid quantum edge computing network. In: SPIE OPTICAL ENGINEERING + APPLICATIONS 2022. vol. PC12238. SPIE, San Diego, CA (2022)
21. McMahan, B., Moore, E., Ramage, D., Hampson, S., Arcas, B.A.y.: Communication-efficient learning of deep networks from decentralized data. In: International Conference on Artificial Intelligence and Statistics (AISTATS 2017). pp. 1273–1282. Fort Lauderdale, FL (2017)
22. Natarajan, B.K.: On learning sets and functions. Machine Learning **4**, 67–97 (1989)
23. Nielsen, M.A., Chuang, I.L.: Quantum computation and quantum information. Cambridge University Press (2010)
24. Padakandla, A., Magner, A.: Pac learning of quantum measurement classes: Sample complexity bounds and universal consistency. In: Camps-Valls, G., Ruiz, F., Valera, I. (eds.) 25th International Conference on Artificial Intelligence and Statistics (AISTATS 2022). vol. 151, pp. 11305–11319. PMLR, Valencia, Spain (2022)
25. Rofougaran, R., Yoo, S., Tseng, H.H., Chen, S.Y.C.: Federated quantum machine learning with differential privacy. In: 49th IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2024). pp. 9811–9815. IEEE, Seoul, South Korea (2024)
26. Suthaharan, S.: Quantum aliasing: a negative influence of data scarcity on quantum machine learning. In: SPIE Photonics Europe 2022. vol. 12133, pp. 96–102. SPIE, Strasbourg, France (2022)
27. Tüysüz, C., Clemente, G., Crippa, A., Hartung, T., Kühn, S., Jansen, K.: Classical splitting of parametrized quantum circuits. Quantum Machine Intelligence **5** (2023)
28. Wolpert, D.H., Macready, W.G.: No free lunch theorems for optimization. IEEE Transactions on Evolutionary Computation **1**(1), 67–82 (1997)
29. Wootters, W.K., Zurek, W.H.: A single quantum cannot be cloned. Nature **299**, 802–803 (1982)
30. Zhong, Y., Chang, H.S., Bienfait, A., Dumur, É., Chou, M.H., Conner, C.R., Grebel, J., Povey, R.G., Yan, H., Schuster, D.I., others: Deterministic multi-qubit entanglement in a quantum network. Nature **590**, 571–575 (2021)