

On the learnability of quantum neural networks

Yuxuan Du¹, Min-Hsiu Hsieh², Tongliang Liu¹, Shan You³, and Dacheng Tao¹

¹UBTECH Sydney AI Centre, School of Computer Science, Faculty of Engineering, The University of Sydney, Australia

²Centre for Quantum Software and Information, Faculty of Engineering and Information Technology, University of Technology Sydney, Australia

³SenseTime

Abstract

We consider the learnability of the quantum neural network (QNN) built on the variational hybrid quantum-classical scheme, which remains largely unknown due to the non-convex optimization landscape, the measurement error, and the unavoidable gate errors introduced by noisy intermediate-scale quantum (NISQ) machines. Our contributions in this paper are multi-fold. First, we derive the utility bounds of QNN towards empirical risk minimization, and show that large gate noise, few quantum measurements, and deep circuit depth will lead to the poor utility bounds. This result also applies to the variational quantum circuits with gradient-based classical optimization, and can be of independent interest. We then prove that QNN can be treated as a differentially private (DP) model. Thirdly, we show that if a concept class can be efficiently learned by QNN, then it can also be effectively learned by QNN even with gate noise. This result implies the same learnability of QNN whether it is implemented on noiseless or noisy quantum machines. We last exhibit that the quantum statistical query (QSQ) model can be effectively simulated by noisy QNN. Since the QSQ model can tackle certain tasks with runtime speedup, our result suggests that the modified QNN implemented on NISQ devices will retain the quantum advantage. Numerical simulations support the theoretical results.

1 Introduction

Deep neural network (DNN) has substantially impacted the field of machine learning in the past decade [1]. Most real-world applications, such as object detection [2, 3], question answering [4, 5], social recommendation [6], among many others, could be accomplished by DNN-based learning algorithms with state-of-the-art performance because of the powerful computational hardware and the flexible architecture of DNN. As shown in Fig. 1 (a), DNN adopts a multi-layer scheme. The inputs were processed through the feature embedding layers $\mathcal{F}_x(\cdot)$, followed by the fully-connected layers $\prod_{\ell} W_{\ell}(\cdot)$, where the choice of each layer and the combination rule can be tailor made for various learning tasks. Training DNN is a process to uncover the intrinsic relation between the input and the output of the given dataset. A huge amount of effort has been dedicated to understanding and explaining the *learnability* of DNN from the perspective of the convergence and the generalization [7, 8, 9, 10, 11]; namely, the capabilities and limitations of DNN learning models.

Quantum machine learning is a central application of quantum computing [12]. With the aim of solving real-world problems beyond the reach of classical computers, firm and steady progress has been developed during the past decade [13, 14, 15, 16]. Among these breakthroughs, a quantum extension of DNN, i.e., the quantum neural network (QNN), which is separately proposed in [17, 18, 19, 20],

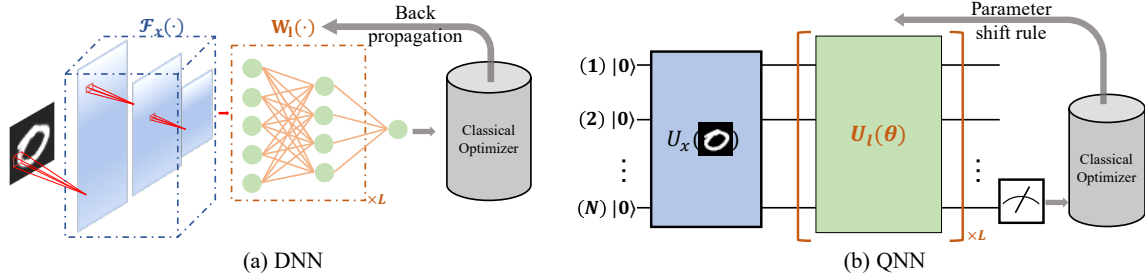


Figure 1: Illustration of DNN and QNN. The left and right panel shows DNN and QNN, respectively. For DNN, the feature embedding layers $\mathcal{F}_x(\cdot)$, which contains a sequence of operations with the arbitrary combination such as convolution and attention, maps the input ‘0’ to the feature space. $W_l(\cdot)$ is the l -th fully-connected layer. For QNN, an encoding quantum circuit U_x maps the classical input ‘0’ to the quantum feature space. $U_l(\theta)$ is the l -th trainable quantum circuit. Classical information for optimization is extracted by quantum measurements.

received great attention due to the huge success of DNN and the superior computational power of quantum machines. As shown in Fig. 1 (b), QNN also adopts the multi-layer architecture, where the inputs were converted into corresponding quantum states by the encoding quantum circuit U_x , followed by trainable quantum circuits $U(\theta) = \prod_{l=1}^L U_l(\theta)$, where θ is the adjustable parameter of quantum gates, and a classical optimizer. There is a close correspondence between DNN and QNN: the feature embedding layers ‘ \mathcal{F}_x ’ of DNN coincide with the encoding quantum circuit U_x of QNN, while the fully-connected layer $W_l(\cdot)$ of DNN coincides with the trainable quantum circuit $U_l(\theta)$ of QNN. Celebrated by the stronger power of quantum circuits to prepare classical distributions [21, 22], QNN could possess a stronger expressive power than its classical counterparts [23] and accelerates a wide range of machine learning problems.

Despite the promising prospects, theoretical results about QNN remain largely unknown. The **difficulties** mainly come from two sides. First, the versatile structures of QNN and their non-convex optimization landscapes, similar to the DNN, heavily challenge the analysis. Second, due to the nature of quantum mechanics, the classical optimizer only receives estimated statistical information with a finite number of measurements, and the error will pile up with the increased number of iterations. Although some studies have overcome partial difficulties from the aspect of vanishing gradients [24, 25], robustness [26, 27], information scrambling [28], memory capacity [29], and no-free lunch theorem [30], the fundamental question, namely, ‘*What is the learnability of QNN*’, is left open.

The importance of exploring QNN’s learnability is further increased in the noisy intermediate-scale quantum (NISQ) era [31, 32], since QNN can be easily built on NISQ machines and its performance is robust against gate noise. Empirical studies have shown that QNN can accomplish various supervised learning tasks, e.g., classification [33, 18, 20], regression [19, 34]. However, no theoretical results can conclude any quantum advantage of these outcomes. To theoretically explain the empirical observations, exploring the learnability of QNN under ERM framework [35] could be very fruitful, because ERM underpins many core results in statistical learning theory and offers learning guarantees for a wide range of supervised learning tasks. Furthermore, it is unclear how gates noise affects the learnability of QNN. This answer substantially affects the feasibility of QNN on NISQ machines to pursue quantum merits.

Problem setup. We follow the convention in statistical learning theory, and examine the learnability of QNN under the framework of empirical risk minimization (ERM) [36] as a first step. In this way, analyzing the learnability of QNN amounts to checking the utility bounds generated by

QNN.

Let $\mathbf{z} = \{\mathbf{z}_j\}_{j=1}^n \in \mathcal{Z}$ be the given dataset with \mathcal{Z} being the sample domain, where the i -th sample $\mathbf{z}_j = (\mathbf{x}_j, y_j)$ includes a feature vector $\mathbf{x}_j \in \mathbb{R}^D$ and a label $y_j \in \mathbb{R}$. ERM aims to find the optimal $\boldsymbol{\theta}^* \in \mathbb{R}^d$ by minimizing the objective function \mathcal{L} within the constraint set $\mathcal{C} \subseteq \mathbb{R}^d$, i.e.,

$$\boldsymbol{\theta}^* = \arg \min_{\boldsymbol{\theta} \in \mathcal{C}} \mathcal{L}(\boldsymbol{\theta}, \mathbf{z}) := \frac{1}{n} \sum_{j=1}^n \ell(y_j, \hat{y}_j) + r(\boldsymbol{\theta}), \quad (1)$$

where \hat{y}_i is the predicted label that is determined by $\boldsymbol{\theta}$ and \mathbf{x}_i , ℓ is the loss function that measures the disparity between true labels $\{y_j\}_{j=1}^n$ and the predicted labels $\{\hat{y}_i\}_{i=1}^n$, and $r(\cdot)$ is a regularizer. To ease the discussion, throughout the paper, we consider the square loss ℓ , and use $r(\boldsymbol{\theta}) = \lambda \|\boldsymbol{\theta}\|_2^2/2$ with $\lambda \geq 0$. Note that our analysis can be easily generalized to other loss functions ℓ that satisfy S -smooth and G -Lipschitz properties as discussed in Sec. 3.

The common optimization rule to tackle ERM is the batch gradient descent method [1]. Depending on the available resources, the sample indices are divided into B disjoint batches $\{\mathcal{B}_i\}_{i=1}^B$ with equal size B_s , namely, $\mathbf{z} = \cup_{j \in \{\mathcal{B}_i\}_{i=1}^B} \mathbf{z}_j$. The optimization rule at the t -th iteration is $\boldsymbol{\theta}^{(t+1)} = \boldsymbol{\theta}^{(t)} - \frac{\eta}{B} \sum_{i=1}^B \nabla \mathcal{L}(\boldsymbol{\theta}^{(t)}, \mathcal{B}_i)$, where η is the learning rate, the gradient $\nabla \mathcal{L}(\cdot)$ is

$$\nabla \mathcal{L}(\boldsymbol{\theta}^{(t)}, \mathcal{B}_i) = \left(\hat{Y}_i^{(t)} - Y_i \right) \frac{\partial \hat{Y}_i^{(t)}}{\partial \boldsymbol{\theta}^{(t)}} + \lambda \boldsymbol{\theta}^{(t)}, \quad (2)$$

$Y_i = \frac{1}{B_s} \sum_{j \in \mathcal{B}_i} y_j$ and $\hat{Y}_i^{(t)} = \frac{1}{B_s} \sum_{j \in \mathcal{B}_i} \hat{y}_j^{(t)}$ are the sum average of the true labels and the predicted labels for the i -th batch \mathcal{B}_i , respectively. When no confusion will occur, we use $\mathcal{L}(\boldsymbol{\theta}^{(t)})$ and $\mathcal{L}_i(\boldsymbol{\theta}^{(t)})$ instead of $\mathcal{L}(\boldsymbol{\theta}^{(t)}, \mathbf{z})$ and $\mathcal{L}(\boldsymbol{\theta}^{(t)}, \mathcal{B}_i)$ in the rest of study.

The training of QNN is similar to those of DNN. In particular, QNN also generated a sum average of the predicted labels, based on $\boldsymbol{\theta}$ and \mathcal{B}_i , after the measurement component in Fig. 1 (b). However, the major difference between the gradient-based optimization of QNN and DNN is as follows. In DNN, the gradient in Eqn. (2) can be easily obtained via backpropagation [1]. However, due to the nature of quantum mechanics, the gradient of a quantum unitary operator (e.g., trainable quantum circuit layer $U_i(\boldsymbol{\theta})$) is, in general, not a legitimate quantum operator anymore [37]. To overcome this shortcoming, the *parameter shift rule* [19, 37] is proposed to estimate the gradients of a quantum unitary operator using K measurements. However, difficulties arise since only approximated $\hat{Y}_i^{(t)}$ and $\partial \hat{Y}_i^{(t)} / \partial \boldsymbol{\theta}^{(t)}$ are available due to a finite number of measurements, and the precision deteriorates when more iterations occur. The detailed steps will be discussed in Sec. 3.

Furthermore, we would like to incorporate the unavoidable gate noise of the trainable quantum circuit $U(\boldsymbol{\theta})$ in our studies. This can be done by considering the worst-case scenario, i.e., modeling the gate noise at each circuit depth to be quantum depolarization noise \mathcal{N}_p [38]. Intuitively, if a quantum state passes through \mathcal{N}_p , with probability $1 - p$, the output remains unchanged; otherwise, all information of the input is lost and the output is the maximally mixed state. Note that the achieved results can be easily extended to a more general noisy channel (See Appendix K for details).

We adopt two standard utility metrics to quantify the performance (learnability) of QNN:

$$R_1(\boldsymbol{\theta}^{(T)}, \mathbf{z}) := \mathbb{E} \left[\|\nabla \mathcal{L}(\boldsymbol{\theta}^{(T)}, \mathbf{z})\|^2 \right], \quad R_2(\boldsymbol{\theta}^{(T)}, \mathbf{z}) := \mathbb{E}[\mathcal{L}(\boldsymbol{\theta}^{(T)}, \mathbf{z})] - \mathcal{L}(\boldsymbol{\theta}^*, \mathbf{z}), \quad (3)$$

where $\boldsymbol{\theta}^{(T)}$ is the output of QNN after T iterations and $\nabla \mathcal{L}(\cdot)$ denotes the gradient of the function $\mathcal{L}(\cdot)$. The metric R_1 evaluates how far QNN is away from the stationary point, $\|\nabla \mathcal{L}(\boldsymbol{\theta}^{(T)}, \mathbf{z})\|^2 = 0$, in expectation [39, 40]. The utility metric R_2 evaluates the expected excess empirical risk [41, 42]. Due

to the hardness to find the global optima in the non-convex landscape, R_2 can only be applied to some special non-convex objective functions, i.e., the objective functions satisfy the Polyak-Lojasiewicz (PL) condition [43, 44]. We will show that, under a mild assumption, the objective function of QNN also meets the PL condition, and R_2 can be employed to analyze its performance.

Contributions. The main contributions of this study are as follows. Our first contribution is deriving QNN’s utility bounds for ERM. As aforementioned, the non-convex optimization landscape, the piled up estimation error due to quantum measurements and the inevitable gate noise, heavily challenge the analysis of QNN’s utility bounds. To the best of our knowledge, this is the first study towards understanding the learnability of QNN with the provable guarantee.

Theorem 1. *QNN outputs $\theta^{(T)} \in \mathbb{R}^d$ after T iterations with utility bounds*

$$R_1 \leq \tilde{O} \left(d, \frac{1}{\frac{1}{BK}K}, \frac{1}{(1-p)^{L_Q}} \right) \text{ and, } R_2 \leq \tilde{O} \left(\frac{1}{K^2B}, d, \frac{1}{(1-p)^{L_Q}} \right),$$

where K is the number of quantum measurements, L_Q is quantum circuit depth, p is the gate noise, and B is the batch size.

Theorem 1 indicates that a larger number of measurements K , a smaller gate noise rate p , a shallower circuit depth L_Q , and a smaller number of trainable parameters d can yield a better utility bounds for both R_1 and R_2 . We remark that the achieved utility bounds R_1 and R_2 are very general, and cover various types of encoding quantum circuits U_x and trainable quantum circuits $U(\theta)$. In particular, our results cover all typical encoding circuits, e.g., amplitude encoding [45, 46, 47], kernel mapping [18, 19, 20], dimension reduction method [48], and basis encoding methods [49, 17], and a diverse architectures of the trainable quantum circuit, as long as it is composed of the parameterized single qubit gates and two qubits gates [50].

Note that the variational hybrid quantum-classical learning models have also been empirically applied to explore fundamental properties of physical systems, e.g., ground energies approximation and thermal averages computation [51, 52]. These problems are generally more sensitive to the global minimum than that of machine learning problems. Therefore the utility bounds in Theorem 1 can serve as a powerful tool to support those results.

A central topic in classical machine learning is exploring whether the noise affects the learnability of a given learning task. A notable example is that the class of parity functions is probably approximately correctly (PAC) learnable; however, learning parity with noise is thought to be computationally hard [53]. Here we lift this essential question from the classical scenario to the quantum scenario: whether there exists any concept class that separates the learnability of the noiseless QNN with noisy QNN, i.e., noiseless QNN uses polynomial samples to learn this concept class, while exponential samples are needed for the NISQ case.

Our second contribution is providing a negative answer towards the above question.

Theorem 2. *If QNN with noiseless gates PAC learns a concept, then there exists a modified QNN with certain types of noisy gates that can also learn this concept using polynomial samples.*

The result of Theorem 2 indicates the same sample complexity between noiseless QNN and noisy QNN to learn a specific concept class. This implies that if QNN achieves certain learning tasks with quantum advantages, then we can implement QNN on NISQ machines with a simple modification to preserve advantages as well.

The key technique used to achieve Theorems 2 is differentially private (DP) learning [54, 55, 56, 57]. The intuition to employ DP is as follows. The behavior of QNN with gate noise resembles learning with noise and DP learning, where a certain type of noise is injected into the learning model. However,

DP learning dispels learning with noise [58], where the former can effectively tackle some tasks that are computationally hard for the latter. Hence, it is beneficial to explore whether QNN with gate noise belongs to a DP learning model instead of learning with noise. The exploration about the DP property of QNN leads to our third contribution.

Lemma 1. *The QNN with gate noise can be treated as a (ϵ, δ) -DP model with $\delta \geq 0$ and*

$$\epsilon = \tilde{O} \left(\sqrt{Td} + Td \left(\frac{(1 - \tilde{p}) + \tilde{p} \frac{\text{Tr}(\Pi)}{D}}{\left(\tilde{p}(1 - \tilde{p})(1 - \frac{\text{Tr}(\Pi)}{D}) \right)^K} \right)^d - Td \right).$$

Together with the fact that non-private and DP algorithms share the same learnability in terms of sample complexity [58], we complete Theorem 2.

Last, we explore the learnability of QNN implemented on NISQ machines. In particular, we aim to find certain tasks that can be achieved by these two learning models with quantum advantages. To reach this goal, we explore whether quantum statistical query learning (QSQ) model can be efficiently simulated by these two learning models, since QSQ can efficiently tackle parity learning, juntas learning, and DNF (disjunctive normal form) learning problems with quantum advantages, whereas these problems are computationally hard for classical SQ models [59]. Our third contribution is exhibiting that QSQ can only be efficiently simulated by QNN with gate noise, and establish a computational separation between the original QNN and modified QNN.

Theorem 3. *A QSQ learning model can be efficiently simulated by QNN with gates noise using polynomial samples.*

1.1 Related work

Previous quantum machine learning literatures that are related to our work can be divided into two groups: quantum learning theory and quantum neural networks. We address that, none of the studies listing below have concerned the learnability of QNN.

For the first group, the studies [60, 61, 62, 63] exhibited that the sample complexity of quantum and classical probably approximately correct learning (PAC) (or agnostic) learning models is equal up to a constant factor under the distribution-independent setting. A recent study [59] generalized the classical statistical query model (SQ) to the quantum statistical query (QSQ) model and compare the learnability among SQ learner, (noisy) quantum PAC learner, QSQ learner, and PPAC learner. However, how to use these results to analyze the learnability of QNN is inexplicit.

For the second group, beyond the hybrid scheme as discussed in this study, there are different schemes and platforms to implement QNN. Specifically, several studies have investigated how to implemented QNN on noiseless quantum machines [64, 65, 66], quantum reservoir [67], and quantum annealers [68]. Since these proposals adopt distinct frameworks, they are incomparable with our results.

2 Preliminary

We unify the notations throughout the whole paper. We denote D as the feature dimension ($\mathbf{x} \in \mathbb{R}^D$), d as the number of training parameters ($\boldsymbol{\theta} \in \mathbb{R}^d$). Define N as the number of qubits and n as the number of training examples. Denote the set $\{1, 2, \dots, m\}$ as $[m]$. A random variable X that follows Bernoulli distribution is denoted as $X \sim \text{Ber}(p)$, i.e., $\Pr(X = 1) = p$ and $\Pr(X = 0) = 1 - p$. With a

slight abuse of notations, we denote ℓ_b as the b -norm, while ℓ (without subscript) is the loss function. We use $O(\cdot)$ (or $\tilde{O}(\cdot)$) to denote the complexity bound (hide poly-logarithmic factors). See Appendix A for details.

Quantum computing. We show basic insights of quantum computing.

Quantum state works in the Hilbert space \mathcal{H} with $\mathcal{H} \cong \mathbb{C}$. Let $|0\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ and $|1\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$ be the standard *basis states* for \mathbb{C}^2 . A quantum bit (**qubit**) lives in a two-dimensional Hilbert space formed by $|0\rangle$ and $|1\rangle$. Multiple qubit basis states follow the tensor products rule, e.g., $|0\rangle \otimes |1\rangle \equiv |1\rangle |0\rangle \in \mathbb{C}^4$ describes a basis state of a 2-qubit system. A *pure state* $|\mathbf{a}\rangle$ with N -qubits follows $|\mathbf{a}\rangle = \sum_{i=1}^d \mathbf{a}_i |i\rangle$ with $d = 2^N$ and $\|\mathbf{a}\|^2 = 1$, where the basis state $|i\rangle \in \{|0\rangle, |1\rangle\}^{\otimes N}$ is also called *computation basis*. $|\mathbf{a}\rangle$ is in *superposition* if $\|\mathbf{a}\|_0 > 1$. The conjugate transpose of $|\mathbf{a}\rangle$ is denoted as $\langle \mathbf{a}|$. We use *density matrix* to describe more general quantum states. Given a mixture of m pure states $\{p_i, |\psi_i\rangle\}_{i=1}^m$ with $p_i \geq 0$ and $\sum_{i=1}^m p_i = 1$, the density matrix ρ is $\rho = \sum_{i=1}^m p_i \rho_i$ with $\rho_i = |\psi_i\rangle \langle \psi_i|$ and $\text{Tr}(\rho) = 1$. There are two main types of quantum operations in quantum computation. The first one is *quantum channel*, which is a completely positive trace-preserving map, e.g., applying a channel \mathcal{N} to a density matrix $\rho \in \mathbb{C}^{d \times d}$ generates the state $\mathcal{N}(\rho) = \sum_a \mathbf{M}_a \rho \mathbf{M}_a^\dagger$ with $\sum_a \mathbf{M}_a \mathbf{M}_a^\dagger = \mathbb{I}_d$. Note that, a quantum gate, which is a unitary matrix, is a special quantum channel. The second one is *quantum measurement*, which extracts classical information from quantum state. An m -outcome measurement, a.k.a. positive-operator-valued measure (**POVM**), is modeled by m positive semidefinite matrices $\{\Pi_b\}_{b=1}^m$ with $\sum_b \Pi_b = \mathbb{I}_d$. Given ρ , the probability to get outcome b is $p_b = \text{Tr}(\Pi_b \rho)$.

Definition 1 (Depolarization channel). *Given a quantum state ρ , the depolarization channel \mathcal{N}_p acts on D -dimensional Hilbert space is defined as $\mathcal{N}_p(\rho) = (1 - p)\rho + p\mathbb{I}/D$.*

Differential privacy. Differential privacy (DP) is a rigorous and standard notion for data privacy, which aims to train an accurate learning model without exposing the precise information in individual training example, e.g., genomic data and medical records for patients [55].

Definition 2 ((ϵ, δ) -DP). *Let $h(\mathbf{z}, \mathbf{z}')$ be the hamming distance. An algorithm \mathcal{M} is (ϵ, δ) -differential private if for any two datasets $\mathbf{z}, \mathbf{z}' \in \mathcal{Z}$ and a distance measure $h(X, X') \leq 1$, and for all measurable sets $\mathcal{O} \subseteq \text{Range}(\mathcal{M})$, the following holds. $\Pr(\mathcal{A}(\mathbf{z}) \in \mathcal{O}) \leq e^\epsilon \Pr(\mathcal{M}(\mathbf{z}') \in \mathcal{O}) + \delta$.*

3 Utility bounds of quantum neural network towards ERM

A well-known consequence in ERM study is that the utility bounds of a given learning model massively depend on what kind of and how much error contained in its gradient [69, 54, 70]. Specifically, when the gradient is perturbed by a sufficiently large amount of noise, the optimization may not converge and the utility bound is poor [71]. Meanwhile, empirical and theoretical evidence has corroborated that, injecting certain types of noise into the gradient does not affect or can even accelerate the convergence [54, 72, 73]. A similar issue also happens to optimize QNN. In particular, when QNN is realized on quantum chips, the presence of sampling error and gate error enables that the classical optimizer only has access to an estimated gradient instead of the analytic gradient. However, theoretical results about how the involved estimation error of gradient affects the optimization remain largely unknown. Moreover, the heuristic study [74] showed that the conclusions based on certain quantum learning models, which are built under the ideal setting that omits the gate error or sample error, may not be applicable to experiments. Therefore, it is crucial to establish the analytical relation between the estimated and analytic gradients, since this relation is not only the precondition to analyze the utility bounds of QNN towards ERM, but can also be used to quantify how the hybrid classical-quantum learning schemes perform on real quantum devices as an independent interest.

We first elaborate the workflow of QNN. As shown in Figure 1 (b), QNN first employs a state preparation unitary $U_{\mathbf{x}}$ to encode classical inputs $\{\mathbf{x}_j | j \in \mathcal{B}_i\}$ into quantum states, followed by the quantum circuit $U(\boldsymbol{\theta})$ with tunable parameter $\boldsymbol{\theta}$ to produce the state $\gamma_{\mathcal{B}_i}$. We refer the interested reader to Appendix B for implementation details of $U_{\mathbf{x}}$ and $U(\boldsymbol{\theta})$. Finally, a two-outcome measurement POVM Π is applied to the state $\gamma_{\mathcal{B}_i}$ and produces the outcome V_i that can be viewed as a binary random variable with the Bernoulli distribution $\text{Ber}(\hat{Y}_i)$, where $\hat{Y}_i := \text{Tr}(\Pi\gamma_{\mathcal{B}_i})$. Denote the obtained statistics, i.e., the sample mean, by $\bar{Y}_i = \frac{1}{K} \sum_{k=1}^K V_k$ after repeating the above procedure K times. The law of quantum mechanics ensures $\bar{Y}_i \rightarrow \hat{Y}_i$ when $K \rightarrow \infty$. However, in reality, only a finite number of measurements is allowed, and this results in the sample error (measurement error).

In addition, the quantum gates in NISQ machines, which are used to implement $U_{\mathbf{x}}$ and $U(\boldsymbol{\theta})$, are prone to having errors [32]. The gate noise can be modelled by applying certain quantum channels to each quantum circuit layer, and we use the depolarization channel \mathcal{N}_p in Definition 1 in the following analysis. Note that our analysis works for more general channels, as discussed in Remark 1. Specifically, with applying \mathcal{N}_p to each layer of quantum circuit, the quantum state before measurement is $\tilde{\gamma}_{\mathcal{B}_i} = \mathcal{N}_p(\gamma_{\mathcal{B}_i})$ instead of $\gamma_{\mathcal{B}_i}$. When the POVM Π is applied to the state $\tilde{\gamma}_{\mathcal{B}_i}$, the obtained outcome V_i follows the distribution $\text{Ber}(\tilde{Y}_i)$ with $\tilde{Y}_i := \text{Tr}(\Pi\tilde{\gamma}_{\mathcal{B}_i})$ instead of $\text{Ber}(\hat{Y}_i)$.

Recall that the updating rule of QNN at the t -th iteration: $\boldsymbol{\theta}^{(t+1)} = \boldsymbol{\theta}^{(t)} - \frac{\eta}{B} \sum_{i=1}^B \nabla \mathcal{L}_i(\boldsymbol{\theta}^{(t)})$, requires the computation of the gradient $\nabla_j \mathcal{L}_i(\boldsymbol{\theta}^{(t)}) = (\hat{Y}_i^{(t)} - Y_i) \partial \hat{Y}_i^{(t)} / \partial \boldsymbol{\theta}_j^{(t)} + \lambda \boldsymbol{\theta}_j^{(t)}$ with $j \in [d]$. In order to obtain the gradient $\nabla_j \mathcal{L}_i(\boldsymbol{\theta}^{(t)})$, the parameter shift rule is developed [19, 37], since the gradient of a quantum unitary operator may not be a legitimate quantum operation and cannot be realized on quantum circuits. Specifically, the parameter shift rule proceeds by separately feeding tunable parameters $\boldsymbol{\theta}^{(t)}$ and $\boldsymbol{\theta}^{(t, \pm j)} := \boldsymbol{\theta}^{(t)} \pm \frac{\pi}{2} \mathbf{e}_j$ to the trainable circuit $U(\boldsymbol{\theta})$, where \mathbf{e}_j is the basis vector with the j -th entry being 1 and zero otherwise. Following the notations used above, we denote $\hat{Y}_i^{(t)}$ and $\hat{Y}_i^{(t, \pm j)}$ as the expectation values of quantum measurements when feeding parameters $\boldsymbol{\theta}^{(t)}$ and $\boldsymbol{\theta}^{(t, \pm j)}$ into trainable quantum circuit $U(\boldsymbol{\theta})$ in the noiseless scenario. The corresponding analytic gradient is

$$\nabla_j \mathcal{L}_i(\boldsymbol{\theta}^{(t)}) = (\hat{Y}_i^{(t)} - Y_i) \frac{\hat{Y}_i^{(t, +j)} - \hat{Y}_i^{(t, -j)}}{2} + \lambda \boldsymbol{\theta}_j^{(t)}.$$

However, in practice, QNN could only generate statistics $\bar{Y}_i^{(t)} = \frac{1}{K} \sum_{k=1}^K V_k^{(t)}$ and $\bar{Y}_i^{(t, \pm j)} = \frac{1}{K} \sum_{k=1}^K V_k^{(t, \pm j)}$, where $V_k^{(t)} \sim \text{Ber}(\hat{Y}_i^{(t)})$ and $V_k^{(t, \pm j)} \sim \text{Ber}(\hat{Y}_i^{(t, \pm j)})$, and $\hat{Y}_i^{(t)}$ and $\hat{Y}_i^{(t, \pm j)}$ refer to the expectation values of quantum measurements when feeding parameters $\boldsymbol{\theta}^{(t)}$ and $\boldsymbol{\theta}^{(t, \pm j)}$ into the noisy trainable quantum circuit $U(\boldsymbol{\theta})$. This leads to the estimated gradient as

$$\nabla_j \bar{\mathcal{L}}_i(\boldsymbol{\theta}^{(t)}) = (\bar{Y}_i^{(t)} - Y_i) \frac{\bar{Y}_i^{(t, +j)} - \bar{Y}_i^{(t, -j)}}{2} + \lambda \boldsymbol{\theta}_j^{(t)}.$$

Our main technical contribution here is showing that the estimated gradient, which is caused by the gates noise and the sampling error, can be related to its optimal gradient, and can be explicitly formulated. An informal result is summarized below (See Theorem D.1 in Appendix D for details).

Theorem 4. *It follows that*

$$\nabla_j \bar{\mathcal{L}}_i(\boldsymbol{\theta}^{(t)}) = (1 - \tilde{p})^2 \nabla_j \mathcal{L}_i(\boldsymbol{\theta}^{(t)}) + C_{j,1}^{(i,t)} + \boldsymbol{\varsigma}_i^{(t,j)},$$

where $\tilde{p} = 1 - (1 - p)^{L_Q}$, L_Q is the circuit depth, the constant $C_{j,1}^{(i,t)}$ only depends on Y_i , $\boldsymbol{\theta}^{(t)}$, and \tilde{p} , and $\boldsymbol{\varsigma}_i^{(t,j)}$ follows the distribution \mathcal{P}_Q that is formed by Y_i , $\boldsymbol{\theta}^{(t)}$, the number of measurements K , and \tilde{p} with zero mean.

The achieved result in Theorem 4 indicates that the estimated gradient $\nabla_j \bar{\mathcal{L}}_i(\boldsymbol{\theta}^{(t)})$ is centralized around the $(1 - \tilde{p})^2 \nabla_j \mathcal{L}_i(\boldsymbol{\theta}^{(t)}) + C_{j,1}^{(i,t)}$ and perturbed by a random variable $\boldsymbol{\varsigma}_i^{(t,j)}$. This enables us to quantitatively measure how far the estimated gradient is away from the analytic gradient, which is the precondition to leverage the optimization theory to analyze the performance of QNN. Moreover, the result of Theorem 4 implies that, compared with the finite measurements, the gate error is more harmful for the QNN’s optimization, which may lead to diverging. In particular, the term $C_{j,1}^{(i,t)}$, which is independent with K , will always exist and induce a biased optimization direction when $\tilde{p} \neq 0$. For the worst case, with $\tilde{p} = 1$, the analytic gradient information is exactly lost. In contrast, K only determines the variance of the distribution \mathcal{P}_Q with zero mean, where classical and quantum literatures [75, 73] have provided the convergence guarantee even if $K = 1$.

Beside the effects of gradient error, the utility bounds also heavily depend on the properties of objective functions. In the following, we show that \mathcal{L} used in QNN satisfies S -smooth, G -Lipschitz, and PL condition. The formal definitions of these concepts and the achieved results are given below.

Definition 3. A function f is S -smooth over a set \mathcal{C} if $\nabla^2 f(\mathbf{u}) \preceq S\mathbb{I}$ with $S > 0$ and $\forall \mathbf{u} \in \mathcal{C}$. A function f is G -Lipschitz over a set \mathcal{C} if for all $\mathbf{u}, \mathbf{w} \in \mathcal{C}$, we have $|f(\mathbf{u}) - f(\mathbf{w})| \leq G\|\mathbf{u} - \mathbf{w}\|_2$. A function f satisfies PL condition if there exists $\mu > 0$ and for every possible $\boldsymbol{\theta} \in \mathcal{C}$, $\|\nabla f(\boldsymbol{\theta})\|^2 \geq 2\mu(f(\boldsymbol{\theta}) - f^*)$, where $f^* = \min_{\boldsymbol{\theta} \in \mathcal{C}} f(\boldsymbol{\theta})$.

Lemma 2. Following the notations in Eqn. (1), $\mathcal{L}(\boldsymbol{\theta})$ is S -smooth with $S = (\frac{3}{2} + \lambda)$ and G -Lipschitz with $G = d(1 + 3\pi\lambda)$. Assuming $\lambda > \frac{1}{\pi}$, \mathcal{L} satisfies PL condition with $\mu = \frac{(-1 + \lambda\pi)^2}{1 + \lambda d(3\pi)^2}$.

The proof of this lemma is provided in Appendix C.

The analysis of the utility bounds R_1 and R_2 of QNN towards ERM, which are summarized in Theorem 1, can be effectively conducted by leveraging Theorem 4 and Lemma 2. Theorem 1 provides the following theoretical guidances to design QNN-based learning algorithms, i.e., a larger amount of measurements K and lager batch size B , smaller depolarizing error p , smaller parameter space d , and shallower quantum circuit L_Q , can yield a better utility bounds R_1 and R_2 .

The full proof of Theorem 1 is provided in Appendix E. The proof strategy of Theorem 1 is as follows. Recall that the utility bound R_1 measures how far the trainable parameter of QNN is away from the stationary point. A well-known result in optimization theory [76] is that the stationary point of a function can be efficiently located by a simple analytic gradient-based algorithm, once the function satisfies the smooth property. Hence, to achieve R_1 , we can utilize the smooth property of \mathcal{L} and the result of Theorem 4, which reformulates the estimated gradient by the analytic gradient, to analyze the stationary convergence of QNN. The key component to achieve R_2 is the PL condition. Recall that the utility bound R_2 evaluates the disparity between the expected and the optimal empirical risk. The study [43] indicates that, if a non-convex function satisfies PL condition, then every stationary point is the global minimum. Alternatively, PL relates the stationary point with the optimal empirical risk, which is determined by the global minimum. Hence, by leveraging the PL condition and the result of R_1 , we can obtain the utility bounds of R_2 .

Remark 1. Theorem 1 can be easily extended to a more general noisy channel \mathcal{E}_{p_1} , i.e., $\mathcal{E}_{p_1}(\rho) = (1 - p_1)\rho + p_2\kappa + p_3\mathbb{I}_D/D$, where $\rho, \kappa \in \mathbb{C}^{D \times D}$ and $p_2 + p_3 = p_1$. See Appendix I for details.

4 The learnability of quantum neural networks with noisy gates

In the analysis of ERM, we exhibit that the gate noise of QNN massively affects its utility bounds. In this section, we aim to understand how noise in QNN affects its learning capabilities in terms

of the sample complexity; namely, whether any concept class that can be probably approximately correctly (PAC) learned by QNN with noiseless gates can also be PAC learned by QNN with noisy gates. If the answer is negative, this concept class is unlikely to be efficiently learned on the NISQ quantum devices. Moreover, it will demonstrate the inequivalent learnability of noiseless QNN and noisy QNN. In the classical literature, the class of parity functions serves as an excellent example to separate the learnability of the PAC learning model with the statistical query (SQ) learning model [53]. Furthermore, there is an even more pressing need to understand what kinds of concept classes can be efficiently learned by QNN with quantum advantages. Towards this question, we explore whether any concept class that is learnable in the quantum statistical query (QSQ) model [59] is also learnable by noiseless and noisy QNN, enlighten by the fact that QSQ model can tackle certain learning tasks that outperform its classical counterpart.

In order to answer the above questions, we attempt to relate noisy QNN with the differentially private (DP) learning model [77], driven by the observation that DP models share a similar behavior with noisy QNN. Specifically, analogous to QNN, DP models involves certain types of noise to achieve the privacy guarantee. If noisy QNN were also a DP model, then we can conclude the same learnability of QNN and noisy QNN, since a concept class that is learnable by a (non-private) algorithm with polynomial sample complexity can also be learned privately using a polynomial number of samples [58]. The Lemma 1 provides an affirmative response, which exhibits that QNN with noisy gates can be treated as a DP learning model (The proof details is given in Appendix F).

The learnability of DP models [58] has been extensively explored in the literature. Two studies [60, 59] separately proved that the sample complexities of classical and quantum (differentially private) PAC learning are equal, up to constant factors. Combining with the fact that PAC = PPAC [58], we can conclude that the sample complexity of PAC, PPAC, quantum PAC, and quantum PPAC learning are equivalent. The conclusion together with Lemma 1 allows us to achieve Theorem 2, i.e., if noiseless QNN PAC learns a concept class, then QNN with gate errors can also learn this concept class using polynomial number of samples. We present the full proof of Theorem 2 in Appendix G.

The result of Theorem 2 indicates that there does not exist a concept class that can be efficiently learned by noiseless QNN, while it is computationally hard for noisy QNN. This result provides a theoretical guarantee to realize QNN on NISQ chips to seek potential quantum advantages.

We further utilize the theoretical results of QSQ model to quantify what kinds of learning problems can be tackled by noisy QNN with quantum advantages. The study [59] shows that QSQ can efficiently tackle parity, juntas, and DNF learning tasks, which are provably hard to learn by the classical statistical query (SQ) models [53]. We proved in Theorem 3 that the QSQ model can be efficiently simulated by QNN, and whose proof is given in Appendix H. Therefore we conclude that these tasks can also be accomplished by QNN with quantum advantage.

All results in this section assume depolarization gate noise; however, they can be extended to a more general model of gate noise. See Appendix I for details.

5 Numerical simulations

We employ the UCI ML hand-written digits datasets [78] to validate the correctness of utility bounds R_1 and R_2 of QNN, as achieved in Section 3 and 4. In the rest of this section, we first introduce the employed dataset and the required preprocessing steps. We then elaborate the employed parameterized quantum circuits that are used in QNN. We last demonstrate our numerical simulation results.

The employed dataset includes in total 1797 hand-written digits images with 10 labels, where each label refers to a digit and each image has 64 attributes. The data preprocessing has three steps.

First, we clean the dataset and only collect images with labels 0 and 1. After cleaning, the total number of images is 360, where the number of examples with label 0 (label 1) is 178 (172). Some collected examples are shown in the left panel of Fig. 2. Alternatively, our simulation focuses on the binary classification task. Second, we utilize a feature dimension reduction technique, i.e., principal component analysis (PCA) [79], to reduce the feature dimension of each data example from 64 to 3. The middle panel of Fig. 2 exhibits the reconstructed hand-written digit images using the reduced data features. Such a step aims to balance the relatively high dimension features of the data example and the limited quantum resources available in present-day. After applying PCA, we denote the employed dataset as $\mathbf{z} = \{(\mathbf{x}_i, y_i)\}_{i=1}^{360}$, where $\mathbf{x}_i \in \mathbb{R}^3$ is the i -th data feature and $y_i \in \{0, 1\}$ is the i -th label. The last step is uniformly and randomly splitting the dataset \mathbf{z} into two groups, i.e., the training dataset \mathbf{z}_t and the test dataset \mathbf{z}_p . The size of the training dataset \mathbf{z}_t and the test dataset \mathbf{z}_p is 280 and 80, respectively.

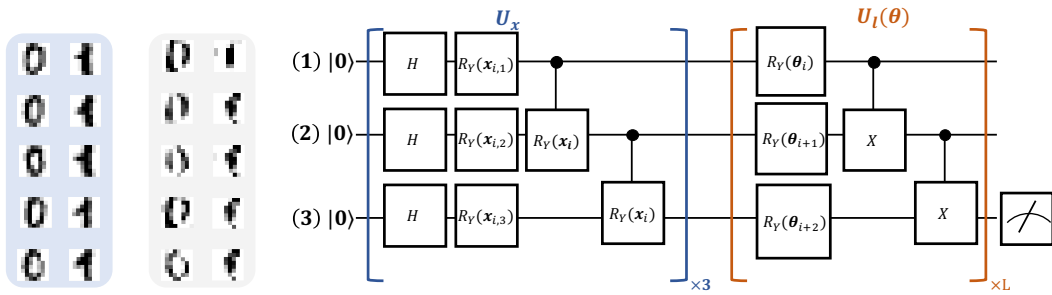


Figure 2: The training examples of hand-written digits and the implementation of quantum circuits. The left panel and middle panel illustrate the original and reconstructed training examples, respectively. The right panel demonstrates the implementation of data encoding circuit and trainable circuit used in QNN. The label ‘x3’ and ‘xL’ means repeating the quantum gates in blue and brown boxes with 3 and L times, respectively.

The construction of parameterized quantum circuit, i.e., the data encoding circuit $U_{\mathbf{x}}$ and the trainable unitary $U(\boldsymbol{\theta})$, follows the proposal [18]. In particular, the data encoding circuit $U_{\mathbf{x}}$ uses the kernel encoding method, and the architecture of trainable unitary $U(\boldsymbol{\theta})$ follows the layer structure. The right panel of Fig. 2 illustrates the implementation of data encoding circuit and trainable circuit used in QNN. Three qubits are employed to build such two circuits. The data encoding circuits $U_{\mathbf{x}}$ is composed of Hadamard gates $H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$, R_Y gates $R_Y(2a) = \begin{pmatrix} \cos(a) & -\sin(a) \\ \sin(a) & \cos(a) \end{pmatrix}$, and controlled- R_Y gates $\text{CRY}(2a) = |0\rangle\langle 0| \otimes \mathbb{I}_2 + |1\rangle\langle 1| \otimes R_Y(2a)$. Specifically, the rotation angle in $R_Y(\mathbf{x})$ is $(\pi - \mathbf{x}_{i,1})(\pi - \mathbf{x}_{i,2})(\pi - \mathbf{x}_{i,3})$. The construction of trainable circuits $U(\boldsymbol{\theta})$ uses R_Y gates and controlled-NOT gates $\text{CX} = |0\rangle\langle 0| \otimes \mathbb{I}_2 + |1\rangle\langle 1| \otimes X$ with $X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$.

We now employ the preprocessed hand-written digits dataset and quantum circuits as described above to study the learnability of QNN under depolarization noise. Specifically, we apply depolarization channel \mathcal{N}_p to every quantum circuit depth, where the depolarization rate is set as $p = 0.0025$. The depth of trainable circuits $U(\boldsymbol{\theta})$ is set as $L = 5$ and $L = 20$, respectively. The corresponding number of trainable parameters is 15 and 60, respectively. We also train QNN without noisy channels \mathcal{N}_p under the setting $L = 5, 20$, which aims to estimate the optimal parameter $\boldsymbol{\theta}^*$ and the minimized objective function \mathcal{L}^* . The number of iterations for all numerical simulations is set as $T = 400$. For QNN, the number of measurements to estimate the expectation value is set as $K = 20$.

The simulation results are shown in Fig. 3. We now elaborate how numerical simulations accord with our theoretical results. Two red arrows indicate the gap between optimal result \mathcal{L}^* and the achieved results $\mathcal{L}(\boldsymbol{\theta}^{(T)})$. With increasing the circuit depth L , the gap becomes large for QNN. Such a phenomenon follows our theoretical result, where a larger d and \tilde{p} lead a poor utility bound R_2 .

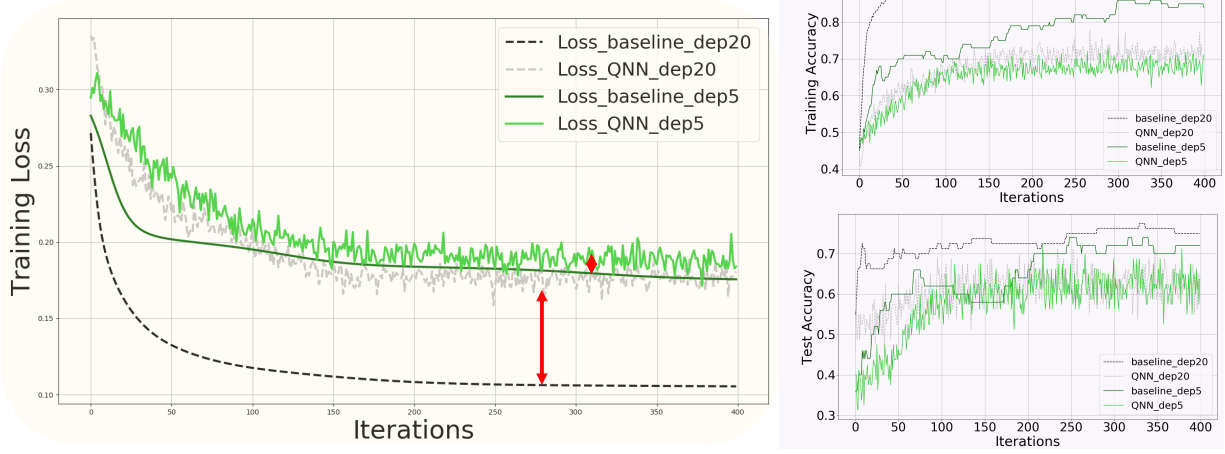


Figure 3: The simulation results of QNN on hand-written digit dataset. The left panel shows the training loss under different hyper-parameters settings. In particular, the label ‘Loss_baseline_dep20’ (‘Loss_baseline_dep5’) refers to the obtained loss under the setting $L = 20$ ($L = 5$), $p = 0$, and $K \rightarrow \infty$, where L , p , and K refer to the circuit depth, depolarization rate, the number of measurements to estimate expectation value used in QNN, respectively. Similarly, the label ‘Loss_QNN_dep20’ (‘Loss_QNN_dep5’) refers to the obtained loss of QNN under the setting $L = 20$ ($L = 5$), $p = 0.0025$, $K = 20$. The upper right and lower right panels separately demonstrate the training accuracy and test accuracy of the quantum classifiers with different hyper-parameters settings.

6 Conclusion

In this study, we explore the learnability of QNN from the aspect of ERM framework and sample complexity. The achieved utility bounds towards ERM indicate that, more measurements, lower noise, and shallower circuit depth contribute to a better performance of QNN. Built on the conclusion that the same learnability between the noiseless QNN and QNN with noisy gates, we obtain the theoretical evidence that supports implementation of QNN on NISQ chips to pursue quantum advantages. Moreover, we demonstrate that QNN with noisy gates can efficiently learn parity, juntas, and DNF with quantum advantages even with gate noise.

Our work also generates plausible new directions for NISQ study that we plan to explore in the future. First, we will use other advanced DP results to analyze various variational hybrid models on NISQ machines with provable guarantees. Second, we aim to tackle private learning tasks with quantum merits because the gate noise of NISQ machines will benefit the design of quantum DP mechanism.

References

- [1] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.
- [2] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017.
- [3] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017.
- [4] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [5] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. Xlnet: Generalized autoregressive pretraining for language understanding. In *Advances in neural information processing systems*, pages 5754–5764, 2019.
- [6] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. Neural collaborative filtering. In *Proceedings of the 26th international conference on world wide web*, pages 173–182, 2017.
- [7] Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via over-parameterization. In *International Conference on Machine Learning*, pages 242–252, 2019.
- [8] Arturs Backurs, Piotr Indyk, and Ludwig Schmidt. On the fine-grained complexity of empirical risk minimization: Kernel methods and neural networks. In *Advances in Neural Information Processing Systems*, pages 4308–4318, 2017.
- [9] Noah Golowich, Alexander Rakhlin, and Ohad Shamir. Size-independent sample complexity of neural networks. *arXiv preprint arXiv:1712.06541*, 2017.
- [10] David Haussler. Decision theoretic generalizations of the pac model for neural net and other learning applications. *Information and computation*, 100(1):78–150, 1992.
- [11] Behnam Neyshabur, Srinadh Bhojanapalli, David McAllester, and Nati Srebro. Exploring generalization in deep learning. In *Advances in Neural Information Processing Systems*, pages 5947–5956, 2017.
- [12] Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. Quantum machine learning. *Nature*, 549(7671):195, 2017.
- [13] Carlo Ciliberto, Mark Herbster, Alessandro Davide Ialongo, Massimiliano Pontil, Andrea Rocchetto, Simone Severini, and Leonard Wossnig. Quantum machine learning: a classical perspective. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 474(2209):20170551, 2018.
- [14] Yuxuan Du, Min-Hsiu Hsieh, and Dacheng Tao. Efficient online quantum generative adversarial learning algorithms with applications. *arXiv preprint arXiv:1904.09602*, 2019.
- [15] Vedran Dunjko and Hans J Briegel. Machine learning & artificial intelligence in the quantum domain: a review of recent progress. *Reports on Progress in Physics*, 81(7):074001, 2018.

- [16] Aram W Harrow and Ashley Montanaro. Quantum computational supremacy. *Nature*, 549(7671):203, 2017.
- [17] Edward Farhi and Hartmut Neven. Classification with quantum neural networks on near term processors. *arXiv preprint arXiv:1802.06002*, 2018.
- [18] Vojtěch Havlíček, Antonio D Córcoles, Kristan Temme, Aram W Harrow, Abhinav Kandala, Jerry M Chow, and Jay M Gambetta. Supervised learning with quantum-enhanced feature spaces. *Nature*, 567(7747):209, 2019.
- [19] Kosuke Mitarai, Makoto Negoro, Masahiro Kitagawa, and Keisuke Fujii. Quantum circuit learning. *arXiv preprint arXiv:1803.00745*, 2018.
- [20] Maria Schuld and Nathan Killoran. Quantum machine learning in feature hilbert spaces. *Physical review letters*, 122(4):040504, 2019.
- [21] Scott Aaronson and Alex Arkhipov. The computational complexity of linear optics. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 333–342. ACM, 2011.
- [22] Michael J Bremner, Richard Jozsa, and Dan J Shepherd. Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 467(2126):459–472, 2011.
- [23] Yuxuan Du, Min-Hsiu Hsieh, Tongliang Liu, and Dacheng Tao. The expressive power of parameterized quantum circuits. *arXiv preprint arXiv:1810.11922*, 2018.
- [24] M Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio, and Patrick J Coles. Cost-function-dependent barren plateaus in shallow quantum neural networks. *arXiv preprint arXiv:2001.00550*, 2020.
- [25] Jarrod R McClean, Sergio Boixo, Vadim N Smelyanskiy, Ryan Babbush, and Hartmut Neven. Barren plateaus in quantum neural network training landscapes. *Nature communications*, 9(1):1–6, 2018.
- [26] Shouvanik Chakrabarti, Huang Yiming, Tongyang Li, Soheil Feizi, and Xiaodi Wu. Quantum wasserstein generative adversarial networks. In *Advances in Neural Information Processing Systems*, pages 6778–6789, 2019.
- [27] Kunal Sharma, Sumeet Khatri, Marco Cerezo, and Patrick J Coles. Noise resilience of variational quantum compiling. *New Journal of Physics*, 22(4):043006, 2020.
- [28] Huitao Shen, Pengfei Zhang, Yi-Zhuang You, and Hui Zhai. Information scrambling in quantum neural networks. *Physical Review Letters*, 124(20):200504, 2020.
- [29] Logan G Wright and Peter L McMahan. The capacity of quantum neural networks. *arXiv preprint arXiv:1908.01364*, 2019.
- [30] Kyle Poland, Kerstin Beer, and Tobias J Osborne. No free lunch for quantum machine learning. *arXiv preprint arXiv:2003.14103*, 2020.
- [31] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando GSL Brandao, David A Buell, et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019.

- [32] John Preskill. Quantum computing in the nisq era and beyond. *Quantum*, 2:79, 2018.
- [33] Carsten Blank, Daniel K Park, June-Koo Kevin Rhee, and Francesco Petruccione. Quantum classifier with tailored quantum kernel. *npj Quantum Information*, 6(1):1–7, 2020.
- [34] Nathan Killoran, Thomas R Bromley, Juan Miguel Arrazola, Maria Schuld, Nicolás Quesada, and Seth Lloyd. Continuous-variable quantum neural networks. *arXiv preprint arXiv:1806.06871*, 2018.
- [35] Vladimir Vapnik. *The nature of statistical learning theory*. Springer science & business media, 2013.
- [36] Vladimir Vapnik. Principles of risk minimization for learning theory. In *Advances in neural information processing systems*, pages 831–838, 1992.
- [37] Maria Schuld, Ville Bergholm, Christian Gogolin, Josh Izaac, and Nathan Killoran. Evaluating analytic gradients on quantum hardware. *Physical Review A*, 99(3):032331, 2019.
- [38] Michael A Nielsen and Isaac L Chuang. *Quantum computation and quantum information*. Cambridge University Press, 2010.
- [39] Vladimir Koltchinskii. *Oracle Inequalities in Empirical Risk Minimization and Sparse Recovery Problems: Ecole d’Eté de Probabilités de Saint-Flour XXXVIII-2008*, volume 2033. Springer Science & Business Media, 2011.
- [40] Jiaqi Zhang, Kai Zheng, Wenlong Mou, and Liwei Wang. Efficient private erm for smooth objectives. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, pages 3922–3928. AAAI Press, 2017.
- [41] Peter L Bartlett, Michael I Jordan, and Jon D McAuliffe. Convexity, classification, and risk bounds. *Journal of the American Statistical Association*, 101(473):138–156, 2006.
- [42] Peter L Bartlett and Shahar Mendelson. Empirical minimization. *Probability theory and related fields*, 135(3):311–334, 2006.
- [43] Yurii Nesterov and Boris T Polyak. Cubic regularization of newton method and its global performance. *Mathematical Programming*, 108(1):177–205, 2006.
- [44] Di Wang, Changyou Chen, and Jinhui Xu. Differentially private empirical risk minimization with non-convex loss functions. In *International Conference on Machine Learning*, pages 6526–6535, 2019.
- [45] Martin Plesch and Āaslav Brukner. Quantum-state preparation with universal gate decompositions. *Physical Review A*, 83(3):032302, 2011.
- [46] Maria Schuld, Mark Fingerhuth, and Francesco Petruccione. Implementing a distance-based classifier with a quantum interference circuit. *arXiv preprint arXiv:1703.10793*, 2017.
- [47] Maria Schuld, Alex Bocharov, Krysta M Svore, and Nathan Wiebe. Circuit-centric quantum classifiers. *Physical Review A*, 101(3):032308, 2020.
- [48] CM Wilson, JS Otterbach, Nikolas Tezak, RS Smith, GE Crooks, and MP da Silva. Quantum kitchen sinks: An algorithm for machine learning on near-term quantum computers. *arXiv preprint arXiv:1806.08321*, 2018.

- [49] Ashish Kapoor, Nathan Wiebe, and Krysta Svore. Quantum perceptron models. In *Advances in Neural Information Processing Systems*, pages 3999–4007, 2016.
- [50] Marcello Benedetti, Erika Lloyd, Stefan Sack, and Mattia Fiorentini. Parameterized quantum circuits as machine learning models. *Quantum Science and Technology*, 4(4):043001, 2019.
- [51] Mario Motta, Chong Sun, Adrian TK Tan, Matthew J O’Rourke, Erika Ye, Austin J Minnich, Fernando GSL Brandão, and Garnet Kin-Lic Chan. Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution. *Nature Physics*, 16(2):205–210, 2020.
- [52] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J Love, Alán Aspuru-Guzik, and Jeremy L O’Brien. A variational eigenvalue solver on a photonic quantum processor. *Nature communications*, 5:4213, 2014.
- [53] Avrim Blum, Adam Kalai, and Hal Wasserman. Noise-tolerant learning, the parity problem, and the statistical query model. *Journal of the ACM (JACM)*, 50(4):506–519, 2003.
- [54] Kamalika Chaudhuri, Claire Monteleoni, and Anand D Sarwate. Differentially private empirical risk minimization. *Journal of Machine Learning Research*, 12(Mar):1069–1109, 2011.
- [55] Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.
- [56] Shenggang Ying, Mingsheng Ying, and Yuan Feng. Quantum privacy-preserving perceptron. *arXiv preprint arXiv:1707.09893*, 2017.
- [57] Li Zhou and Mingsheng Ying. Differential privacy in quantum computation. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pages 249–262. IEEE, 2017.
- [58] Shiva Prasad Kasiviswanathan, Homin K Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. What can we learn privately? *SIAM Journal on Computing*, 40(3):793–826, 2011.
- [59] Srinivasan Arunachalam, Alex B Grilo, and Henry Yuen. Quantum statistical query learning. *arXiv preprint arXiv:2002.08240*, 2020.
- [60] Srinivasan Arunachalam and Ronald de Wolf. Guest column: A survey of quantum learning theory. *ACM SIGACT News*, 48(2):41–67, 2017.
- [61] Alp Atici and Rocco A Servedio. Improved bounds on quantum learning algorithms. *Quantum Information Processing*, 4(5):355–386, 2005.
- [62] Ethan Bernstein and Umesh Vazirani. Quantum complexity theory. *SIAM Journal on computing*, 26(5):1411–1473, 1997.
- [63] Rocco A Servedio and Steven J Gortler. Equivalences and separations between quantum and classical learnability. *SIAM Journal on Computing*, 33(5):1067–1092, 2004.
- [64] Kerstin Beer, Dmytro Bondarenko, Terry Farrelly, Tobias J Osborne, Robert Salzmann, Daniel Scheiermann, and Ramona Wolf. Training deep quantum neural networks. *Nature Communications*, 11(1):1–6, 2020.
- [65] Iordanis Kerenidis, Jonas Landman, and Anupam Prakash. Quantum algorithms for deep convolutional neural networks. In *International Conference on Learning Representations*, 2020.

- [66] Nathan Wiebe, Ashish Kapoor, and Krysta M Svore. Quantum deep learning. *arXiv preprint arXiv:1412.3489*, 2014.
- [67] Sanjib Ghosh, Andrzej Opala, Michał Matuszewski, Tomasz Paterek, and Timothy CH Liew. Quantum reservoir processing. *npj Quantum Information*, 5(1):1–6, 2019.
- [68] Bartłomiej Gardas, Marek M. Rams, and Jacek Dziarmaga. Quantum neural networks to simulate many-body quantum systems. *Phys. Rev. B*, 98:184304, Nov 2018.
- [69] Raef Bassily, Adam Smith, and Abhradeep Thakurta. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 464–473. IEEE, 2014.
- [70] Song Mei, Yu Bai, and Andrea Montanari. The landscape of empirical risk for nonconvex losses. *Ann. Statist.*, 46(6A):2747–2774, 12 2018.
- [71] Stephen Boyd and Lieven Vandenbergh. *Convex optimization*. Cambridge university press, 2004.
- [72] Yuanzhi Li and Yang Yuan. Convergence analysis of two-layer neural networks with relu activation. In *Advances in neural information processing systems*, pages 597–607, 2017.
- [73] Mo Zhou, Tianyi Liu, Yan Li, Dachao Lin, Enlu Zhou, and Tuo Zhao. Towards understanding the importance of noise in training neural networks. *arXiv preprint arXiv:1909.03172*, 2019.
- [74] Kevin J. Sung, Matthew P. Harrigan, Nicholas C. Rubin, Zhang Jiang, Ryan Babbush, and Jarrod R. McClean. An exploration of practical optimizers for variational quantum algorithms on superconducting qubit processors, 2020.
- [75] Ryan Sweke, Frederik Wilde, Johannes Meyer, Maria Schuld, Paul K Fährmann, Barthélémy Meynard-Piganeau, and Jens Eisert. Stochastic gradient descent for hybrid quantum-classical optimization. *arXiv preprint arXiv:1910.01155*, 2019.
- [76] Chi Jin, Rong Ge, Praneeth Netrapalli, Sham M Kakade, and Michael I Jordan. How to escape saddle points efficiently. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1724–1732. JMLR. org, 2017.
- [77] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pages 265–284. Springer, 2006.
- [78] Dheeru Dua and Casey Graff. UCI machine learning repository, 2017.
- [79] Svante Wold, Kim Esbensen, and Paul Geladi. Principal component analysis. *Chemometrics and intelligent laboratory systems*, 2(1-3):37–52, 1987.
- [80] Yuxuan Du, Min-Hsiu Hsieh, Tongliang Liu, and Dacheng Tao. Implementable quantum classifier for nonlinear data. *arXiv preprint arXiv:1809.06056*, 2018.
- [81] Abhinav Kandala, Antonio Mezzacapo, Kristan Temme, Maika Takita, Markus Brink, Jerry M Chow, and Jay M Gambetta. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature*, 549(7671):242–246, 2017.

- [82] William Huggins, Piyush Patil, Bradley Mitchell, K Birgitta Whaley, and E Miles Stoudenmire. Towards quantum machine learning with tensor networks. *Quantum Science and technology*, 4(2):024001, 2019.
- [83] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm. *arXiv preprint arXiv:1411.4028*, 2014.
- [84] Yuxuan Du, Min-Hsiu Hsieh, Tongliang Liu, Dacheng Tao, and Nana Liu. Quantum noise protects quantum classifiers against adversaries. *arXiv preprint arXiv:2003.09416*, 2020.

The organization of the appendix is as follows. In Appendix A, we unify the notations used in the whole appendix. In Appendix B, we elaborate the implementation details of the quantum encoding circuit $U_{\mathbf{x}}$ and the trainable quantum circuit $U(\boldsymbol{\theta})$ used in QNN. In Appendix C, we present the proof of Lemma 2, which quantifies the properties of the objective function with respect to the optimization theory. Then, in Appendix D, we exhibit the proof of Theorem 4, as the precondition to achieve utility bounds of QNN. In Appendix E, we exhibit the proofs details of Theorem 1 that achieves the utility bounds of QNN towards ERM. The following four sections explain the learnability of QNN from the perspective of sample complexity. Specifically, in Appendix F, we provide the proof details of Lemma 1. Next, in Appendix G and H, we separately prove Theorem 2 and Theorem 3. Eventually, in Appendix I, we generalize all achieved results to a more general quantum channel.

A The summary of notations

Here we unify the notations used in the appendices. A random variable X that follows Delta distribution is denoted as $X \sim \text{Del}(x_0)$, i.e., $\Pr(X = x_0) = 1$ and $\Pr(X \neq x_0) = 0$. A random variable X that follows uniform distribution is denoted as $X \sim U(a, b)$, where $P(X = x_0) = 1/(b - a)$ with $a \leq x_0 \leq b$. We denote the ℓ_p norm of \mathbf{v} as $\|\mathbf{v}\|_p$. In particular, $\|\mathbf{v}\|$ refers to the ℓ_2 norm.

B Implementation details of encoding circuit and trainable circuit of QNN

The selection of encoding circuits $U_{\mathbf{x}}$ and trainable circuit $U(\boldsymbol{\theta})$ is flexible in QNN. We now separately explain the implementation details of these two circuits supported by QNN.

Encoding circuit $U_{\mathbf{x}}$. The typical encoding circuits can be divided into four categories. A common feature of these encoding methods is that their implementation only costs low circuit depth, driven by the restricted quantum resources. The first category is the direct amplitude encoding [80, 45, 46, 47]. Specifically, the encoder circuit satisfies $U_{\mathbf{x}} : \mathcal{B}_i \rightarrow \frac{1}{\sqrt{B_s}} \sum_{b=1}^{B_s} \sum_{j=1}^D \hat{\mathbf{x}}_{b,j}^{(i)} |b\rangle |j\rangle$ with $\hat{\mathbf{x}}_{b,j}^{(i)} = \mathbf{x}_{b,j}^{(i)} / \|\mathbf{x}_{b,j}^{(i)}\|$. This method requires a low feature dimension D , since the quantum gates complexity to build $U_{\mathbf{x}}$ is $O(D)$. The second category is the kernel mapping [18, 19, 20], where \mathcal{B}_i is encoded into a set of single-qubit gates with a specified arrangements, e.g., $U_{\mathbf{x}}(\mathcal{B}_i) = \sum_{b=1}^{B_s} (|b\rangle \langle b| \otimes_{j=1}^D \text{R}_Y(\mathbf{x}_{b,j}^{(i)}))$. The third category is the dimension reduction method proposed by [48]. Specifically, instead of encoding \mathcal{B}_i , the amplitude or kernel encoder circuits $U_{\mathbf{x}}$ is exploited to encode a projected features $g(\mathcal{B}_i) \in \mathbb{R}^{B_s \times D'}$, where $g(\cdot)$ is a predefined function and $D' \ll D$. The fourth category is the basis encoding [60, 59, 17], which is broadly used in quantum learning theory. Specifically, the encoding circuit $U_{\mathbf{x}}$ is employed to prepare a quantum example $|\psi\rangle = \sum_{\mathbf{x} \in \{0,1\}^N} \sqrt{\mathcal{D}(\mathbf{x})} |\mathbf{x}, c(\mathbf{x})\rangle$, where $\mathcal{D}(\mathbf{x})$ is the data distribution over \mathbf{x} , $c(\mathbf{x})$ corresponds to the label of the bit-string \mathbf{x} [60, 61]. In most cases, the distribution $\mathcal{D}(\mathbf{x})$ is uniform. Hence, the state $|\psi\rangle$ can be efficiently prepared by setting $B = 1$, and applying Hadamard gates and control-not gates [38] to the initial state $|0\rangle^{\otimes N+1}$.

Trainable quantum circuits $U(\boldsymbol{\theta})$. The trainable quantum circuits, a.k.a, parameterized quantum circuits [50, 23], used in QNN can be written as a product of layers of unitaries in the form $U(\boldsymbol{\theta}) = \prod_{l=1}^L U_l(\boldsymbol{\theta}_l)$, where $U_l(\boldsymbol{\theta}_l)$ is composed of parameterized single-qubit gates and fixed two-qubits gates. Each trainable layer can be decomposed into $U_l(\boldsymbol{\theta}_l) = (\bigotimes_{k=1}^N U_{l,k}(\boldsymbol{\theta}_l)) U_{eng}$, where $U_{l,k}(\boldsymbol{\theta}_l)$ represents the composition of trainable single-qubit gates and U_{eng} refers to entanglement layer that contains two-qubits gates. Depending on the detailed architecture, the implementation of $U_l(\boldsymbol{\theta}_l)$ can be categorized into three classes. The first class is the hardware-efficient circuit

architecture, where the selection of $U_k(\boldsymbol{\theta}_l)$ and U_{eng} is according to the given NISQ machine that has the specific sparse qubit-to-qubit connectivity and a specified set of quantum gates [24, 81, 25]. The second class is the tensor network inspired architecture. In particular, the layout of quantum gates is following different tensor networks, e.g., the matrix product state, the tree tensor network, and the multi-scale entanglement renormalization ansatz (MERA) [82]. The third class is the Hamiltonian based architecture, where the entanglement layer U_{eng} refers to a specific Hamiltonian, e.g., the study [19] employs $U_{eng} = e^{-iHT}$ with $H = \sum_{j=1}^N a_j X_j + \sum_{j=1}^N \sum_{k=1}^{j-1} J_{jk} Z_j Z_k$. Notably, almost all quantum approximate optimization algorithms follow the Hamiltonian based architecture [83].

C Proof of Lemma 2

The Lemma 2 indicates that the objective function $\mathcal{L}(\boldsymbol{\theta})$ used in QNN satisfies S -smooth and G -Lipschitz properties. Moreover, it also satisfies the Polyak-Lojasiewicz (PL) condition under the assumption with $\lambda \geq 1/\pi$. To ease the discussion, we first formulate the explicit form of $\mathcal{L}(\boldsymbol{\theta})$. Without loss of generality, we set $B = n$, where each batch \mathcal{B}_i only contains the i -th input \mathbf{x}_i . Denote the prepared quantum states as $\{\rho_{\mathcal{B}_i}\}_{i=1}^n$ i.e., $\rho_{\mathcal{B}_i} = |\phi_{\mathcal{B}_i}\rangle\langle\phi_{\mathcal{B}_i}|$ and $|\phi_{\mathcal{B}_i}\rangle \stackrel{U_{\mathbf{x}}}{\leftarrow} \{\mathbf{x}_i\}$ refers to the quantum example corresponding to the classical input batch \mathcal{B}_i (or equivalently, \mathbf{x}_i). The explicit form of the objective function is

$$\mathcal{L}(\boldsymbol{\theta}, \mathbf{z}) = \frac{1}{n} \sum_{i=1}^n (\hat{y}_i - y_i)^2 + \frac{\lambda}{2} \|\boldsymbol{\theta}\|_2^2, \quad (4)$$

where $\hat{y}_i = \text{Tr}(\Pi U(\boldsymbol{\theta}) \rho_{\mathcal{B}_i} U(\boldsymbol{\theta})^\dagger)$ refers to the prediction of QNN given the i -th input \mathbf{x}_i , $U(\boldsymbol{\theta})$ is the trainable circuit, Π is the employed two-outcome POVM, and y_i is the true label of the i -th input. Moreover, since the tunable parameters $\boldsymbol{\theta}$ in QNN refer to the rotation angles, we set its range as $\boldsymbol{\theta} \in [\pi, 3\pi]^d$.

Proof of Lemma 2. We employ the three lemmas presented below to prove Lemma 2, whose proofs are given in the following subsections.

Lemma 3. *The objective function \mathcal{L} is S -smooth with $S = (3/2 + \lambda)$.*

Lemma 4. *The objective function \mathcal{L} is G -Lipschitz with $G = d(1 + 3\pi\lambda)$.*

Lemma 5. *Assume $\lambda > \frac{1}{\pi}$. The objective function \mathcal{L} satisfies PL condition with $\mu = \frac{(-1+\lambda\pi)^2}{1+\lambda d(3\pi)^2}$.*

In conjunction with the results of Lemma 3, 4, and 5, the proof of Lemma 2 is completed. \square

C.1 Proof of S -smooth, Lemma 3

Proof of Lemma 3. Recall the function $\mathcal{L}(\boldsymbol{\theta})$ is S -smooth if

$$\nabla^2 \mathcal{L}(\boldsymbol{\theta}) \preceq S \mathbb{I}, \quad (5)$$

with $S > 0$. In other words, to obtain S , we need to obtain the upper bound of the second derivative of $\mathcal{L}(\boldsymbol{\theta})$, i.e., $S \geq \|\nabla^2 \mathcal{L}(\boldsymbol{\theta})\|_\infty$.

Following the notation used in Eqn. (4), the gradient for the parameter $\boldsymbol{\theta}_j$ is

$$\frac{\partial \mathcal{L}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}_j}$$

$$\begin{aligned}
&= \frac{2}{n} \sum_{i=1}^n (\hat{y}_i - y_i) \frac{\partial \hat{y}_i}{\partial \theta_j} + \frac{\lambda}{2} \frac{\partial \|\boldsymbol{\theta}\|_2^2}{\partial \theta_j} \\
&= \frac{2}{n} \sum_{i=1}^n (\hat{y}_i - y_i) \frac{\hat{y}_i^{(+j)} - \hat{y}_i^{(-j)}}{2} + \lambda \theta_j \\
&\leq 1 + 3\lambda\pi,
\end{aligned} \tag{6}$$

where $\hat{y}_i^{(\pm j)} = \text{Tr}(\Pi U(\boldsymbol{\theta} \pm \frac{\pi}{2} \mathbf{e}_j) \rho_{\mathcal{B}_i} U(\boldsymbol{\theta} \pm \frac{\pi}{2} \mathbf{e}_j)^\dagger)$, the second equality employs the conclusion of the parameter shift rule with $\frac{\partial \hat{y}_i}{\partial \theta_j} = \frac{\hat{y}_i^{(+j)} - \hat{y}_i^{(-j)}}{2}$ [19, 37], and the last inequality uses the facts $\pi \leq \theta_j \leq 3\pi$, $(\hat{y}_i - y_i) \leq 1$, and $\hat{y}_i^{(+j)} - \hat{y}_i^{(-j)} \leq 1$, since $\hat{y}_i, y_i, \hat{y}_i^{(\pm j)} \in [0, 1]$.

The upper bound of the derivative $\frac{\partial^2 \mathcal{L}(\boldsymbol{\theta})}{\partial \theta_j \partial \theta_k}$ can be derived using the results of Eqn. (6). In particular,

$$\begin{aligned}
\frac{\partial^2 \mathcal{L}(\boldsymbol{\theta})}{\partial \theta_j \partial \theta_k} &= \frac{\partial \left(\frac{\partial \mathcal{L}(\boldsymbol{\theta})}{\partial \theta_j} \right)}{\partial \theta_k} = \frac{1}{n} \sum_{i=1}^n \frac{\partial \left((\hat{y}_i - y_i) \left(\hat{y}_i^{(+j)} - \hat{y}_i^{(-j)} \right) + \lambda \theta_j \right)}{\partial \theta_k} \\
&= \frac{1}{n} \sum_{i=1}^n \left[\frac{\partial \hat{y}_i}{\partial \theta_k} \left(\hat{y}_i^{(+j)} - \hat{y}_i^{(-j)} \right) + (\hat{y}_i - y_i) \frac{\partial \left(\hat{y}_i^{(+j)} - \hat{y}_i^{(-j)} \right)}{\partial \theta_k} + \lambda \right] \\
&\leq \frac{3}{2} + \lambda,
\end{aligned} \tag{7}$$

where the first equality comes from the last equality of Eqn. (6), and the last inequality employs $(\hat{y}_i - y_i) \leq 1$, $\hat{y}_i^{(+j)} - \hat{y}_i^{(-j)} \leq 1$, and

$$\frac{\partial \hat{y}_i}{\partial \theta_k}, \frac{\partial \hat{y}_i^{(+j)}}{\partial \theta_k}, \frac{\partial \hat{y}_i^{(-j)}}{\partial \theta_k} \in [-1/2, 1/2],$$

supported by the parameter shift rule and $\hat{y}_i, \hat{y}_i^{(\pm j)} \in [0, 1]$.

The result of Eqn. (7) implies that $\|\nabla^2 \mathcal{L}\|_\infty \leq \frac{3}{2} + \lambda$. In conjunction with Eqn. (5), the objective function is S -smooth with $S = \frac{3}{2} + \lambda$. \square

C.2 Proof of G -Lipschitz, Lemma 4

Proof of Lemma 4. Recall a function $f(\mathbf{x})$ is G -Lipschitz if it satisfies

$$|f(\mathbf{b}) - f(\mathbf{a})| \leq G \|\mathbf{b} - \mathbf{a}\|. \tag{8}$$

Moreover, the mean value theorem gives that, if $f: \mathbb{R}^d \rightarrow \mathbb{R}$ is differentiable and $[\mathbf{a}, \mathbf{b}] \subseteq \mathbb{R}^d$, then $\exists \mathbf{c} \in (\mathbf{a}, \mathbf{b})$ such that

$$f(\mathbf{b}) - f(\mathbf{a}) = \langle \nabla f(\mathbf{c}), \mathbf{b} - \mathbf{a} \rangle. \tag{9}$$

Combining Eqn. (8) and (9), the G -Lipschitz condition in Eqn. (8) is equivalent to

$$|\langle \nabla f(\mathbf{c}), \mathbf{b} - \mathbf{a} \rangle| \leq G \|\mathbf{b} - \mathbf{a}\|. \tag{10}$$

We now replace f , \mathbf{b} , and \mathbf{a} used in Eqn. (10) with \mathcal{L} , $\boldsymbol{\theta}^{(1)}$, and $\boldsymbol{\theta}^{(2)}$ to prove that the objective function \mathcal{L} is G -Lipschitz. Specifically, we need to find a real value G that satisfies

$$\left| \langle \nabla \mathcal{L}(\boldsymbol{\theta}), \boldsymbol{\theta}^{(1)} - \boldsymbol{\theta}^{(2)} \rangle \right| \leq G \|\boldsymbol{\theta}^{(1)} - \boldsymbol{\theta}^{(2)}\|, \tag{11}$$

where $\boldsymbol{\theta} \in (\boldsymbol{\theta}^{(2)}, \boldsymbol{\theta}^{(1)})$.

The upper bound of the term $\langle \nabla \mathcal{L}(\boldsymbol{\theta}), \boldsymbol{\theta}^{(1)} - \boldsymbol{\theta}^{(2)} \rangle$ is

$$\langle \nabla \mathcal{L}(\boldsymbol{\theta}), \boldsymbol{\theta}^{(1)} - \boldsymbol{\theta}^{(2)} \rangle \leq \|\nabla \mathcal{L}(\boldsymbol{\theta})\| \|\boldsymbol{\theta}^{(1)} - \boldsymbol{\theta}^{(2)}\| \leq d \|\nabla \mathcal{L}(\boldsymbol{\theta})\|_{\infty} \|\boldsymbol{\theta}^{(1)} - \boldsymbol{\theta}^{(2)}\|. \quad (12)$$

In conjunction with Eqn. (11) and (12), G -Lipschitz of \mathcal{L} requests

$$d \|\nabla \mathcal{L}(\boldsymbol{\theta})\|_{\infty} \leq G. \quad (13)$$

By leveraging the result of Eqn. (6) with $\nabla_j \mathcal{L}(\boldsymbol{\theta}) \leq 1 + 3\lambda\pi$, we obtain the upper bound of the left side in Eqn. (13) is

$$d \|\nabla \mathcal{L}(\boldsymbol{\theta})\|_{\infty} \leq d(1 + 3\pi\lambda). \quad (14)$$

This leads to the objective function \mathcal{L} of QNN satisfying G -Lipschitz with $G = d(1 + 3\pi\lambda)$. \square

C.3 Proof of PL condition, Lemma 5

Proof of Lemma 5. Recall the definition of Polyak-Lojasiewicz as formulated in Definition 3, it requires that the objective function \mathcal{L} satisfies

$$\|\nabla \mathcal{L}(\boldsymbol{\theta})\|^2 \geq 2\mu(\mathcal{L}(\boldsymbol{\theta}) - \mathcal{L}^*), \quad (15)$$

where $\mathcal{L}^* = \min_{\boldsymbol{\theta} \in \mathcal{C}} \mathcal{L}(\boldsymbol{\theta})$.

We first derive a lower bound of $\|\nabla \mathcal{L}(\boldsymbol{\theta})\|^2$. In particular, we have

$$\|\nabla \mathcal{L}(\boldsymbol{\theta})\|^2 = \sum_{j=1}^d (\nabla_j \mathcal{L}(\boldsymbol{\theta}_j))^2 \geq \max_j (\nabla_j \mathcal{L}(\boldsymbol{\theta}))^2. \quad (16)$$

The lower bound of $\max_j (\nabla_j \mathcal{L}(\boldsymbol{\theta}))^2$ as shown in Eqn. (16) follows

$$\max_j (\nabla_j \mathcal{L}(\boldsymbol{\theta}))^2 \geq (-1 + \lambda\pi)^2, \quad (17)$$

where the last inequality is achieved by exploiting the last second line of Eqn. (6), and the facts $\boldsymbol{\theta}_j \in [\pi, 3\pi]$ and $\hat{y}_i, y_i, \hat{y}_i^{(\pm j)} \in [0, 1]$, i.e.,

$$\nabla_j \mathcal{L}(\boldsymbol{\theta}) = \frac{2}{n} \sum_{i=1}^n (\hat{y}_i - y_i) \frac{\hat{y}_i^{(+j)} - \hat{y}_i^{(-j)}}{2} + \lambda \boldsymbol{\theta}_j \geq -1 + \lambda\pi.$$

Combining the assumption $\lambda \geq 1/\pi$ and the above results, the lower bound of Eqn. (16) satisfies

$$\|\nabla \mathcal{L}(\boldsymbol{\theta})\|^2 \geq (-1 + \lambda\pi)^2 > 0.$$

We then derive the upper bound of the term $(\mathcal{L}(\boldsymbol{\theta}) - \mathcal{L}^*)$ in Eqn. (15). In particular, we have

$$\mathcal{L}(\boldsymbol{\theta}) - \mathcal{L}^* \leq \mathcal{L}(\boldsymbol{\theta}) + 0 \leq 1 + \lambda d (3\pi)^2, \quad (18)$$

where the first inequality comes from the definitions of \mathcal{L}^* , i.e.,

$$-\mathcal{L}^* = -\frac{1}{n} \sum_{i=1}^n (\hat{y}_i^* - y_i)^2 - \frac{\lambda}{2} \|\boldsymbol{\theta}\|^2 \leq 0,$$

with $\hat{y}_i^* = \text{Tr}(\Pi U(\boldsymbol{\theta}^*) \rho_i U(\boldsymbol{\theta}^*)^\dagger)$, and the second inequality employs the definition of $\mathcal{L}(\boldsymbol{\theta})$ with

$$\mathcal{L}(\boldsymbol{\theta}) = \frac{1}{n} \sum_{i=1}^n (\hat{y}_i - y_i)^2 + \frac{\lambda}{2} \|\boldsymbol{\theta}\|^2 \leq 1 + \frac{\lambda}{2} \|\boldsymbol{\theta}\|^2 ,$$

and $\frac{\lambda}{2} \|\boldsymbol{\theta}\|^2 \leq \frac{\lambda}{2} d \|\boldsymbol{\theta}\|_\infty^2 = (3\pi)^2 \lambda d / 2$.

By combining Eqn. (17) and (18) with Eqn. (15), we obtain the following relation

$$\|\nabla \mathcal{L}(\boldsymbol{\theta})\|^2 \geq (-1 + \lambda\pi)^2 \geq 2\mu(1 + \lambda d(3\pi)^2) \geq 2\mu(\mathcal{L}(\boldsymbol{\theta}) - \mathcal{L}^*) . \quad (19)$$

The above relation indicates that the objection function $\mathcal{L}(\boldsymbol{\theta})$ satisfies PL condition with

$$\mu = \frac{(-1 + \lambda\pi)^2}{1 + \lambda d(3\pi)^2} .$$

□

D Proof of Theorem 4

Theorem 4 establishes the relation between the analytic gradient $\nabla_j \mathcal{L}_i(\boldsymbol{\theta}^{(t)})$ and the estimated gradient $\nabla_j \bar{\mathcal{L}}_i(\boldsymbol{\theta}^{(t)})$ of QNN. Its formal description is as follows.

Theorem 5 (The formal description of Theorem 4). *Denote $\tilde{p} = 1 - (1 - p)^{L_Q}$ with L_Q being the quantum circuit depth. At the t -th iteration, we define five constants with*

$$C_{j,a}^{(i,t)} = \begin{cases} (1 - \tilde{p})\tilde{p}(1/2 - Y_i)(\hat{Y}_i^{(t,+j)} - \hat{Y}_i^{(t,-j)}) - (2\tilde{p} - \tilde{p}^2)\lambda\boldsymbol{\theta}_j^{(t)} , & a = 1 \\ (1 - \tilde{p})(\hat{Y}_i^{(t,+j)} - \hat{Y}_i^{(t,-j)}) , & a = 2 \\ ((1 - \tilde{p})\hat{Y}_i^{(t)} + \tilde{p}/2 - Y_i) , & a = 3 \\ \frac{-(1-\tilde{p})(\hat{Y}_i^{(t)})^2 + (1-\tilde{p})^2\hat{Y}_i^{(t)} + \frac{\tilde{p}}{2} - \frac{\tilde{p}^2}{4}}{K} , & a = 4 \\ \frac{-(1-\tilde{p})(\hat{Y}_i^{(t,+j)})^2 + (\hat{Y}_i^{(t,-j)})^2 + (1-\tilde{p})^2(\hat{Y}_i^{(t,+j)} + \hat{Y}_i^{(t,-j)}) + \tilde{p} - \frac{\tilde{p}^2}{2}}{K} , & a = 5 , \end{cases}$$

where $\hat{Y}_i^{(t,\pm j)} = \text{Tr}(\Pi U(\boldsymbol{\theta} \pm \mathbf{e}_j) \rho_{\mathcal{B}_i} U(\boldsymbol{\theta} \pm \mathbf{e}_j)^\dagger)$, K refers to the number of quantum measurements, and $\hat{Y}_i^{(t)}$ and Y_i are the sum average of the predicted and true labels for the i -th batch \mathcal{B}_i .

The relation between the estimated and analytic gradients follows

$$\nabla_j \bar{\mathcal{L}}_i(\boldsymbol{\theta}^{(t)}) = (1 - \tilde{p})^2 \nabla_j \mathcal{L}_i(\boldsymbol{\theta}^{(t)}) + C_{j,1}^{(i,t)} + \boldsymbol{\varsigma}_i^{(t,j)}$$

with $\boldsymbol{\varsigma}_i^{(t,j)} = C_{j,2}^{(i,t)} \xi_i^{(t)} + C_{j,3}^{(i,t)} \xi_i^{(t,j)} + \xi^{(t)} \xi_i^{(t,j)}$, where $\xi_i^{(t)}$ and $\xi_i^{(t,j)}$ are two random variables with zero mean and variances $C_{j,4}^{(i,t)}$ and $C_{j,5}^{(i,t)}$, respectively.

The intuition to achieve Theorem 5 is as follows. As explained in the main text, the discrepancy between the estimated gradient $\nabla_j \bar{\mathcal{L}}_i(\boldsymbol{\theta}^{(t)})$ and the analytic gradient $\nabla_j \mathcal{L}_i(\boldsymbol{\theta}^{(t)})$ is caused by the difference between the estimated results $\bar{Y}_i^{(t)}$ (or $\bar{Y}_i^{(t,\pm j)}$) and the expected results $\hat{Y}_i^{(t)}$ (or $\hat{Y}_i^{(t,\pm j)}$), due to the involved depolarization noise \mathcal{N}_p and the finite number of measurements K . Specifically, the noisy channel \mathcal{N}_p shifts the expectation values, and the finite number of measurements K turns the output of quantum circuit from the determination to be random. Under the above observation, the estimated gradients $\nabla_j \bar{\mathcal{L}}_i(\boldsymbol{\theta}^{(t)})$ can be treated as the random variable that is formed

by three random variables $\bar{Y}_i^{(t)}$ and $\bar{Y}_i^{(t,\pm j)}$, where the probability distributions of $\bar{Y}_i^{(t)}$ and $\bar{Y}_i^{(t,\pm j)}$ are determined by K , \mathcal{N}_p , $\hat{Y}_i^{(t)}$, and $\hat{Y}_i^{(t,\pm j)}$. Therefore, to explicitly build the relation between $\nabla_j \bar{\mathcal{L}}_i(\boldsymbol{\theta}^{(t)})$ and $\nabla_j \mathcal{L}_i(\boldsymbol{\theta}^{(t)})$, we should first formulate the distribution of the estimated gradients using $\bar{Y}_i^{(t)}$ and $\bar{Y}_i^{(t,\pm j)}$, and then connect the obtained distribution with the analytic gradients. The following lemma summarizes the distribution of the estimated gradients using $\bar{Y}_i^{(t)}$ and $\bar{Y}_i^{(t,\pm j)}$, whose proof is given in Subsection D.1.

Lemma 6. *The mean $\nu_i^{(t)}$ and variance $(\sigma_i^{(t)})^2$ of the estimated result $\bar{Y}_i^{(t)}$ are*

$$\begin{aligned} \nu^{(t)} &= (1 - \tilde{p})\hat{Y}_i^{(t)} + \tilde{p}\frac{\text{Tr}(\Pi)}{D}, \\ (\sigma_i^{(t)})^2 &= \frac{-(1 - \tilde{p})^2(\hat{Y}_i^{(t)})^2 + (1 - \tilde{p})\left(1 - 2\tilde{p}\frac{\text{Tr}(\Pi)}{D}\right)\hat{Y}_i^{(t)} + \tilde{p}\frac{\text{Tr}(\Pi)}{D} - \tilde{p}^2\frac{(\text{Tr}(\Pi))^2}{D^2}}{K}. \end{aligned} \quad (20)$$

The mean $\nu_i^{(t,\pm j)}$ and variance $(\sigma_i^{(t,\pm j)})^2$ of the estimated results $\bar{Y}_i^{(t,\pm j)}$ are

$$\begin{aligned} \nu^{(t,\pm j)} &= (1 - \tilde{p})\hat{Y}_i^{(t,\pm j)} + \tilde{p}\frac{\text{Tr}(\Pi)}{D}, \\ (\sigma_i^{(t,\pm j)})^2 &= \frac{-(1 - \tilde{p})^2(\hat{Y}_i^{(t,\pm j)})^2 + (1 - \tilde{p})\left(1 - 2\tilde{p}\frac{\text{Tr}(\Pi)}{D}\right)\hat{Y}_i^{(t,\pm j)} + \tilde{p}\frac{\text{Tr}(\Pi)}{D} - \tilde{p}^2\frac{(\text{Tr}(\Pi))^2}{D^2}}{K}. \end{aligned} \quad (21)$$

Proof of Theorem 5. We now utilize the established relations as shown in Lemma 6 to obtain the relation between the estimated and the analytic gradients. Recall that, at the t -th iteration, given the input \mathcal{B}_i and K measurements, the estimated gradient for j -th parameter $\boldsymbol{\theta}_j$ of noisy QNN is

$$\nabla_j \bar{\mathcal{L}}_i(\boldsymbol{\theta}^{(t)}) = (\bar{Y}_i^{(t)} - Y_i) \left(\bar{Y}_i^{(t,+j)} - \bar{Y}_i^{(t,-j)} \right) + \lambda \boldsymbol{\theta}_j. \quad (22)$$

Combining Lemma 6 and Eqn. (22), the term $\Delta_i^{(t,j)} := \bar{Y}_i^{(t,+j)} - \bar{Y}_i^{(t,-j)}$ in Eqn. (22) can be treated as the difference of two random variables. The term $(\bar{Y}_i^{(t)} - Y_i)$ in Eqn. (22) can also be treated as a random variable. We now separately investigate their moment properties.

The term $\Delta_i^{(t,j)}$. Following the notations used in Lemma 6, the mean and variance of the term $\Delta_i^{(t,j)}$ are $\nu_i^{(t,+j)} - \nu_i^{(t,-j)}$ and $(\sigma_i^{(t,j)})^2 = (\sigma_i^{(t,+j)})^2 + (\sigma_i^{(t,-j)})^2$, supported by the definition of moments and the independent relation between $\bar{Y}_i^{(t,+j)}$ and $\bar{Y}_i^{(t,-j)}$.

By leveraging the explicit form of $\nu_i^{(t,\pm j)}$, the random variable $\Delta_i^{(t,j)}$ can be rewritten as

$$\Delta_i^{(t,j)} = (1 - \tilde{p})(\hat{Y}^{(t,+j)} - \hat{Y}^{(t,-j)}) + \xi^{(t,j)}, \quad (23)$$

where $\xi^{(t,j)}$ is a random variable with zero mean and variance $(\sigma_i^{(t,j)})^2$.

The term $(\bar{Y}_i^{(t)} - Y_i)$. Following the notations used in Lemma 6, an equivalent representation of $(\bar{Y}_i^{(t)} - Y_i)$ is

$$(\bar{Y}_i^{(t)} - Y_i) = (1 - \tilde{p})\hat{Y}_i^{(t)} + \tilde{p}\frac{\text{Tr}(\Pi)}{D} + \xi^{(t)} - Y_i, \quad (24)$$

where $\xi^{(t)}$ is a random variable with zero mean and variance $(\sigma_i^{(t)})^2$.

The reformulated terms as shown in Eqn. (23) and Eqn. (24) indicate that the estimated result $\nabla_j \bar{\mathcal{L}}_i(\boldsymbol{\theta}^{(t)})$ can be rewritten as

$$\nabla_j \bar{\mathcal{L}}_i(\boldsymbol{\theta}^{(t)})$$

$$\begin{aligned}
&= (\bar{Y}_i^{(t)} - Y_i)(\bar{Y}_i^{(t,+j)} - \bar{Y}_i^{(t,-j)}) + \lambda \boldsymbol{\theta}_j^{(t)} \\
&= \left((1 - \tilde{p}) \hat{Y}_i^{(t)} + \tilde{p} \frac{\text{Tr}(\Pi)}{D} - Y_i \right) (1 - \tilde{p})(\hat{Y}^{(t,+j)} - \hat{Y}^{(t,-j)}) + \left((1 - \tilde{p}) \hat{Y}_i^{(t)} + \tilde{p} \frac{\text{Tr}(\Pi)}{D} - Y_i \right) \xi^{(t,j)} \\
&\quad + (1 - \tilde{p})(\hat{Y}^{(t,+j)} - \hat{Y}^{(t,-j)}) \xi^{(t)} + \xi^{(t)} \xi^{(t,j)} + \lambda \boldsymbol{\theta}_j^{(t)} \\
&= (1 - \tilde{p})^2 \nabla_j \mathcal{L}_i(\boldsymbol{\theta}^{(t)}) + (1 - \tilde{p}) \tilde{p} \left(\frac{\text{Tr}(\Pi)}{D} - Y_i \right) (\hat{Y}^{(t,+j)} - \hat{Y}^{(t,-j)}) + (2\tilde{p} - \tilde{p}^2) \lambda \boldsymbol{\theta}_j^{(t)} \\
&\quad + (1 - \tilde{p})(\hat{Y}^{(t,+j)} - \hat{Y}^{(t,-j)}) \xi^{(t)} + \left((1 - \tilde{p}) \hat{Y}_i^{(t)} + \tilde{p} \frac{\text{Tr}(\Pi)}{D} - Y_i \right) \xi^{(t,j)} + \xi^{(t)} \xi^{(t,j)}. \tag{25}
\end{aligned}$$

Combining the above equation and the explicit expression of $\xi^{(t)}$ and $\xi^{(t,j)}$, we obtain the relation between the estimated and the analytic gradients. Specifically, the estimated gradient can be formulated as

$$\nabla_j \bar{\mathcal{L}}_i(\boldsymbol{\theta}^{(t)}) = (1 - \tilde{p})^2 \nabla_j \mathcal{L}_i(\boldsymbol{\theta}^{(t)}) + C_{j,1}^{(i,t)} + \boldsymbol{\varsigma}_i^{(t,j)},$$

where $\boldsymbol{\varsigma}_i^{(t,j)} = C_{j,2}^{(i,t)} \xi_i^{(t)} + C_{j,3}^{(i,t)} \xi_i^{(t,j)} + \xi^{(t)} \xi_i^{(t,j)}$, the first three constants $\{C_{j,1}^{(i,t)}\}_{i=1}^3$ are defined as

$$C_{j,a}^{(i,t)} = \begin{cases} (1 - \tilde{p}) \tilde{p} \left(\frac{\text{Tr}(\Pi)}{D} - Y_i \right) (\hat{Y}^{(t,+j)} - \hat{Y}^{(t,-j)}) + (2\tilde{p} - \tilde{p}^2) \lambda \boldsymbol{\theta}_j^{(t)}, & a = 1 \\ (1 - \tilde{p})(\hat{Y}_i^{(t,+j)} - \hat{Y}_i^{(t,-j)}), & a = 2 \\ \left((1 - \tilde{p}) \hat{Y}_i^{(t)} + \tilde{p} \frac{\text{Tr}(\Pi)}{D} - Y_i \right), & a = 3, \end{cases}$$

and the last two constants, which separately correspond to the variance $(\sigma_i^{(t)})^2$ and $(\sigma_i^{(t,j)})^2$ of the random variables $\xi_i^{(t)}$ and $\xi_i^{(t,j)}$, are

$$C_{j,a}^{(i,t)} = \begin{cases} \frac{-(1-\tilde{p})^2(\hat{Y}_i^{(t)})^2 + (1-\tilde{p})\left(1-2\tilde{p}\frac{\text{Tr}(\Pi)}{D}\right)\hat{Y}_i^{(t)} + \tilde{p}\frac{\text{Tr}(\Pi)}{D} - \tilde{p}^2\frac{(\text{Tr}(\Pi))^2}{D^2}}{K}, & a = 4 \\ \frac{-(1-\tilde{p})^2((\hat{Y}_i^{(t,+j)})^2 + (\hat{Y}_i^{(t,-j)})^2) + (1-\tilde{p})\left(1-2\tilde{p}\frac{\text{Tr}(\Pi)}{D}\right)(\hat{Y}_i^{(t,+j)} + \hat{Y}_i^{(t,-j)}) + 2\tilde{p}\frac{\text{Tr}(\Pi)}{D} - 2\tilde{p}^2\frac{(\text{Tr}(\Pi))^2}{D^2}}{K}, & a = 5. \end{cases}$$

□

D.1 Proof of Lemma 6

To achieve Lemma 6, we first simplify the learning model of QNN with the depolarization noise. In particular, all noisy channels \mathcal{N}_p , which are separately applied to each quantum circuit depth, can be merged together to a specific circuit depth and presented by a new depolarization channel $\mathcal{N}_{\tilde{p}}$.

Lemma 7. *Let \mathcal{N}_p be the depolarization channel. There always exists a depolarization channel $\mathcal{N}_{\tilde{p}}$ with $\tilde{p} = 1 - (1 - p)^{L_Q}$ that satisfies $\mathcal{N}_p(U_L(\boldsymbol{\theta}) \dots U_2(\boldsymbol{\theta}) \mathcal{N}_p(U_1(\boldsymbol{\theta}) \rho U_1(\boldsymbol{\theta})^\dagger) U_2(\boldsymbol{\theta})^\dagger \dots U_L(\boldsymbol{\theta})^\dagger) = \mathcal{N}_{\tilde{p}}(U(\boldsymbol{\theta}) \rho U(\boldsymbol{\theta})^\dagger)$, where ρ is the input quantum state.*

Proof of Lemma 7. Denote $\rho^{(k)}$ as $\rho^{(k)} = \prod_{l=1}^k U_l(\boldsymbol{\theta}) \rho U_l(\boldsymbol{\theta})^\dagger$. Applying \mathcal{N}_p to $\rho^{(1)}$ gives

$$\mathcal{N}_p(\rho^{(1)}) = (1 - p)\rho^{(1)} + p \frac{\mathbb{I}_D}{D}, \tag{26}$$

where D refers to the dimensions of Hilbert space interacted with \mathcal{N}_p .

Supporting by the above equation, applying $U_2(\boldsymbol{\theta})$ to the state $\mathcal{N}_p(\rho^{(1)})$ gives

$$U_2(\boldsymbol{\theta}) \mathcal{N}_p(\rho^{(1)}) U_2(\boldsymbol{\theta})^\dagger = (1 - p)\rho^{(2)} + p \frac{\mathbb{I}_D}{D}. \tag{27}$$

Then interacting \mathcal{N}_p with the state $U_2(\boldsymbol{\theta})\mathcal{N}_p(\rho^{(1)})U_2(\boldsymbol{\theta})^\dagger$ gives

$$\mathcal{N}_p(U_2(\boldsymbol{\theta})\mathcal{N}_p(\rho^{(1)})U_2(\boldsymbol{\theta})^\dagger) = (1-p)^2\rho^{(2)} + (1-p)p\frac{\mathbb{I}_D}{D} + p\frac{\mathbb{I}_D}{D} = (1-p)^2\rho^{(2)} + (1-(1-p)^2)\frac{\mathbb{I}_D}{D}. \quad (28)$$

By induction, suppose at k -th step, the generated state is

$$\rho^{(k)} = (1-p)^k\rho^{(k)} + (1-(1-p)^k)\frac{\mathbb{I}_D}{D}. \quad (29)$$

Then applying $U_{k+1}(\boldsymbol{\theta})$ followed by \mathcal{N}_p gives

$$\rho^{(k+1)} = \mathcal{N}_p\left(U_{k+1}(\boldsymbol{\theta})\rho^{(k)}U_{k+1}(\boldsymbol{\theta})^\dagger\right) = (1-p)^{k+1}\rho^{(k+1)} + (1-(1-p)^{k+1})\frac{\mathbb{I}_D}{D}. \quad (30)$$

According to the formula of depolarization channel, an immediate observation is that the noisy QNN is equivalent to applying a single depolarization channel $\mathcal{N}_{\tilde{p}}$ at the last circuit depth L , i.e.,

$$\mathcal{N}_{\tilde{p}}(\rho) = (1-p)^L\rho^{(L)} + (1-(1-p)^L)\frac{\mathbb{I}}{D}, \quad (31)$$

where

$$\tilde{p} = 1 - (1-p)^L. \quad (32)$$

□

We then use the simplified QNN given by Lemma 7 to explore the relation between the generated statistic $\bar{Y}_i^{(t)}$ and the expectation value $\hat{Y}^{(t)}$ (the same rule applies to connect $\bar{Y}_i^{(t,\pm j)}$ with $\hat{Y}^{(t,\pm j)}$).

At the t -th iteration, given the tunable parameters $\boldsymbol{\theta}^{(t)}$ and inputs \mathcal{B}_i , the ensemble corresponding to the generated state of QNN before taking quantum measurements is $\{p_l, \gamma_{i,l}^{(t)}\}_{l=1}^2$, i.e., $p_1 = 1 - \tilde{p}$ with $\gamma_{i,1}^{(t)} = U(\boldsymbol{\theta}^{(t)})\rho_{\mathcal{B}_i}U(\boldsymbol{\theta}^{(t)})^\dagger$ and $p_2 = \tilde{p}$ with $\gamma_{i,2}^{(t)} = \mathbb{I}_D/D$. After applying a two-outcome POVM Π to measure such an ensemble K times, the generated statistics (sample mean) is $\bar{Y}_i^{(t)} = \frac{1}{K} \sum_{k=1}^K V_k^{(t)}$, where each measured outcome $V_k^{(t)}$ with $k \in [K]$ is a random variable that satisfies Fact 1.

Fact 1. $V_k^{(t)}$ is a random variable that follows the distribution $\mathcal{P}_{Q'}(V_k^{(t)}) = \sum_{c=1}^2 \Pr(z=c) \Pr(V_k^{(t)}|z=c)$. The explicit formula of $\mathcal{P}_{Q'}$ is

1. $\Pr(z=1) = 1 - \tilde{p}$ with $V_k^{(t)}|z=1 \sim \text{Ber}(\hat{Y}_i^{(t)})$ and $\hat{Y}_i^{(t)} = \text{Tr}(\Pi\gamma_{i,1}^{(t)})$;
2. $\Pr(z=2) = \tilde{p}$ with $V_k^{(t)}|z=2 \sim \text{Ber}(\frac{\text{Tr}(\Pi)}{D})$ with $\frac{\text{Tr}(\Pi)}{D} = \text{Tr}(\Pi\gamma_{i,2}^{(t)})$.

Fact 1 implies that the mean and variance of $V_k^{(t)}$ are

$$(1-\tilde{p})\hat{Y}_i^{(t)} + \tilde{p}\frac{\text{Tr}(\Pi)}{D} \text{ and } -(1-\tilde{p})^2(\hat{Y}_i^{(t)})^2 + (1-\tilde{p})\left(1-2\tilde{p}\frac{\text{Tr}(\Pi)}{D}\right)\hat{Y}_i^{(t)} + \tilde{p}\frac{\text{Tr}(\Pi)}{D} - \tilde{p}^2\frac{(\text{Tr}(\Pi))^2}{D^2},$$

respectively. Moreover, since each outcome $V_k^{(t)}$ follows the distribution $\mathcal{P}_{Q'}$, the mean $\nu_i^{(t)}$ and the variance $(\sigma_i^{(t)})^2$ of the sample mean $\bar{Y}_i^{(t)}$ are

$$\begin{aligned} \nu^{(t)} &= (1-\tilde{p})\hat{Y}_i^{(t)} + \tilde{p}\frac{\text{Tr}(\Pi)}{D}, \\ (\sigma_i^{(t)})^2 &= \frac{-(1-\tilde{p})^2(\hat{Y}_i^{(t)})^2 + (1-\tilde{p})\left(1-2\tilde{p}\frac{\text{Tr}(\Pi)}{D}\right)\hat{Y}_i^{(t)} + \tilde{p}\frac{\text{Tr}(\Pi)}{D} - \tilde{p}^2\frac{(\text{Tr}(\Pi))^2}{D^2}}{K}. \end{aligned} \quad (33)$$

Following the same routine, where the mean $\nu_i^{(t,\pm j)}$ and the variance $(\sigma_i^{(t,\pm j)})^2$ of the sample mean $\bar{Y}_i^{(t,\pm j)}$ satisfy

$$\begin{aligned} \nu^{(t,\pm j)} &= (1 - \tilde{p})\hat{Y}_i^{(t,\pm j)} + \tilde{p}\frac{\text{Tr}(\Pi)}{D}, \\ (\sigma_i^{(t,\pm j)})^2 &= \frac{-(1 - \tilde{p})^2(\hat{Y}_i^{(t,\pm j)})^2 + (1 - \tilde{p})\left(1 - 2\tilde{p}\frac{\text{Tr}(\Pi)}{D}\right)\hat{Y}_i^{(t,\pm j)} + \tilde{p}\frac{\text{Tr}(\Pi)}{D} - \tilde{p}^2\frac{(\text{Tr}(\Pi))^2}{D^2}}{K}. \end{aligned} \quad (34)$$

E Proof of Theorem 1

Theorem 1 quantifies the utility bounds R_1 and R_2 of QNN under the depolarization noise towards ERM framework. For ease of illustration, we restate Theorem 1 below.

Theorem 6 (Restate of Theorem 1). *QNN outputs $\theta^{(T)} \in \mathbb{R}^d$ after T iterations with utility bounds $R_1 \leq \tilde{O}\left(d, \frac{1}{BK}, \frac{1}{(1-p)^{L_Q}}\right)$ and $R_2 \leq \tilde{O}\left(\frac{1}{K^2B}, d, \frac{1}{(1-p)^{L_Q}}\right)$, where K is the number of quantum measurements, L_Q is the quantum circuit depth, p is the gate noise, and B is the number of batches.*

The high level idea to achieve the utility bounds R_1 and R_2 is as follows. Recall that R_1 measures how far the trainable parameter of QNN is away from the stationary point. A well-known result in optimization theory [76] is that when a function satisfies the smooth property, its stationary point can be efficiently located by a simple gradient-based algorithm. By leveraging this observation and the relation between the estimated and analytic gradients as achieved in Theorem 5, we can quantify how the estimated gradients of QNN converge to the stationary point, which corresponds to the utility bound R_1 .

Recall that the utility bound R_2 evaluates the disparity between the expected empirical risk and the optimal risk that is determined by the global minimum. To achieve R_2 , we utilize the result of the study [43], which claims that if a non-convex function satisfies PL condition, then every stationary point is the global minimum. Since the objective function used in QNN satisfies PL condition as shown in Lemma 2, we can effectively combine the PL condition with the result of R_1 to obtain the utility bound R_2 .

Proof of Theorem 6. We employ the following two theorems to achieve Theorem 6, whose proofs are given in Subsections E.1 and E.2, respectively.

Theorem 7. *Given the dataset \mathbf{z} , QNN outputs $\theta^{(T)}$ after T iterations with utility bound*

$$R_1 \leq \frac{2S(1 + 90\lambda d)}{T(1 - \tilde{p})^2} + \frac{(2\tilde{p} - \tilde{p}^2)(2G + d)(1 + 10\lambda)^2}{(1 - \tilde{p})^2} + \frac{6dK + 8d}{(1 - \tilde{p})^2BK^2}.$$

Theorem 8. *Given the dataset \mathbf{z} , QNN outputs $\theta^{(T)}$ after T iterations with utility bound*

$$R_2 \leq (1 + 90\lambda d) \exp\left(-\frac{\mu(1 - \tilde{p})^2 T}{S}\right) + T \frac{(2\tilde{p} - \tilde{p}^2)(G + 2d)(1 + 10\lambda)^2 BK^2 + 6dK + 8d}{2SBK^2}.$$

As for R_1 , with setting $T \leftarrow \infty$ and after the simplification, the utility bound as shown in Theorem 7 follows

$$R_1 \leq \tilde{O}\left(\frac{1}{BK}, \frac{1}{(1-p)^{L_Q}}, d\right). \quad (35)$$

As for R_2 , with setting $T = \mathcal{O}\left(\frac{S}{\mu(1-\tilde{p})^2} \ln\left(\frac{(1+90\lambda d)2SBK^2}{(2\tilde{p}-\tilde{p}^2)(G+2d)(1+10\lambda)^2BK^2+6dK+8d}\right)\right)$ and after simplification, the utility bound as shown in Theorem 8 follows

$$R_2 \leq \tilde{O}\left(\frac{1}{K^2B}, d, \frac{1}{(1-p)^{L_Q}}\right). \quad (36)$$

□

E.1 Proof of Theorem 7: The utility bound R_1

The proof of Theorem 7 employs the following Lemma, where its proof is given in Subsection E.3.

Lemma 8. *Taking expectation over the randomness of $\xi_i^{(t)}$ and $\xi_i^{(t,j)}$ in the estimated gradient $\nabla_j \bar{\mathcal{L}}(\boldsymbol{\theta}^{(t)})$ as formulated in Theorem 5, the term $\frac{1}{2S} \sum_{j=1}^d \mathbb{E}_{\xi_i^{(t)}, \xi_i^{(t,j)}} \left[(\nabla_j \bar{\mathcal{L}}(\boldsymbol{\theta}^{(t)}))^2 \right]$ with S being the smooth parameter is upper bounded by*

$$\frac{(1-\tilde{p})^4}{2S} \|\nabla \mathcal{L}(\boldsymbol{\theta}^{(t)})\|^2 + \frac{(1-\tilde{p})^2 G}{2S} \max_{i,j} C_{j,1}^{(i,t)} + \frac{d}{2S} \max_{i,j} \left(C_{j,1}^{(i,t)} \right)^2 + \frac{6dK+8d}{2SBK^2}.$$

Proof of Theorem 7. Recall that the optimization rule of noisy QNN at the t -th iteration follows

$$\boldsymbol{\theta}^{(t+1)} = \boldsymbol{\theta}^{(t)} - \eta \nabla \bar{\mathcal{L}}(\boldsymbol{\theta}^{(t)}). \quad (37)$$

Since the objective function $\mathcal{L}(\boldsymbol{\theta})$ is S -smooth, as indicated in Lemma 2, we have

$$\mathcal{L}(\boldsymbol{\theta}^{(t+1)}) - \mathcal{L}(\boldsymbol{\theta}^{(t)}) \leq \langle \nabla \mathcal{L}(\boldsymbol{\theta}^{(t)}), \boldsymbol{\theta}^{(t+1)} - \boldsymbol{\theta}^{(t)} \rangle + \frac{S}{2} \|\boldsymbol{\theta}^{(t+1)} - \boldsymbol{\theta}^{(t)}\|^2. \quad (38)$$

Combine the above two equations and setting $\eta = 1/S$, we have

$$\begin{aligned} & \mathcal{L}(\boldsymbol{\theta}^{(t+1)}) - \mathcal{L}(\boldsymbol{\theta}^{(t)}) \\ & \leq \langle \nabla \mathcal{L}(\boldsymbol{\theta}^{(t)}), \boldsymbol{\theta}^{(t+1)} - \boldsymbol{\theta}^{(t)} \rangle + \frac{S}{2} \|\boldsymbol{\theta}^{(t+1)} - \boldsymbol{\theta}^{(t)}\|^2 \\ & = -\frac{1}{S} \langle \nabla \mathcal{L}(\boldsymbol{\theta}^{(t+1)}), \nabla \bar{\mathcal{L}}(\boldsymbol{\theta}^{(t)}) \rangle + \frac{1}{2S} \|\nabla \bar{\mathcal{L}}(\boldsymbol{\theta}^{(t)})\|^2 \\ & = -\frac{1}{S} \sum_{j=1}^d \left(\nabla_j \mathcal{L}(\boldsymbol{\theta}^{(t+1)}) \nabla_j \bar{\mathcal{L}}(\boldsymbol{\theta}^{(t)}) \right) + \frac{1}{2S} \sum_{j=1}^d \left(\nabla_j \bar{\mathcal{L}}(\boldsymbol{\theta}^{(t)}) \right)^2. \end{aligned} \quad (39)$$

Recall the definition of the estimated gradient is $\nabla_j \bar{\mathcal{L}}(\boldsymbol{\theta}^{(t)}) = \frac{1}{B} \sum_{i=1}^B \nabla_j \bar{\mathcal{L}}_i(\boldsymbol{\theta}^{(t)})$ and the explicit expression of $\nabla_j \bar{\mathcal{L}}_i(\boldsymbol{\theta}^{(t)})$ is

$$\nabla_j \bar{\mathcal{L}}_i(\boldsymbol{\theta}^{(t)}) = (1-\tilde{p})^2 \nabla_j \mathcal{L}_i(\boldsymbol{\theta}^{(t)}) + C_{j,1}^{(i,t)} + C_{j,2}^{(i,t)} \xi_i^{(t)} + C_{j,3}^{(i,t)} \xi_i^{(t,j)} + \xi_i^{(t)} \xi_i^{(t,j)}.$$

Alternatively, the gradient for the j -th parameter $\nabla_j \bar{\mathcal{L}}(\boldsymbol{\theta}^{(t)})$ follows

$$\nabla_j \bar{\mathcal{L}}(\boldsymbol{\theta}^{(t)}) = \frac{1}{B} \sum_{i=1}^B (1-\tilde{p})^2 \nabla_j \mathcal{L}_i(\boldsymbol{\theta}^{(t)}) + C_{j,1}^{(i,t)} + C_{j,2}^{(i,t)} \xi_i^{(t)} + C_{j,3}^{(i,t)} \xi_i^{(t,j)} + \xi_i^{(t)} \xi_i^{(t,j)}. \quad (40)$$

Combining Eqn. (39) with Eqn. (40) and taking expectation over $\xi_i^{(t)}$ and $\xi_i^{(t,j)}$, we obtain

$$\mathbb{E}_{\xi_i^{(t)}, \xi_i^{(t,j)}} [\mathcal{L}(\boldsymbol{\theta}^{(t+1)}) - \mathcal{L}(\boldsymbol{\theta}^{(t)})]$$

$$\begin{aligned}
&\leq -\frac{1}{S}(1-\tilde{p})^2\|\nabla\mathcal{L}(\boldsymbol{\theta}^{(t)})\|^2 - \frac{1}{S}\sum_{j=1}^d\nabla_j\mathcal{L}(\boldsymbol{\theta}^{(t)})\left(\frac{1}{B}\sum_{i=1}^BC_{j,1}^{(i,t)}\right) \\
&\quad - \frac{1}{S}\sum_{j=1}^d\nabla_j\mathcal{L}(\boldsymbol{\theta}^{(t)})\frac{1}{B}\sum_{i=1}^B\mathbb{E}_{\xi_i^{(t)}}\left[C_{j,2}^{(i,t)}\xi_i^{(t)}\right] - \frac{1}{S}\sum_{j=1}^d\nabla_j\mathcal{L}(\boldsymbol{\theta}^{(t)})\frac{1}{B}\sum_{i=1}^B\mathbb{E}_{\xi_i^{(t,j)}}\left[C_{j,3}^{(i,t)}\xi_i^{(t,j)}\right] \\
&\quad - \frac{1}{S}\sum_{j=1}^d\nabla_j\mathcal{L}(\boldsymbol{\theta}^{(t)})\frac{1}{B}\sum_{i=1}^B\mathbb{E}_{\xi_i^{(t)},\xi_i^{(t,j)}}\left[\xi_i^{(t)}\xi_i^{(t,j)}\right] + \frac{1}{2S}\sum_{j=1}^d\mathbb{E}_{\xi_i^{(t)},\xi_i^{(t,j)}}\left[\left(\nabla_j\bar{\mathcal{L}}(\boldsymbol{\theta}^{(t)})\right)^2\right] \\
&\leq -\frac{1}{S}(1-\tilde{p})^2\|\nabla\mathcal{L}(\boldsymbol{\theta}^{(t)})\|^2 + \frac{G}{2S}\max_{i,j}C_{j,1}^{(i,t)} + \frac{1}{2S}\sum_{j=1}^d\mathbb{E}_{\xi_i^{(t)},\xi_i^{(t,j)}}\left[\left(\nabla_j\bar{\mathcal{L}}(\boldsymbol{\theta}^{(t)})\right)^2\right]. \tag{41}
\end{aligned}$$

The first inequality uses the result of Eqn. (40). The second inequality uses $\mathbb{E}[\xi_i^{(t)}] = 0$, $\mathbb{E}[\xi_i^{(t,j)}] = 0$ as shown in Theorem 5, and $-G/d \leq \nabla_j\mathcal{L}(\boldsymbol{\theta}^{(t)}) \leq G/d$ supported by G -Lipschitz property.

By leveraging Lemma 8, Eqn. (41) can be further simplified as

$$\begin{aligned}
&\mathbb{E}_{\xi_i^{(t)},\xi_i^{(t,j)}}[\mathcal{L}(\boldsymbol{\theta}^{(t+1)}) - \mathcal{L}(\boldsymbol{\theta}^{(t)})] \\
&\leq -\frac{1}{S}(1-\tilde{p})^2\|\nabla\mathcal{L}(\boldsymbol{\theta}^{(t)})\|^2 + \frac{G}{2S}\max_{i,j}C_{j,1}^{(i,t)} + \frac{(1-\tilde{p})^4}{2SB}\|\nabla_j\mathcal{L}(\boldsymbol{\theta}^{(t)})\|^2 \\
&\quad + \frac{(1-\tilde{p})^2G}{2S}\max_{i,j}C_{j,1}^{(i,t)} + \frac{d}{2S}\max_{i,j}\left(C_{j,1}^{(i,t)}\right)^2 + \frac{6dK+8d}{2SBK^2} \\
&\leq -\frac{1}{2S}(1-\tilde{p})^2\|\nabla\mathcal{L}(\boldsymbol{\theta}^{(t)})\|^2 + \frac{2G+d}{2S}(2-\tilde{p})\tilde{p}(1+10\lambda)^2 + \frac{6dK+8d}{2SBK^2}. \tag{42}
\end{aligned}$$

The first inequalities comes from Lemma 8, and the second inequality employs $\frac{(1-\tilde{p})^4}{2SB} \leq \frac{(1-\tilde{p})^2}{2S}$ and the following result

$$\begin{aligned}
&\frac{G}{2S}\max_{i,j}C_{j,1}^{(i,t)} + \frac{(1-\tilde{p})^2G}{2S}\max_{i,j}C_{j,1}^{(i,t)} + \frac{d}{2S}\max_{i,j}\left(C_{j,1}^{(i,t)}\right)^2 \\
&\leq \frac{(1+(1-\tilde{p})^2)G}{2S}(2-\tilde{p})\tilde{p}(1+10\lambda) + \frac{d}{2S}(2-\tilde{p})\tilde{p}(1+10\lambda)^2 \\
&\leq \frac{2G+d}{2S}(2-\tilde{p})\tilde{p}(1+10\lambda)^2, \tag{43}
\end{aligned}$$

where the first inequality uses the upper bound of $C_{j,1}^{(i,t)}$ and $(C_{j,1}^{(i,t)})^2$, i.e., $\max_{i,j}C_{j,1}^{(i,t)} \leq (1-\tilde{p})\tilde{p} + 10(2-\tilde{p})\tilde{p}\lambda \leq (2-\tilde{p})\tilde{p}(1+10\lambda)$ and $\max_{i,j}\left(C_{j,1}^{(i,t)}\right)^2 \leq ((2-\tilde{p})\tilde{p}(1+10\lambda))^2 \leq (2-\tilde{p})\tilde{p}(1+10\lambda)^2$, and the second inequality uses $(1-\tilde{p})^2 \leq 1$.

An equivalent representation of Eqn. (42) is

$$\|\nabla\mathcal{L}(\boldsymbol{\theta}^{(t)})\|^2 \leq 2S\frac{\mathcal{L}(\boldsymbol{\theta}^{(t)}) - \mathbb{E}_{\xi_i^{(t)},\xi_i^{(t,j)}}[\mathcal{L}(\boldsymbol{\theta}^{(t+1)})]}{(1-\tilde{p})^2} + \frac{(2\tilde{p}-\tilde{p}^2)(2G+d)(1+10\lambda)^2}{(1-\tilde{p})^2} + \frac{6dK+8d}{(1-\tilde{p})^2BK^2}. \tag{44}$$

By induction, with summing over $t = 0, \dots, T - 1$ and taking expectation of Eqn. (44), we obtain

$$\begin{aligned}
& \mathbb{E}_t \left[\|\nabla \mathcal{L}(\boldsymbol{\theta}^{(t)})\|^2 \right] \\
& \leq 2S \frac{\mathcal{L}(\boldsymbol{\theta}^{(0)}) - \mathbb{E}_{\xi_i^{(T)}, \xi_i^{(T,j)}}[\mathcal{L}(\boldsymbol{\theta}^{(T)})]}{T(1-\tilde{p})^2} + \frac{(2\tilde{p} - \tilde{p}^2)(2G+d)(1+10\lambda)^2}{(1-\tilde{p})^2} + \frac{6dK+8d}{(1-\tilde{p})^2 BK^2} \\
& \leq \frac{2S + 2S\lambda d(3\pi)^2}{T(1-\tilde{p})^2} + \frac{(2\tilde{p} - \tilde{p}^2)(2G+d)(1+10\lambda)^2}{(1-\tilde{p})^2} + \frac{6dK+8d}{(1-\tilde{p})^2 BK^2} \\
& \leq \frac{2S(1+90\lambda d)}{T(1-\tilde{p})^2} + \frac{(2\tilde{p} - \tilde{p}^2)(2G+d)(1+10\lambda)^2}{(1-\tilde{p})^2} + \frac{6dK+8d}{(1-\tilde{p})^2 BK^2}, \tag{45}
\end{aligned}$$

where the second inequality uses $\mathcal{L}(\boldsymbol{\theta}^{(0)}) - \mathbb{E}_{\xi_i^{(T)}, \xi_i^{(T,j)}}[\mathcal{L}(\boldsymbol{\theta}^{(T)})] \leq \mathcal{L}(\boldsymbol{\theta}^{(0)}) - \mathcal{L}^*$, $\mathcal{L}^* > 0$ and $\mathcal{L}(\boldsymbol{\theta}^{(0)}) \leq 1 + \lambda d(3\pi)^2$. \square

E.2 Proof of Theorem 8: The utility bound R_2

Proof of Theorem 8. The proof of Theorem 8 is similar with that of Theorem 7. In particular, following the same routine, we obtain the result of Eqn.(42), i.e.,

$$\begin{aligned}
& \mathbb{E}_{\xi_i^{(t)}, \xi_i^{(t,j)}}[\mathcal{L}(\boldsymbol{\theta}^{(t+1)}) - \mathcal{L}(\boldsymbol{\theta}^{(t)})] \\
& \leq -\frac{1}{2S}(1-\tilde{p})^2 \|\nabla \mathcal{L}(\boldsymbol{\theta}^{(t)})\|^2 + \frac{2G+d}{2S}(2-\tilde{p})\tilde{p}(1+10\lambda)^2 + \frac{6dK+8d}{2SBK^2}. \tag{46}
\end{aligned}$$

Then, we call the conclusion of PL condition as formulated in Lemma 2 and acquire

$$\begin{aligned}
& \mathbb{E}_{\xi_i^{(t)}, \xi_i^{(t,j)}}[\mathcal{L}(\boldsymbol{\theta}^{(t+1)}) - \mathcal{L}(\boldsymbol{\theta}^{(t)})] \\
& \leq -\frac{\mu(1-\tilde{p})^2}{S}(\mathcal{L}(\boldsymbol{\theta}^{(t)}) - \mathcal{L}^*) + \frac{2G+d}{2S}(2-\tilde{p})\tilde{p}(1+10\lambda)^2 + \frac{6dK+8d}{2SBK^2}. \tag{47}
\end{aligned}$$

An equivalent reformulation of Eqn. (47) is

$$\begin{aligned}
& \mathbb{E}_{\boldsymbol{\zeta}^{(t)}}[\mathcal{L}(\boldsymbol{\theta}^{(t+1)})] - \mathcal{L}^* \\
& \leq \left(1 - \frac{\mu(1-\tilde{p})^2}{S}\right)(\mathcal{L}(\boldsymbol{\theta}^{(t)}) - \mathcal{L}^*) + \frac{2G+d}{2S}(2-\tilde{p})\tilde{p}(1+10\lambda)^2 + \frac{6dK+8d}{2SBK^2}. \tag{48}
\end{aligned}$$

By induction, with summing over $t = 0, \dots, T$ and taking expectation, we obtain

$$\begin{aligned}
& \mathbb{E}_{\boldsymbol{\zeta}^{(t)}}[\mathcal{L}(\boldsymbol{\theta}^{(T)})] - \mathcal{L}^* \\
& \leq \left(1 - \frac{\mu(1-\tilde{p})^2}{S}\right)^T (\mathcal{L}(\boldsymbol{\theta}^{(0)}) - \mathcal{L}^*) + T \frac{2G+d}{2S}(2-\tilde{p})\tilde{p}(1+10\lambda)^2 + T \frac{6dK+8d}{2SBK^2} \\
& \leq (1+90\lambda d) \exp\left(-\frac{\mu(1-\tilde{p})^2 T}{S}\right) + T \frac{(2\tilde{p} - \tilde{p}^2)(G+2d)(1+10\lambda)^2 BK^2 + 6dK+8d}{2SBK^2}, \tag{49}
\end{aligned}$$

where the second inequality uses $\mathcal{L}(\boldsymbol{\theta}^{(0)}) - \mathcal{L}^* \leq 1 + 90\lambda d$ and $1 + x \leq e^x$ for all real x . \square

E.3 Proof of Lemma 8

Proof of Lemma 8. As shown in Theorem 5, the explicit formula of the estimated gradient is

$$\nabla_j \bar{\mathcal{L}}(\boldsymbol{\theta}^{(t)}) = \frac{1}{B} \sum_{i=1}^B (1 - \tilde{p})^2 \nabla_j \mathcal{L}_i(\boldsymbol{\theta}^{(t)}) + C_{j,1}^{(i,t)} + C_{j,2}^{(i,t)} \xi_i^{(t)} + C_{j,3}^{(i,t)} \xi_i^{(t,j)} + \xi_i^{(t)} \xi_i^{(t,j)}. \quad (50)$$

By using the above result, we obtain

$$\begin{aligned} & \frac{1}{2S} \sum_{j=1}^d \mathbb{E}_{\xi_i^{(t)}, \xi_i^{(t,j)}} \left[\left(\nabla_j \bar{\mathcal{L}}(\boldsymbol{\theta}^{(t)}) \right)^2 \right] \\ & \leq \frac{(1 - \tilde{p})^4}{2S} \|\nabla \mathcal{L}(\boldsymbol{\theta}^{(t)})\|^2 + \frac{(1 - \tilde{p})^2}{2SB} \sum_{j=1}^d \nabla_j \mathcal{L}(\boldsymbol{\theta}^{(t)}) \left(\sum_{i=1}^B C_{j,1}^{(i,t)} \right) + \frac{(1 - \tilde{p})^2}{SB} \sum_{j=1}^d \nabla_j \mathcal{L}(\boldsymbol{\theta}^{(t)}) \sum_{i=1}^B \mathbb{E}_{\xi_i^{(t)}} [\xi_i^{(t)}] \\ & \quad + \frac{(1 - \tilde{p})^2}{SB} \sum_{j=1}^d \nabla_j \mathcal{L}(\boldsymbol{\theta}^{(t)}) \sum_{i=1}^B \mathbb{E}_{\xi_i^{(t,j)}} [\xi_i^{(t,j)}] + \frac{(1 - \tilde{p})^2}{SB} \sum_{j=1}^d \nabla_j \mathcal{L}(\boldsymbol{\theta}^{(t)}) \sum_{i=1}^B \mathbb{E}_{\xi_i^{(t)}, \xi_i^{(t,j)}} [\xi_i^{(t)} \xi_i^{(t,j)}] \\ & \quad + \frac{d}{2SB^2} \left(\sum_{i=1}^B C_{j,1}^{(i,t)} \right)^2 + \frac{1}{2S} \sum_{j=1}^d \mathbb{E}_{\xi_i^{(t)}} [\xi_i^{(t)}] + \frac{1}{2S} \sum_{j=1}^d \mathbb{E}_{\xi_i^{(t,j)}} [\xi_i^{(t,j)}] + \frac{1}{2S} \sum_{j=1}^d \mathbb{E}_{\xi_i^{(t)}, \xi_i^{(t,j)}} [\xi_i^{(t)} \xi_i^{(t,j)}] \\ & \quad + \frac{1}{2SB^2} \sum_{j=1}^d \sum_{i=1}^B \mathbb{E}_{\xi_i^{(t)}} [(\xi_i^{(t)})^2] + \frac{1}{SB^2} \sum_{j=1}^d \sum_{i=1}^B \left(\mathbb{E}_{\xi_i^{(t)}, \xi_i^{(t,j)}} [\xi_i^{(t)} \xi_i^{(t,j)}] + \mathbb{E}_{\xi_i^{(t)}, \xi_i^{(t,j)}} [(\xi_i^{(t)})^2 \xi_i^{(t,j)}] \right) \\ & \quad + \frac{1}{2SB^2} \sum_{j=1}^d \sum_{i=1}^B \mathbb{E}_{\xi_i^{(t,j)}} [(\xi_i^{(t,j)})^2] + \frac{1}{SB^2} \sum_{j=1}^d \sum_{i=1}^B \mathbb{E}_{\xi_i^{(t)}, \xi_i^{(t,j)}} [\xi_i^{(t)} (\xi_i^{(t,j)})^2] + \\ & \quad + \frac{1}{2SB^2} \sum_{j=1}^d \sum_{i=1}^B \mathbb{E}_{\xi_i^{(t)}, \xi_i^{(t,j)}} [(\xi_i^{(t)})^2 (\xi_i^{(t,j)})^2] \\ & \leq \frac{(1 - \tilde{p})^4}{2S} \|\nabla \mathcal{L}(\boldsymbol{\theta}^{(t)})\|^2 + \frac{(1 - \tilde{p})^2 G}{2S} \max_{i,j} C_{j,1}^{(i,t)} + \frac{d}{2S} \max_{i,j} \left(C_{j,1}^{(i,t)} \right)^2 \\ & \quad + \frac{dC_{j,4,\max}^{(t)}}{2SB} + \frac{dC_{j,5,\max}^{(t,j)}}{2SB} + \frac{dC_{j,4,\max}^{(t)} C_{j,5,\max}^{(t,j)}}{2SB}. \end{aligned} \quad (51)$$

The first and second inequalities uses $C_{j,2}^{(i,t)} \leq 1$, $C_{j,3}^{(i,t)} \leq 1$, $\mathbb{E}[\xi_i^{(t)}] = 0$, $\mathbb{E}[\xi_i^{(t,j)}] = 0$, and $-G/d \leq \nabla_j \mathcal{L}(\boldsymbol{\theta}^{(t)}) \leq G/d$ supported by G -Lipschitz property. The term $C_{j,4,\max}^{(t)}$ refers to $C_{j,4,\max}^{(t)} = \max_i C_{j,4}^{(i,t)}$. Similarly, the term $C_{j,5,\max}^{(t,j)}$ refers to $C_{j,5,\max}^{(t,j)} = \max_i C_{j,5}^{(i,t)}$.

Since Theorem 5 indicates that

$$C_{j,4,\max}^{(t)} \leq \frac{(1 - \tilde{p}) \left(1 - 2\tilde{p} \frac{\text{Tr}(\Pi)}{D} \right)}{K} + \tilde{p} \frac{\text{Tr}(\Pi)}{DK} \leq \frac{2}{K},$$

and

$$C_{j,5,\max}^{(t,j)} \leq \frac{(1 - \tilde{p}) \left(1 - 2\tilde{p} \frac{\text{Tr}(\Pi)}{D} \right) (\hat{Y}_i^{(t,+j)} + \hat{Y}_i^{(t,-j)}) + 2\tilde{p} \frac{\text{Tr}(\Pi)}{D}}{K} \leq \frac{4}{K},$$

we obtain

$$\begin{aligned} & \frac{1}{2S} \sum_{j=1}^d \mathbb{E}_{\xi_i^{(t)}, \xi_i^{(t,j)}} \left[\left(\nabla_j \bar{\mathcal{L}}(\boldsymbol{\theta}^{(t)}) \right)^2 \right] \\ & \leq \frac{(1-\tilde{p})^4}{2S} \|\nabla \mathcal{L}(\boldsymbol{\theta}^{(t)})\|^2 + \frac{(1-\tilde{p})^2 G}{2S} \max_{i,j} C_{j,1}^{(i,t)} + \frac{d}{2S} \max_{i,j} \left(C_{j,1}^{(i,t)} \right)^2 + \frac{6dK+8d}{2SBK^2}. \end{aligned} \quad (52)$$

□

F Proof of Lemma 1

As shown in Theorem 5, the estimated gradient is center around the analytic gradients and is perturbed by the random noise $\boldsymbol{\varsigma}_i^{(t,j)}$ that follows the certain distribution. This behavior resembles a class of differentially private (DP) learning algorithm [55], where a certain type of noise is attached to the gradients to achieve the privacy and utility guarantees. Driven by the similarity between noisy QNN and DP models, here we investigate whether noisy QNN can be treated as a DP learning model.

The proof of Lemma 1 leverages the composition property of DP model as summarized below.

Proposition 1 (Composition property, [77]). *Suppose that a mechanism \mathcal{M} consists of a sequence of adaptive (ϵ, δ) -differentially private mechanisms $\mathcal{M}_1, \dots, \mathcal{M}_k$, where $\epsilon, \delta \geq 0$ $\mathcal{M}_i : \prod_{j=1}^{i-1} \mathcal{R}_j \times \mathbb{R}^D \rightarrow \mathcal{R}_i$. Then the mechanism \mathcal{M} satisfies $(\epsilon', k\delta + \delta')$ -differentially private with $\delta' \geq 0$ and*

$$\epsilon' = \sqrt{2k \ln(1/\delta')} \epsilon + k\epsilon(e^\epsilon - 1).$$

Proof of Lemma 1. Recall that, for noisy QNN, the estimated gradient of j -th parameter at t -th is

$$\nabla_j \bar{\mathcal{L}}_i(\boldsymbol{\theta}^{(t)}) = (\bar{Y}_i^{(t)} - Y_i) \left(\bar{Y}_i^{(t,+j)} - \bar{Y}_i^{(t,-j)} \right) + \lambda \boldsymbol{\theta}_j^{(t)}. \quad (53)$$

The composition property of DP as shown in Proposition 1 indicates that, if the mechanism $\mathcal{M}(\boldsymbol{\theta}_j^{(t)}, \mathcal{B}_i)$ that corresponds to the quantum circuits as shown in Fig. 1 (b), which is used to output $\bar{Y}_i^{(t)}$ and $\bar{Y}_i^{(t,\pm j)}$, satisfies DP property, then QNN with noisy gates also achieves the DP promise. Alternatively, to guarantee the DP property of QNN, we should prove that the random mechanism $\mathcal{M}(\boldsymbol{\theta}_j^{(t)}, \mathcal{B}_i)$ is a DP model. Without loss of generality, here we focus on the setting with $B = 1$ and $\mathcal{B}_i = \mathbf{z}$, since the privacy keeps unchanged when we vary B from 1 to N .

As explained in Theorem 5, the randomness of the mechanism $\mathcal{M}(\boldsymbol{\theta}_j^{(t)}, \mathbf{z})$ comes from the gate noise and finite number of measurements. For K quantum measurements, the possible values of sample mean $\bar{Y}_i^{(t)}$ is discrete, i.e., $\bar{Y}_i^{(t)} \in \{0, 1/K, \dots, 1\}$. By employing the properties of sample mean and Bernoulli random variables, the distribution of $\bar{Y}_i^{(t)}$ follows

$$\Pr(\bar{Y}_i^{(t)} = y) = \binom{K}{Ky} q^{Ky} (1-q)^{K-Ky}, \quad (54)$$

where $q = (1-\tilde{p}) \text{Tr}(\hat{Y}_i^{(t)}) + \tilde{p} \text{Tr}(\Pi)/D$.

In conjunction with Eqn. (54) and the definition of DP as formulated in Definition 2, the random algorithm $\mathcal{M}(\boldsymbol{\theta}_j^{(t)}, \mathbf{z})$ is DP if the following relation is satisfied, i.e.,

$$\frac{\Pr(\bar{Y}_i^{(t)} = y)}{\Pr(\bar{Y}_i'^{(t)} = y)} \leq e^{\epsilon'}, \quad (55)$$

where $\bar{Y}_i'^{(t)}$ refers to the sample mean of QNN given the tunable parameters $\boldsymbol{\theta}^{(t)}$ and the neighborhood dataset \mathbf{z}' . Combining Eqn. (54) and Eqn. (55), we obtain

$$\frac{\Pr(\bar{Y}_i^{(t)} = y)}{\Pr(\bar{Y}_i'^{(t)} = y)} = \frac{\binom{K}{Ky} q^{Ky} (1-q)^{K-Ky}}{\binom{K}{Ky} q'^{Ky} (1-q')^{K-Ky}} \leq \frac{q}{(q'(1-q'))^K}, \quad (56)$$

where $q' = (1 - \tilde{p}) \text{Tr}(\hat{Y}_i'^{(t)}) + \tilde{p} \text{Tr}(\Pi)/D$, the inequality uses the facts $q^{Ky}(1-q)^{K-Ky} \leq q$ and $q'^{Ky}(1-q')^{K-Ky} \geq (q'(1-q'))^K$.

By replacing q and q' with their explicit expressions, Eqn. (56) can be further simplified as

$$\frac{\Pr(\bar{Y}_i^{(t)} = y)}{\Pr(\bar{Y}_i'^{(t)} = y)} \leq \frac{(1 - \tilde{p}) + \tilde{p} \frac{\text{Tr}(\Pi)}{D}}{\left(\tilde{p}(1 - \tilde{p}) \left(1 - \frac{\text{Tr}(\Pi)}{D}\right)\right)^K}, \quad (57)$$

where the nominator employs $\text{Tr}(\hat{Y}_i^{(t)}) \leq 1$ and the denominator uses the fact $\tilde{p} \text{Tr}(\Pi)/D \leq q' \leq (1 - \tilde{p}) + \tilde{p} \text{Tr}(\Pi)/D$.

The result achieved in Eqn. (56) indicates that the mechanism $\mathcal{M}(\boldsymbol{\theta}_j^{(t)}, \mathbf{z})$ is a DP model, where

$$\frac{(1 - \tilde{p}) + \tilde{p} \frac{\text{Tr}(\Pi)}{D}}{\left(\tilde{p}(1 - \tilde{p}) \left(1 - \frac{\text{Tr}(\Pi)}{D}\right)\right)^K} = e^{\epsilon'} \Leftrightarrow \epsilon' = \ln \left(\frac{(1 - \tilde{p}) + \tilde{p} \frac{\text{Tr}(\Pi)}{D}}{\left(\tilde{p}(1 - \tilde{p}) \left(1 - \frac{\text{Tr}(\Pi)}{D}\right)\right)^K} \right). \quad (58)$$

We then use the DP property of the mechanism \mathcal{M} to derive the privacy parameter of QNN at the t -th iteration. By leveraging Proposition 1, the privacy parameters (ϵ'', δ'') of QNN to generate the estimated gradient of the j -th parameter $\nabla_j \mathcal{L}_i$ is

$$\epsilon'' = \sqrt{6 \ln(1/\delta'')} \epsilon' + 3\epsilon' (e^{\epsilon'} - 1). \quad (59)$$

Since the d trainable parameters of $\nabla \mathcal{L}_i$ are independent with each other, the definition of DP requests that, given two neighborhood input datasets \mathbf{z} and \mathbf{z}' , the following relation should be satisfied at the t -th iteration,

$$\prod_{j=1}^d \max_{r_j} \frac{\Pr(\mathcal{M}(\boldsymbol{\theta}_j^{(t)}, \mathbf{z}) = r_j)}{\Pr(\mathcal{M}(\boldsymbol{\theta}_j^{(t)}, \mathbf{z}') = r_j)} \leq e^{\epsilon'''} + \delta'''. \quad (60)$$

In conjunction with Eqn. (59) and Eqn. (60), we obtain

$$\epsilon''' = d\epsilon'', \text{ and } \delta''' = d\delta''. \quad (61)$$

Since the mechanism of QNN that is used to generate $\nabla \mathcal{L}_i$ at the t -th iteration satisfies the (ϵ''', δ''') -DP property, we can utilize Proposition 1 again to show QNN with T iterations is also an (ϵ, δ) -DP model, i.e.,

$$\begin{aligned} \epsilon &= \sqrt{2T \ln(1/\bar{\delta})} d\epsilon'' + T d\epsilon'' (e^{d\epsilon''} - 1) \\ &= \tilde{O} \left(\sqrt{Td} + Td \left(\frac{(1 - \tilde{p}) + \tilde{p} \frac{\text{Tr}(\Pi)}{D}}{\left(\tilde{p}(1 - \tilde{p}) \left(1 - \frac{\text{Tr}(\Pi)}{D}\right)\right)^K} \right)^d - Td \right). \end{aligned} \quad (62)$$

□

G Proof of Theorem 2

The key ingredients to achieve Theorems 2 are classical and quantum differentially private (DP) learning techniques [54, 84, 55, 57], and quantum PAC and QSQ learning models [60, 59]. The intuition to employ DP is as follows. The behavior of QNN with gate noise resembles DP learning, where a certain type of noise is injected into the learning model. Moreover, a recent study [59] proved that if a learning problem is quantum PAC learnable, then it is also quantum privately PAC (PPAC) learnable. Such an observation implies that if QNN with gate noise belongs to the DP learning model, then we can conclude the same learnability between noiseless QNN and QNN with gate noise.

To incorporate the achieved result of QNN with other quantum learning theory conclusions, the quantum examples discussed below concentrate on a specific type as formulated in Definition 4, which is broadly employed in quantum PAC learning and quantum statistical query (QSQ) learning. Note that the quantum encoding circuit $U_{\mathbf{x}}$ can efficiently prepare such quantum examples, as explained in Appendix B.

Definition 4 (Quantum example). *Let $c^* : \{0, 1\}^N \rightarrow \{0, 1\}$ be an unknown concept sampled from a known concept class \mathcal{C} . Denote the labeled examples as $(\mathbf{x}, c^*(\mathbf{x}))$, where \mathbf{x} is drawn from some unknown distribution \mathcal{D} . The quantum example is defined as $|\psi_{c^*}\rangle = \sum_{\mathbf{x} \in \{0, 1\}^N} \sqrt{\mathcal{D}(\mathbf{x})} |\mathbf{x}\rangle |c^*(\mathbf{x})\rangle$.*

Proof of Theorem 2. Given access to quantum examples $|\psi_{c^*}\rangle$ as formulated in Definition 4, we can leverage the results of quantum learning theory [60, 59] to exploit the learnability of noiseless QNN and QNN with noisy gates. In particular, the two studies [60, 59] proved quantum PAC = PAC and quantum PPAC = PPAC. Since a well known classical result [58] is PAC = PPAC, we obtain the following relationship in terms of sample complexity, i.e.,

$$\text{quantum PAC} = \text{quantum PPAC} = \text{PAC} = \text{PPAC} . \quad (63)$$

Eqn. (63) indicates that the learnable concept classes for non-private learning model and the DP learning model are same. Consequently, if a concept is PAC learnable by a QNN, then such a concept is also PAC learnable by a DP learning model, i.e., QNN with noisy gates, where its DP property has been proved in Lemma 1. \square

H Proof of Theorem 3

Theorem 3 quantifies the required query complexity of QNN with noisy gates to simulate one query of QSQ model. The definition of quantum statistical query (QSQ) learning model and its relevant theoretical results [59] as shown below.

Definition 5 (QSQ). *Let $\mathcal{C} \subseteq \{c : \{0, 1\}^N \rightarrow \{0, 1\}\}$ be a concept class and $\mathcal{D} : \{0, 1\}^N \rightarrow \{0, 1\}$ be a distribution. A quantum statistical query oracle $\text{Qstat}(\Pi, \tau)$ for some $c^* \in \mathcal{C}$ receives as inputs a tolerance $\tau \geq 0$ and an observable $\mathbb{M} \in (\mathbb{C}^2)^{\otimes N+1} \times (\mathbb{C}^2)^{\otimes N+1}$, and outputs a number α satisfying*

$$|\alpha - \langle \psi_{c^*} | \mathbb{M} | \psi_{c^*} \rangle| \leq \tau ,$$

where $\psi_{c^*} = \sum_{\mathbf{x} \in \{0, 1\}^N} \sqrt{\mathcal{D}(\mathbf{x})} |\mathbf{x}, c^*(\mathbf{x})\rangle$ refers to the quantum example.

Definition 6 (ε -learning). *Let $\mathcal{C} \subseteq \{c : \{0, 1\}^N \rightarrow \{0, 1\}\}$ be a concept class and $\mathcal{D} : \{0, 1\}^N \rightarrow \{0, 1\}$ be a distribution. We say that \mathcal{C} can be ε -learned in the QSQ model with Q queries, if there is an algorithm \mathcal{A} such that for every $c^* \in \mathcal{C}$, \mathcal{A} makes at most Q Qstat queries and outputs a hypothesis h satisfying $\Pr_{\mathbf{x} \sim \mathcal{D}}[h(\mathbf{x}) \neq c^*(\mathbf{x})] \leq \varepsilon$.*

The key technique to achieve Theorem 3 is the concentration inequality, which bounds the deviation of a random variable that corresponds to the output of QNN from a certain number. In particular, with treating the output α in QSQ model as the sample mean of QNN, the relation $\alpha - \langle \psi_{c^*} | \mathbb{M} | \psi_{c^*} \rangle \leq \tau$ evaluates how is the probability when the distance between the sample mean α and its expectation $\langle \psi_{c^*} | \mathbb{M} | \psi_{c^*} \rangle$ is within τ . Such a question can be effectively answered by using concentration inequality.

Lemma 9 (Modified from Lemma 4.2, 4.3, and 4.5 in [59]). *Let \mathcal{C} be the concept class of parities, k -juntas, or $\text{poly}(n)$ -sized DNFs (Disjunctive Normal Forms), then there exists a $\text{poly}(n)$ queries QSQ algorithm with tolerance $\tau = \tilde{O}(\varepsilon)$ that ε -learns \mathcal{C} under the uniform distribution. All of these concepts are computationally hard for SQ models.*

Proof of Theorem 3. Following the notations used in Definitions 4 and 5, supposed that the encoding circuits $U_{\mathbf{x}}$ prepares the quantum example $|\psi_{c^*}\rangle$ and the trainable unitary $U(\boldsymbol{\theta})$ is identity $\mathbb{I}_{2^{\otimes N+1}}$. Then, with applying the observable \mathbb{M} to the generated state of QNN, the expectation value of quantum measurements under the depolarization noise setting $\mathcal{N}_{\tilde{p}}$ follows $\tilde{\nu} = (1 - \tilde{p})\nu + \frac{1}{2^{N+1}}$ with $\nu = \langle \psi_{c^*} | \mathbb{M} | \psi_{c^*} \rangle$, supported by Lemma 7. The measurement outcome V_k is a random variable that satisfies $V_k \sim \text{Ber}(\tilde{\nu})$.

By the Chernoff-Hoeffding bound for real-valued variables, we obtain the relation between the sample mean $\frac{1}{K} \sum_{k=1}^K V_k$ with K measurements and the target result $\tilde{\nu}$, i.e.,

$$\Pr \left(\left| \frac{1}{K} \sum_{i=1}^K V_k - \tilde{\nu} \right| \geq \frac{\delta}{2} \right) \leq 2 \exp(-\delta^2 K/2). \quad (64)$$

Moreover, the distance between the target result ν and the shifted expectation values $\tilde{\nu}$ follows

$$|\nu - \tilde{\nu}| \leq \tilde{p}\nu + \frac{\text{Tr}(\mathbb{M})}{2^{N+1}}. \quad (65)$$

In conjunction with the above two equations, we obtain, with probability at least $1 - 2 \exp(-\delta^2 n/2)$

$$\left| \frac{1}{K} \sum_{k=1}^K V_k - \nu \right| = \left| \frac{1}{K} \sum_{k=1}^K V_k - \tilde{\nu} + \tilde{\nu} - \nu \right| \leq \tilde{p}\nu + \frac{\text{Tr}(\mathbb{M})}{2^{N+1}} + \frac{\delta}{2}. \quad (66)$$

Note that, to guarantee the term $\tilde{p}\nu + \frac{\text{Tr}(\mathbb{M})}{2^{N+1}} + \frac{\delta}{2}$ is upper bounded by τ , the parameter \tilde{p} should satisfy

$$\tilde{p} \leq \frac{\tau - \frac{\delta}{2} - \frac{\text{Tr}(\mathbb{M})}{2^{N+1}}}{\nu}. \quad (67)$$

Under the assumption of Eqn. (67), with setting $\delta \leq 2(\tau - \tilde{p}\nu - \frac{\text{Tr}(\mathbb{M})}{2^{N+1}})$, QNN simulates the QSQ model as formulated in Definition 5, i.e.,

$$\left| \frac{1}{K} \sum_{k=1}^K V_k - \nu \right| \leq \tau.$$

Under this setting, the relation between the number of measurements K and the successful probability b obeys

$$\Pr \left(\left| \frac{1}{K} \sum_{k=1}^K V_k - \tilde{\nu} \right| \geq \left(\tau - \tilde{p}\nu - \frac{\text{Tr}(\mathbb{M})}{2^{N+1}} \right) \right) \leq 2 \exp \left(-2 \left(\tau - \tilde{p}\nu - \frac{\text{Tr}(\mathbb{M})}{2^{N+1}} \right)^2 K \right) = b. \quad (68)$$

After simplification, we conclude that, when $\tilde{p} \leq \frac{\tau - \frac{\delta}{2} - \frac{\text{Tr}(\mathbb{M})}{2^{N+1}}}{\nu}$, with the successful probability at least $1 - b$, the required number of measurements to attain $\left| \frac{1}{K} \sum_{k=1}^K V_k - \nu \right| \leq \tau$ is

$$K = \frac{\ln\left(\frac{2}{b}\right)}{2\left(\tau - \tilde{p}\nu - \frac{\text{Tr}(\mathbb{M})}{2^{N+1}}\right)^2}. \quad (69)$$

□

I Generalization the results to more general quantum channels

In this section, we generalize the achieved results in main text from the depolarization channel to a more general channel \mathcal{E}_{p_1} , i.e.,

$$\mathcal{E}_{p_1}(\rho) = (1 - p_1)\rho + p_2\kappa + p_3\mathbb{I}_D/D, \quad (70)$$

where $\rho, \kappa \in \mathbb{C}^{D \times D}$, κ is a mixed state that can either be correlated or uncorrelated with ρ , and $p_2 + p_3 = p_1$ with $p_1, p_2 \geq 0$ and $p_3 > 0$. It is worth noting that the quantum channel \mathcal{E}_{p_1} is sufficiently universal, which covers most Pauli channels associated with the depolarization channel [38, 27].

The outline of this section is as follows. In Subsection I.1, we discuss the utility bounds of QNN under the general channel setting. Then, in Subsection I.2, we analyze the DP property of QNN when it is perturbed by the general channel. Last, in Subsection I.3, we quantify the learnability of QNN under the general channel setting from the perspective of sample complexity.

I.1 Utility bounds of QNN

Analogous to the depolarization channel setting, we first simplify the noisy QNN model to ease analysis. Specifically, after applying \mathcal{E}_{p_1} to each circuit depth, the generated state follows

$$\begin{aligned} & \mathcal{E}_{p_1}(U_L(\boldsymbol{\theta}) \dots U_2(\boldsymbol{\theta}) \mathcal{E}_{p_1}(U_1(\boldsymbol{\theta}) \rho U_1(\boldsymbol{\theta})^\dagger) U_2(\boldsymbol{\theta})^\dagger \dots U_L(\boldsymbol{\theta})^\dagger) \\ & = (1 - p_1)^{L_Q} (U(\boldsymbol{\theta}) U_{\mathbf{x}}) \rho (U(\boldsymbol{\theta}) U_{\mathbf{x}})^\dagger + p'_2 \kappa + p_3 \frac{L_Q}{D} \mathbb{I}_D, \end{aligned} \quad (71)$$

where $(1 - p_1)^{L_Q} + p'_2 + p_3 \frac{L_Q}{D} = 1$, and κ is a mixed state that can either be correlated or uncorrelated with $(U(\boldsymbol{\theta}) U_{\mathbf{x}}) \rho (U(\boldsymbol{\theta}) U_{\mathbf{x}})^\dagger$. Without confusion, we set $\tilde{p} = 1 - (1 - p_1)^{L_Q}$.

We now employ the simplified model, i.e., the right hand side of Eqn. (71), to establish the relation between the estimated gradients $\nabla_j \bar{\mathcal{L}}_i(\boldsymbol{\theta}^{(t)})$ and the analytic gradients $\nabla_j \mathcal{L}_i(\boldsymbol{\theta}^{(t)})$. Recall that

$$\nabla_j \bar{\mathcal{L}}_i(\boldsymbol{\theta}^{(t)}) = (\bar{Y}_i^{(t)} - Y_i) \left(\bar{Y}_i^{(t, +j)} - \bar{Y}_i^{(t, -j)} \right) + \lambda \boldsymbol{\theta}_j^{(t)},$$

where $\bar{Y}_i^{(t)} = \sum_{k=1}^K V_k^{(t)} / K$ and $\bar{Y}_i^{(t, \pm j)} = \sum_{k=1}^K V_k^{(t, \pm j)} / K$ refer to the sample means when feeding $\boldsymbol{\theta}^{(t)}$ and $\boldsymbol{\theta}^{(t, \pm j)}$ into the trainable circuit. As with depolarization channel, the sample mean $\bar{Y}_i^{(t)}$ or $\bar{Y}_i^{(t, \pm j)}$ is a random variable follows certain distribution. In particular, following the notations used in Theorem 5, the mean and variance of $\bar{Y}_i^{(t)}$ follows

$$\begin{cases} \nu^{(t)} = (1 - \tilde{p}) \hat{Y}_i^{(t)} + p'_2 \text{Tr}(\Pi \kappa^{(t)}) + \frac{p_3}{2} \frac{L_Q}{D}, \\ \sigma^{(t)} = -\frac{\left((1 - \tilde{p}) \hat{Y}_i^{(t)} + p'_2 \text{Tr}(\Pi \kappa^{(t)}) \right)^2}{K} + \frac{(1 - p_3) \left((1 - \tilde{p}) \hat{Y}_i^{(t)} + p'_2 \text{Tr}(\Pi \kappa^{(t)}) \right)}{K} + \frac{p_3}{2} \frac{L_Q}{D} - \frac{(p_3 \frac{L_Q}{D})^2}{4}. \end{cases}$$

Similarly, the mean and variance of $\bar{Y}_i^{(t,\pm j)}$ follows

$$\begin{cases} \nu^{(t,\pm j)} = (1 - \tilde{p})\hat{Y}_i^{(t,\pm j)} + p'_2 \text{Tr}(\Pi\kappa^{(t,\pm j)}) + \frac{p_3^{LQ}}{2}, \\ \sigma^{(t,\pm j)} = -\frac{\left((1-\tilde{p})\hat{Y}_i^{(t,\pm j)} + p'_2 \text{Tr}(\Pi\kappa^{(t,\pm j)})\right)^2}{K} + \frac{(1-p_3^{LQ})\left((1-\tilde{p})\hat{Y}_i^{(t,\pm j)} + p'_2 \text{Tr}(\Pi\kappa^{(t,\pm j)})\right)}{K} + \frac{p_3^{LQ}}{2} - \frac{(p_3^{LQ})^2}{4}. \end{cases}$$

By expanding the sample means using their explicit forms as shown above, we obtain the relation between the estimated and analytic gradients, i.e.,

$$\nabla_j \bar{\mathcal{L}}_i(\boldsymbol{\theta}^{(t)}) = (1 - \tilde{p})^2 \nabla_j \mathcal{L}_i(\boldsymbol{\theta}^{(t)}) + C_{j,1}^{(i,t)} + \boldsymbol{\varsigma}_i^{(t,j)}, \quad (72)$$

where $\boldsymbol{\varsigma}_i^{t,j} = C_{j,2}^{(i,t)} \xi_i^{(t)} + C_{j,3}^{(i,t)} \xi_i^{(t,j)} + \xi_i^{(t)} \xi_i^{(t,j)}$, and two random variables $\xi_i^{(t)}$ and $\xi_i^{(t,j)}$ have zero means and their variances are $C_{j,4}^{(i,t)}$ and $C_{j,5}^{(i,t)}$, respectively. The explicit formula of the five parameters $\{C_{j,a}^{(i,t)}\}_{a=1}^5$ is

$$\begin{cases} C_{j,1}^{(i,t)} = \left(p'_2 \text{Tr}(\Pi\kappa^{(t)}) + \frac{p_3^{LQ}}{2} - \tilde{p}Y_i \right) (1 - \tilde{p})(\hat{Y}_i^{(t,+j)} - \hat{Y}_i^{(t,-j)}) \\ \quad + p'_2(1 - \tilde{p})(\hat{Y}_i^{(t)} - Y_i)(\text{Tr}(\Pi\kappa^{(t,+j)}) - \text{Tr}(\Pi\kappa^{(t,-j)})) \\ \quad + \left(p'_2 \text{Tr}(\Pi\kappa^{(t)}) + \frac{p_3^{LQ}}{2} - \tilde{p}Y_i \right) (\text{Tr}(\Pi\kappa^{(t,+j)}) - \text{Tr}(\Pi\kappa^{(t,-j)})) + (1 - (1 - \tilde{p})^2)\lambda\boldsymbol{\theta}_j^{(t)}, \\ C_{j,2}^{(i,t)} = \left((1 - \tilde{p})(\hat{Y}_i^{(t,+j)} - \hat{Y}_i^{(t,-j)}) + p'_2(\text{Tr}(\Pi\kappa^{(t,+j)}) - \text{Tr}(\Pi\kappa^{(t,-j)})) \right), \\ C_{j,3}^{(i,t)} = \left((1 - \tilde{p})(\hat{Y}_i^{(t)} - Y_i) + \left(p'_2 \text{Tr}(\Pi\kappa^{(t)}) + \frac{p_3^{LQ}}{2} - \tilde{p}Y_i \right) \right), \\ C_{j,4}^{(i,t)} = -\frac{\left((1-\tilde{p})\hat{Y}_i^{(t)} + p'_2 \text{Tr}(\Pi\kappa^{(t)})\right)^2}{K} + \frac{(1-p_3^{LQ})\left((1-\tilde{p})\hat{Y}_i^{(t)} + p'_2 \text{Tr}(\Pi\kappa^{(t)})\right)}{K} + \frac{p_3^{LQ}}{2K} - \frac{(p_3^{LQ})^2}{4K}, \\ C_{j,5}^{(i,t)} = -\frac{\left((1-\tilde{p})\hat{Y}_i^{(t,+j)} + p'_2 \text{Tr}(\Pi\kappa^{(t,+j)})\right)^2}{K} - \frac{\left((1-\tilde{p})\hat{Y}_i^{(t,-j)} + p'_2 \text{Tr}(\Pi\kappa^{(t,-j)})\right)^2}{K} \\ \quad + \frac{(1-p_3^{LQ})\left((1-\tilde{p})(\hat{Y}_i^{(t,+j)} - \hat{Y}_i^{(t,-j)}) + p'_2(\text{Tr}(\Pi\kappa^{(t,+j)}) - \text{Tr}(\Pi\kappa^{(t,-j)}))\right)}{K} + \frac{p_3^{LQ}}{K} - \frac{(p_3^{LQ})^2}{2K}. \end{cases}$$

We next use the relation between the estimated and analytic gradients to separately quantify the utility bounds R_1 and R_2 of QNN under the noisy channel \mathcal{E}_{p_1} setting.

Utility bound R_1 . As with Eqn.(41), with taking expectation over $\xi_i^{(t)}$ and $\xi_i^{(t,j)}$, we obtain

$$\begin{aligned} & \mathbb{E}_{\xi_i^{(t)}, \xi_i^{(t,j)}}[\mathcal{L}(\boldsymbol{\theta}^{(t+1)}) - \mathcal{L}(\boldsymbol{\theta}^{(t)})] \\ & \leq -\frac{1}{S}(1 - \tilde{p})^2 \|\nabla \mathcal{L}(\boldsymbol{\theta}^{(t)})\|^2 + \frac{G}{2S} \left(\frac{1}{B} \sum_{i=1}^B C_{j,1}^{(i,t)} \right) + \frac{1}{2S} \sum_{j=1}^d \mathbb{E}_{\xi_i^{(t)}, \xi_i^{(t,j)}} \left[\left(\nabla_j \bar{\mathcal{L}}(\boldsymbol{\theta}^{(t)}) \right)^2 \right], \quad (73) \end{aligned}$$

where the inequality employs $\mathbb{E}[\xi_i^{(t)}] = 0$, $\mathbb{E}[\xi_i^{(t,j)}] = 0$, and $-G/d \leq \nabla_j \mathcal{L}(\boldsymbol{\theta}^{(t)}) \leq G/d$.

For the term $\frac{1}{2S} \sum_{j=1}^d \mathbb{E}_{\xi_i^{(t)}, \xi_i^{(t,j)}} \left[\left(\nabla_j \bar{\mathcal{L}}(\boldsymbol{\theta}^{(t)}) \right)^2 \right]$ in the above equation, its upper bound satisfies

$$\begin{aligned} \frac{1}{2S} \sum_{j=1}^d \mathbb{E}_{\xi_i^{(t)}, \xi_i^{(t,j)}} \left[\left(\nabla_j \bar{\mathcal{L}}(\boldsymbol{\theta}^{(t)}) \right)^2 \right] & \leq \frac{(1 - \tilde{p})^4}{2S} \|\nabla \mathcal{L}(\boldsymbol{\theta}^{(t)})\|^2 + \frac{(1 - \tilde{p})^2 G}{2SB} \sum_{i=1}^B C_1^{(i,t)} \\ & \quad + \frac{d}{2SB^2} \left(\sum_{i=1}^B C_1^{(i,t)} \right)^2 + d \frac{\sigma_{\max}^{(t)} + \sigma_{\max}^{(t,j)} + \sigma_{\max}^{(t)} \sigma_{\max}^{(t,j)}}{SB}, \quad (74) \end{aligned}$$

where the first and second inequalities uses $C_2^{(i,t)} \leq 2$, $C_3^{(i,t)} \leq 2$, $\mathbb{E}[\xi_i^{(t)}] = 0$, and $\mathbb{E}[\xi_i^{(t,j)}] = 0$. The term $\sigma_{\max}^{(t)}$ refers to $\sigma_{\max}^{(t)} = \max_i \sigma_i^{(t)} \leq 3/K$. Similarly, the term $\sigma_{\max}^{(t,j)}$ refers to $\sigma_{\max}^{(t,j)} = \max_i \sigma_i^{(t,+j)} + \sigma_i^{(t,-j)} \leq 3/K$.

In conjunction with the above two equations, we achieve

$$\begin{aligned} & \mathbb{E}_{\xi_i^{(t)}, \xi_i^{(t,j)}} [\mathcal{L}(\boldsymbol{\theta}^{(t+1)}) - \mathcal{L}(\boldsymbol{\theta}^{(t)})] \\ & \leq -\frac{1}{2S}(1-\tilde{p})^2 \|\nabla \mathcal{L}(\boldsymbol{\theta}^{(t)})\|^2 + \frac{(2G+d)(5+3(1-(1-\tilde{p})^2)\lambda\pi)}{2S} + \frac{6dK+9d}{SBK^2}, \end{aligned} \quad (75)$$

where the inequality uses $C_{j,1}^{(i,t)} \leq 5+3(1-(1-\tilde{p})^2)\lambda\pi$.

After rewriting and taking induction, we have

$$\|\nabla \mathcal{L}(\boldsymbol{\theta}^{(t)})\|^2 \leq 2S \frac{1+9\lambda d}{T(1-\tilde{p})^2} + \frac{(2G+d)(5+3(1-(1-\tilde{p})^2)\lambda\pi)}{(1-\tilde{p})^2} + \frac{12dK+18d}{(1-\tilde{p})^2 BK^2}. \quad (76)$$

With setting $T \rightarrow \infty$, we achieve the utility bound R_1 , i.e.,

$$R_1 \leq \tilde{O} \left(\frac{1}{(1-\tilde{p})^2}, d, \frac{1}{BK} \right). \quad (77)$$

Utility bound R_2 . With combining Eqn. (75) and PL condition, we obtain

$$\begin{aligned} & \mathbb{E}_{\xi_i^{(t)}, \xi_i^{(t,j)}} [\mathcal{L}(\boldsymbol{\theta}^{(t+1)}) - \mathcal{L}(\boldsymbol{\theta}^{(t)})] \\ & \leq -\frac{\mu(1-\tilde{p})^2}{S} (\mathcal{L}(\boldsymbol{\theta}^{(t)}) - \mathcal{L}^*) + \frac{(2G+d)(5+3(1-(1-\tilde{p})^2)\lambda\pi)}{2S} + \frac{6dK+9d}{SBK^2}. \end{aligned} \quad (78)$$

After rewriting and induction, we have

$$\mathbb{E}_{\xi^{(t)}} [\mathcal{L}(\boldsymbol{\theta}^{(T)})] - \mathcal{L}^* \leq 15\lambda d \exp \left(-\frac{\mu(1-\tilde{p})^2 T}{S} \right) + T \frac{(2G+d)(5+3(1-(1-\tilde{p})^2)\lambda\pi)}{2S} + T \frac{6dK+9d}{SBK^2}. \quad (79)$$

With setting $T = O \left(\frac{S}{\mu(1-\tilde{p})^2} \ln \left(\frac{30\lambda d SBK^2}{(2G+d)(5+3(1-(1-\tilde{p})^2)\lambda\pi) BK^2 + 12dK + 18d} \right) \right)$, the utility bound is

$$R_2 \leq O \left(\frac{1}{(1-\tilde{p})^2}, \frac{1}{SBK^2}, d \right). \quad (80)$$

I.2 Differential privacy of QNN

Analogous to QNN with depolarization noise, the DP property of QNN perturbed by the channel \mathcal{E}_{p_1} is determined by the DP property of the mechanism $\mathcal{M}(\boldsymbol{\theta}_j, \mathbf{z})$ to output $Y_i^{(t)}$, supported by the composition property as shown in Proposition 1.

As discussed in Subsection I.1, the distribution of sample mean $Y_i^{(t)}$ is similar to the depolarization case, where the only difference is that the values of mean and variance of the random variable are varied. In other words, the mechanism $\mathcal{M}(\boldsymbol{\theta}_j, \mathbf{z})$ used in QNN perturbed by the channel \mathcal{E}_{p_1} is also DP. This observation promises that QNN perturbed by the channel \mathcal{E}_{p_1} is a DP learning model.

I.3 Learnability of QNN

The generalization of Theorem 2. Celebrated by the DP property of QNN as discussed above, we can effectively generalize the result of Theorem 2 to the noisy channel \mathcal{E}_{p_1} setting, i.e., if QNN with noiseless gates PAC learns a concept, then QNN perturbed by the noisy channel \mathcal{E}_{p_1} can also learn such a concept using polynomial samples.

The generalization of Theorem 3. Analogous to the depolarization noise setting, the distance between the target result $\nu = \text{Tr}(\mathbb{M}(U(\boldsymbol{\theta})U_{\mathbf{x}})\rho(U(\boldsymbol{\theta})U_{\mathbf{x}})^\dagger)$ and the shifted expectation value $\tilde{\nu} = (1 - \tilde{p})\nu + p'_2 \text{Tr}(\mathbb{M}\kappa) + p'_3 \text{Tr}(\mathbb{M})/D$ of QNN under the noisy channel \mathcal{E}_{p_1} follows $|\nu - \tilde{\nu}| \leq \tilde{p}\nu + p'_2 + p'_3 \text{Tr}(\mathbb{M})/D$. Then by employing Chernoff-Hoeffding bound, we achieve, with probability at least $1 - 2\exp(-\delta^2 n/2)$,

$$\left| \frac{1}{k} \sum_{k=1}^K V_k - \nu \right| \leq \tilde{p}\nu + p'_2 + \frac{p'_3}{D} + \frac{\delta}{2}.$$

Assuming $\tilde{p} \leq \frac{\tau - p'_2 - \frac{p'_3}{D} - \frac{\delta}{2}}{\nu}$, with setting $\delta = 2(\tau - \tilde{p}\nu - p'_2 - p'_3/D)$, the relation between the number of measurements K and the successful probability b obeys

$$\Pr \left(\left| \frac{1}{K} \sum_{k=1}^K V_k - \tilde{\nu} \right| \geq \left(\tau - \tilde{p}\nu - p'_2 - \frac{p'_3}{D} \right) \right) \leq 2 \exp \left(-2 \left(\tau - \tilde{p}\nu - p'_2 - \frac{p'_3}{D} \right)^2 K \right) = b. \quad (81)$$

After simplification, we conclude that, when $\tilde{p} \leq \frac{\tau - p'_2 - \frac{p'_3}{D} - \frac{\delta}{2}}{\nu}$, with the successful probability at least $1 - b$, the required number of measurements to attain $\left| \frac{1}{K} \sum_{k=1}^K V_k - \nu \right| \leq \tau$ is

$$K = \frac{\ln \left(\frac{2}{b} \right)}{2 \left(\tau - \tilde{p}\nu - p'_2 - \frac{p'_3}{D} \right)^2}. \quad (82)$$