Q-ShiftDP: A Differentially Private Parameter-Shift Rule for Quantum Machine Learning

Hoang M. Ngo
University of Florida
hoang.ngo@ufl.edu

Nhat Hoang-Xuan
University of Florida
hoangx@ufl.edu

Quan Nguyen
University of Florida
quan.nguyen1@ufl.edu

Nguyen Do
University of Florida
nguyen.do@ufl.edu

Incheol Shin
Pukyong National University
icshin@pknu.ac.kr

My T. Thai*
University of Florida
mythai@cise.ufl.edu

Abstract

Quantum Machine Learning (QML) promises significant computational advantages, but preserving training data privacy remains challenging. Classical approaches like differentially private stochastic gradient descent (DP-SGD) add noise to gradients but fail to exploit the unique properties of quantum gradient estimation. In this work, we introduce the Differentially Private Parameter-Shift Rule (Q-ShiftDP), the first privacy mechanism tailored to QML. By leveraging the inherent boundedness and stochasticity of quantum gradients computed via the parameter-shift rule, Q-ShiftDP enables tighter sensitivity analysis and reduces noise requirements. We combine carefully calibrated Gaussian noise with intrinsic quantum noise to provide formal privacy and utility guarantees, and show that harnessing quantum noise further improves the privacy–utility trade-off. Experiments on benchmark datasets demonstrate that Q-ShiftDP consistently outperforms classical DP methods in QML.

1 INTRODUCTION

Quantum Machine Learning (QML) has emerged as a promising avenue for advancing machine learning by harnessing unique quantum phenomena such as superposition and entanglement. These quantum properties enable QML to learn high-dimensional and complex patterns that are intractable for classical machine learning (ML) algorithms Lloyd et al. (2013); Cai et al. (2015); Schuld and Killoran (2019). As a result, QML potentially achieve advantages in learning tasks such as regression, classification and reinforcement learning.

Privacy remains a fundamental concern for all learning models, both classical and quantum. In classical ML, Differential Privacy (DP) has been established as the standard framework for providing rigorous privacy guarantees Dwork (2006a). Incorporating differential privacy (DP) into machine learning models typically involves injecting carefully calibrated noise into the learning process. Among recent related work, gradient perturbation has been shown to provide the best trade-off between privacy and utility in deep learning models Jarin and Eshete (2022). Notable ones are DP-SGD Abadi et al. (2016a), adaptive clipping Thakkar et al. (2019), and personalized privacy budgets Boenisch et al. (2023).

Recent studies have extended the concept of DP to the quantum domain, giving rise to Quantum Differential Privacy (QDP) Zhou and Ying (2017); Hirche et al. (2023). This line of work is particularly important for QML, as QML models are designed to capture highly expressive, high-dimensional representations that may inadvertently encode sensitive information about individual training samples Rajawat et al. (2023); Thakkar et al. (2024); Marengo and Santamato (2025); Bukkarayasamudram et al. (2025). Thus, rigorously enforcing DP in QML is a practical prerequisite for deploying quantum learning algorithms in real-world, data-sensitive applications. However, as in classical settings, this approach can suffer from sig-

Proceedings of the 29th International Conference on Artificial Intelligence and Statistics (AISTATS) 2026, Tangier, Morocco. PMLR: Volume 300. Copyright 2026 by the author(s). *Corresponding author.

nificant utility degradation. There are two key differences in quantum gradients that make direct application of DP-SGD inefficient in QML. First, quantum gradients are estimated statistically via the parameter-shift rule Wierichs et al. (2022), rather than being exactly computed. This inherent stochasticity can itself serve as a source of privacy, thereby reducing the need for additional Gaussian noise. Second, unlike classical gradients which are unbounded, quantum gradients are naturally bounded due to the properties of quantum operators and observables. Studying this bound allows for precise sensitivity analysis and avoids unnecessary gradient clipping, which can reduce model performance.

In this work, we propose Q-ShiftDP, a Differentially Private Parameter-Shift Rule (PSR) mechanism designed to protect the privacy of QML models. To our knowledge, this is the first approach that explicitly exploits the unique characteristics of quantum gradients computed via PSR to efficiently protect the training process in QML. Specifically, we provide a rigorous theoretical analysis of the ℓ_2-sensitivity of quantum gradients. Based on this, we establish both privacy and utility guarantees for Q-ShiftDP by combining the intrinsic variance of estimated quantum gradients with carefully calibrated Gaussian noise. Furthermore, we demonstrate how the privacy-utility tradeoff of Q-ShiftDP can be further improved in two key ways: first, by treating physical depolarizing noise as a privacy-enhancing resource, and second, by introducing an adaptive technique to tailor the amount of artificial noise based on empirically estimated per-sample variance. To assess the effectiveness of Q-ShiftDP, we conduct experiments on three benchmark datasets: Bars & Stripes, Binary Blobs and Downscaled MNIST Bowles et al. (2024); Recio-Armengol et al. (2025). The empirical results demonstrate that Q-ShiftDP consistently achieves higher utility than input-perturbation based DP mechanisms in QML.

2 BACKGROUNDS

In this section, we first review quantum states, operations, and measurements with an emphasis on expectation values and shot noise. We then describe variational quantum machine learning (VQML) and the parameter-shift rule (PSR) for gradient estimation. Finally, we summarize differential privacy and classical DP-SGD, and illustrate why a direct transfer to QML is suboptimal. Additional material and related works appear in Appendices A and C.

2.1 Fundamentals of Quantum Information

Quantum States and Operations. Quantum computing systems are built upon qubits which are the fundamental units of quantum information. The state of a qubit can exist in the superposition of $|0\rangle$ and $|1\rangle$. For example, a single-qubit pure state can be written as $|\psi\rangle = a\,|0\rangle + b\,|1\rangle$ where $a, b \in \mathbb{C}$ and the normalization condition $|a|^2 + |b|^2 = 1$ holds. Multi-qubit systems of n qubits occupy a 2^n-dimensional Hilbert space formed by the tensor product of individual states.

Quantum Measurement and Observables. A measurement in a quantum system extracts classical information from a quantum state $|\psi\rangle$ by probing a physical property, such as energy, position, or spin. Each such property, known as an *observable*, is represented mathematically by a Hermitian operator O.

The possible outcomes of measuring this observable are its eigenvalues $\{\lambda_i\}$. Based on spectral theorem Moretti (2020), an observable O can be decomposed in terms of its projectors as $O = \sum_i \lambda_i P_i$ where P_i is the projection operator onto the eigenspace corresponding to λ_i. These projection operators satisfy $\sum_i P_i = I$. The expectation value of the observable O, denoted as $\langle O\rangle$ represents the statistical average outcome of many identical measurements performed on identically prepared systems. This average is calculated by summing each possible outcome weighted by its probability:

$$\begin{aligned}\langle O\rangle &= \sum_i \lambda_i \Pr(\lambda_i) = \sum_i \lambda_i \langle\psi| P_i |\psi\rangle \\ &= \langle\psi| \left(\sum_i \lambda_i P_i\right) |\psi\rangle = \langle\psi| O |\psi\rangle .\end{aligned} \tag{1}$$

Importantly, $\langle O\rangle$ is always bounded by the minimum and maximum eigenvalues of O, a property that later enables tight sensitivity analysis.

Finite-Shot Estimation and Shot Noise.

In practice, the expectation value $\langle O\rangle$ cannot be accessed exactly. Instead, it is estimated using N_s independent measurement *shots*. A single shot yields one of its eigenvalues λ_i. To obtain statistical information, we perform N_s independent shots on identically prepared copies of the state $|\psi\rangle$. Each shot produces an outcome $\hat{a} \in \{\lambda_i\}$ with probability $\Pr(\hat{a} = \lambda_i) = \langle\psi| P_i |\psi\rangle$. After N_s shots, we collect N_s independent outcomes $\{\hat{a}_1, \ldots, \hat{a}_{N_s}\}$. The empirical estimator of the expectation value is then $\bar{A}_{N_s} = \frac{1}{N_s}\sum_{i=1}^{N_s} \hat{a}_i$. The variance of this estimator decreases with the number of shots and is given by

$$\mathrm{Var}[\bar{A}_{N_s}] = \frac{\sigma^2_{\text{shot}}}{N_s} \tag{2}$$

where σ^2_{shot} is the variance of a single-shot outcome distribution. This intrinsic statistical noise, known as *shot noise*, plays a central role in both optimization dynamics and privacy guarantees in QML.

2.2 Variational Quantum Machine Learning

Variational Quantum Machine Learning (VQML) is a quantum-classical algorithm in which a quantum circuit with tunable parameters is optimized using classical techniques. This approach is well-suited for noisy intermediate-scale quantum (NISQ) devices, as it harnesses the expressive power of quantum computing to learn complex patterns while relying on efficient classical optimizers to overcome the limitations of current quantum hardware. In this section, we first introduce Parameterized Quantum Circuits (PQCs), which serve as the cornerstone of VQML. We then present a technique for computing gradients of PQCs, called Parameter-Shift Rule.

Parameterized Quantum Circuit (PQC) Models. A VQML model is trained by defining a cost function $C(\boldsymbol{\theta})$ that depends on the parameters $\boldsymbol{\theta}$ of a PQC, and then minimizing it through iterative updates. A PQC typically consists of three components: (i) an encoding that embeds classical data x into a quantum state, (ii) a variational ansatz $U(\boldsymbol{\theta})$ with tunable parameters $\boldsymbol{\theta}$, and (iii) measurement operators $\{\hat{O}_i\}$ that extract classical information from the quantum state. The variational ansatz can be expressed as a unitary transformation $U(\boldsymbol{\theta}) = \prod_{k=1}^{K} e^{i\theta_k G_k}$ where $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_K)$ are trainable parameters and $\{G_k\}$ are Hermitian generators. Given an input x encoded into a quantum state $|\psi_0(x)\rangle$, the PQC produces:

$$|\psi(\boldsymbol{\theta}, x)\rangle = \hat{U}(\boldsymbol{\theta})\,|\psi_0(x)\rangle \tag{3}$$

Finally, the cost function is computed by measuring an observable $\hat{O}_y$ associated with label y, defined as:

$$C(\boldsymbol{\theta}) = \mathbb{E}_{(x,y)\sim S}\left[\langle\psi(\boldsymbol{\theta}, x)|\hat{O}_y|\psi(\boldsymbol{\theta}, x)\rangle\right] \tag{4}$$

Gradient Evaluation via Parameter-Shift Rule. Computing the gradients of $C(\boldsymbol{\theta})$ is essential for training VQML models. However, unlike classical ML, direct backpropagation is not applicable to PQCs because quantum measurements are inherently probabilistic and intermediate quantum states cannot be accessed or copied due to the no-cloning theorem. Instead, gradients are estimated using the a technique, called parameter-shift rule Li et al. (2017); Mitarai et al. (2018); Wierichs et al. (2022).

Consider a variational ansatz $U(\boldsymbol{\theta}) = \prod_{k=1}^{K} e^{i\theta_k G_k}$ where G_k has only two eigenvalues. Let the difference between two eigenvalues be Ω_k. We can calculate the deviation of $f(\boldsymbol{\theta}) = \langle\psi(\boldsymbol{\theta}, x)|\hat{O}_y|\psi(\boldsymbol{\theta}, x)\rangle$ w.r.t θ_k by parameter-shift rule as

$$\frac{\partial f}{\partial \theta_k} = \frac{\Omega_k}{2\sin(\Omega_k s)}\Big[f(\theta_k + s) - f(\theta_k - s)\Big] \tag{5}$$

where s is freely chosen from $(0, \pi)$. A common choice in practice is $s = \pi/(2\Omega_k)$, which simplifies the formula because $\sin(\Omega_k s) = 1$. This approach provides an unbiased gradient evaluation while requiring only two evaluations of the quantum circuit per parameter. $f(\theta_k + s)$ and $f(\theta_k - s)$ are practically implemented by a positive-shifted and a negative-shifted circuits.

2.3 Differential Privacy

Foundations of Differential Privacy. Differential Privacy (DP) provides a rigorous framework that limits how much influence a single data point can have on the output of an algorithm Dwork (2006a).

Definition 1 ((ε, δ)-Differential Privacy)**.** *A randomized mechanism $\mathcal{M} : \mathcal{D} \to \mathcal{R}$ satisfies (ε, δ)-differential privacy if for any two adjacient datasets D_1 and D_2 that differs by a single element, and for any subset of outputs $S \subseteq \mathcal{R}$, the following inequality holds:*

$$\Pr[\mathcal{M}(D_1) \in S] \leq e^{\varepsilon}\Pr[\mathcal{M}(D_2) \in S] + \delta \tag{6}$$

In this definition, $\varepsilon \geq 0$ quantifies the maximum privacy loss. Specifically, it controls how much the output distribution may change when one data point is modified. In addition, the parameter $\delta \in [0, 1)$ allows for a small probability of failure in this guarantee. Smaller values of ε and δ lead to stronger privacy protection.

Differentially Private Stochastic Gradient Descent (DP-SGD). The dominant approach for private training in classical ML is DP-SGD. Differentially Private SGD (DP-SGD) Abadi et al. (2016a). DP-SGD enforces privacy by modifying the standard SGD update through two key mechanisms:

1. *Gradient clipping:* Each per-sample gradient is clipped to a fixed norm bound C. This ensures that sensitivity of the gradients is bounded.

2. *Noise addition:* After clipping, Gaussian noise with variance calibrated to C, ε, and δ is added to the aggregated gradient. The noise masks the exact contribution of each individual, thereby ensuring differential privacy.

By carefully balancing the clipping bound and noise scale, DP-SGD achieves a trade-off between privacy and model accuracy. Privacy guarantees are tracked across training iterations using a privacy accountant,

which accumulates the total privacy loss (ε, δ) over the course of training.

In contrast, gradients in VQML, computed via the PSR, are inherently bounded due to the finite spectrum of quantum observables and are stochastic because of finite-shot measurement noise. As we show next, exploiting these structural properties enables tighter sensitivity analysis and substantially reduces the amount of additional noise required to achieve DP.

3 DIFFERENTIALLY PRIVATE PARAMETER-SHIFT RULE

This section details our central contribution, the Q-ShiftDP algorithm, which is a differentially private algorithm to protect the privacy of the quantum training process, and its theoretical underpinnings. First, we present the core algorithm and establish theoretical guarantees. This includes the main privacy guarantee, that illustrates how the inherent shot noise from measurements can be used to reduce the need for externally added artificial noise, and the utility analysis. Next, we extend this analysis to noisy quantum environments where physical depolarizing noise can be a source for privacy by providing a guaranteed, non-trivial lower bound on the single-shot variance. Finally, we show how to empirically estimate the single-shot noise variance by measurement outcomes. That allows Q-ShiftDP to dynamically tailor the amount of artificial noise added. Due to the page limit, all proofs in this paper are included in Appendix B.

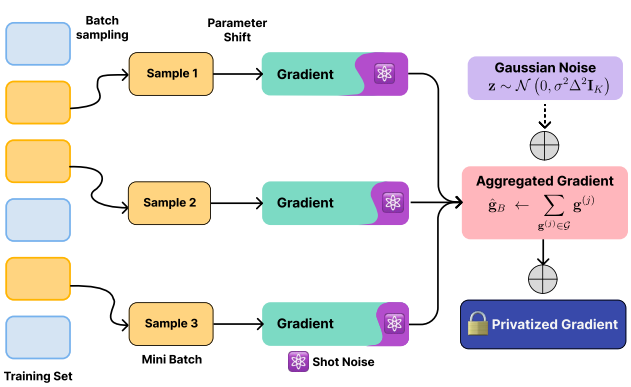

Figure 1: Overall framework of Q-ShiftDP. Each mini-batch is processed via the parameter-shift rule, yielding inherently noisy and bounded gradients. The aggregated gradient is then perturbed with Gaussian noise $z \sim \mathcal{N}(0, \sigma^2\Delta^2\mathbf{I}_K)$ to ensure (ε, δ)-DP.

3.1 Proposed Algorithm

To design an efficient differentially private algorithm for training QML models, it is essential to account for the unique characteristics of QML training, which differ from those in classical machine learning. Two key distinctions are: (i) quantum gradients are inherently bounded, and (ii) gradients are estimated through statistical sampling rather than computed exactly.

We first investigate the boundedness of gradients in QML. According to PSR, gradients are computed as the difference between two shifted expectation values. Thus, it is crucial to understand the range of possible expectation values produced by a quantum circuit. Lemma 1 shows that the expectation value is always bounded by the extremal eigenvalues of the observable.

Lemma 1. *Given the objective function $C(\boldsymbol{\theta}) = \langle\psi_0|\hat{U}(\boldsymbol{\theta})^\dagger\hat{O}\hat{U}(\boldsymbol{\theta})|\psi_0\rangle$ where $\hat{O}$ is an observable with a minimum and maximum eigenvalues λ_{min} and λ_{max} respectively and $\hat{U}(\boldsymbol{\theta})$ is a parametric unitary. The value of $C(\boldsymbol{\theta})$ is bounded by these same extremal eigenvalues for all possible parameter values $\boldsymbol{\theta}$.*

$$\lambda_{\min} \leq C(\boldsymbol{\theta}) \leq \lambda_{\max} \tag{7}$$

Next, we establish a bound on the ℓ_2-sensitivity of the gradients, as stated in Lemma 2. The bound Δ is determined by the extremal eigenvalues of the observable and the frequencies Ω_k of the generators G_k, which can be specified prior to training. Thus, the ℓ_2-sensitivity of the gradients can be controlled directly, without the need for gradient clipping. Based on this property, we design a differentially private algorithm for QML training, presented in Algorithm 1.

Lemma 2 (ℓ_2-Sensitivity of the Quantum Gradient)**.** *We define an objective function $C(\boldsymbol{\theta}) = \left\langle \psi_0(x) \middle| \hat{U}(\boldsymbol{\theta})^\dagger\hat{O}\hat{U}(\boldsymbol{\theta}) \middle| \psi_0(x) \right\rangle$ where the observable $\hat{O}$ has eigenvalues in range from $\lambda_{\min}$ to $\lambda_{\max}$ and parametric unitary $\hat{U}(\boldsymbol{\theta}) = \prod_{k=1}^{K} e^{i\theta_k G_k}$ with frequencies Ω_k. Suppose the gradient of $C(\boldsymbol{\theta})$ with respect to parameter θ_k is computed using the parameter-shift rule:*

$$\mathbf{g}_k = \tfrac{\Omega_k}{2}\,(r_+ - r_-), \tag{8}$$

*where r_+ and r_- are expectation values of the observable under a shifted amount $s = \frac{\pi}{2\Omega_k}$. The ℓ_2-**norm** of this gradient is strictly bounded by the sensitivity Δ:*

$$\|\mathbf{g}_k\|_2 \leq \Delta \quad \textit{where} \quad \Delta = \frac{\lambda_{\max} - \lambda_{\min}}{2}\sqrt{\sum_{k=1}^{K}\Omega_k^2}$$

Algorithm 1 takes as input: a data encoding map $x \mapsto |\psi_0(x)\rangle$, a parametric unitary $\hat{U}(\boldsymbol{\theta}) = \prod_{k=1}^{K} e^{i\theta_k G_k}$ with generator frequencies $\{\Omega_k\}$, a set of observables $\{\hat{O}_y\}$ whose eigenvalues lie in $[\lambda_{\min}, \lambda_{\max}]$, a

private dataset S, as well as training hyperparameters including the learning rate lr, the number of steps T, batch size B, and the noise multiplier σ. The algorithm starts by computing the theoretical ℓ_2-sensitivity bound Δ of the gradients based on Lemma 2. For each training step, a mini-batch B is sampled from the private dataset (for simplicity, we use B to denote both the mini-batch and its size). For each data point $(x_j, y_j) \in B$, we initialize a gradient vector $\mathbf{g}^{(j)}$. Then, for every parameter θ_k, the k-th component of the gradient is estimated using the parameter-shift rule: first, the shift amount $s = \pi/(2\Omega_k)$ is set; next, the intermediate state $|\phi\rangle$ and operator $\hat{A}$ are defined. Two shifted circuits, including the positive-shifted circuit $\hat{A}\hat{U}_k(\theta_k + s)|\phi\rangle$ and the negative-shifted circuit $\hat{A}\hat{U}_k(\theta_k - s)|\phi\rangle$, are then executed in N_s shots, each to obtain empirical estimates r_+ and r_-. The gradient component is computed as $\mathbf{g}_k^{(j)} = \frac{\Omega_k}{2}(r_+ - r_-)$. After all parameters are processed, $\mathbf{g}^{(j)}$ is appended to the gradient list $\mathcal{G}$. Once all samples in the batch are processed, the gradients are averaged, Gaussian noise $\mathbf{z} \sim \mathcal{N}(0, \sigma^2\Delta^2\mathbf{I}_K)$ is added, and the noisy gradient $\tilde{\mathbf{g}}$ is used to update the parameters with learning rate $\mathbf{lr}$. After T iterations, the algorithm outputs the DP parameters $\boldsymbol{\theta}_T$.

Lemma 3 establishes the per-iteration privacy guarantee of Algorithm 1. In this context, we account for the second key distinction between quantum and classical training: quantum gradients are estimated statistically. The statistical noise inherent in quantum measurements, denoted as *shot noise*, can serve as an additional source of privacy. Shot noise is given by the single-shot variance divided by the number of shots, N_s. By the central limit theorem, when the number of shots N_s is sufficiently large, the estimated gradient components follow a Gaussian distribution with variance $\sigma^2_{\text{shot}}/N_s$, where σ^2_{shot} represents a guaranteed lower bound on the single-shot variance across all possible inputs in all positive-shifted and negative-shifted circuits. Shot noise combines with the injected Gaussian noise, which is calibrated based on the ℓ_2-sensitivity, to yield the final guarantee.

Lemma 3. *Suppose N_s is large enough so that the finite-shot estimators $\hat{r}_+$ and $\hat{r}_-$ are independent random variables drawn from Gaussian distributions with the variance at least $\sigma^2_{\text{shot}}/N_s$. Then, Algorithm 1 ensures $(\varepsilon_0, \delta_0) - \text{DP}$ per iteration with:*

$$\varepsilon_0 = \frac{\sqrt{2\ln\frac{1.25}{\delta_0}}}{\sqrt{\frac{2B\sigma^2_{\text{shot}}}{N_s(\lambda_{\max}-\lambda_{\min})^2} + \sigma^2}} \tag{9}$$

We emphasize that Lemma 3 characterizes the per-iteration privacy and does not account for mini-batch subsampling or the cumulative privacy loss over T training steps. In Theorem 1, we extend this result to provide a lower bound on the artificial Gaussian noise required to achieve a desired overall privacy budget. Importantly, the inherent shot noise from finite-shot gradient estimation reduces the need for additional artificial noise: the noisier the estimation, the less extra artificial noise must be injected. We note that if gradients were computed exactly (i.e., $\sigma^2_{\text{shot}}/N_s \approx 0$), the amount of required artificial noise is equivalent with that of DP-SGD for classical ML Abadi et al. (2016a).

Theorem 1. *Consider Algorithm 1 (Q-ShiftDP) applied to a private dataset S of size N, with mini-batch size B, sampling rate $q = B/N$, T training iterations, and per-sample ℓ_2-gradient sensitivity Δ (from Lemma 2). Let the desired total privacy guarantee be (ε, δ), for any $\delta \in (0, 1)$ and for a sufficiently small ε (specifically, $\varepsilon < c_1 q^2 T$ for a universal constant c_1). The algorithm guarantees (ε, δ)-DP if the chosen explicit noise multiplier σ satisfies:*

$$\sigma^2 \geq \max\left(0, C_{DP} - \frac{2B\sigma^2_{\text{shot}}}{N_s(\lambda_{\max}-\lambda_{\min})^2}\right) \tag{10}$$

where $C_{DP} = \left(c_2 \frac{q\sqrt{T\log(1/\delta)}}{\varepsilon}\right)^2$ with c_2 is a constant.

Finally, we analyze the utility of Q-ShiftDP by quantifying the discrepancy between the statistically estimated gradient $\tilde{\mathbf{g}}$ (obtained via subsampling and measurement) and the true dataset gradient $\nabla C_S(\boldsymbol{\theta})$. Theorem 2 provides an upper bound on the expected squared ℓ_2 error of this estimate. The bound is derived from three sources of errors: (i) sampling error due to mini-batch selection, (ii) statistical error from finite-shot measurements, and (iii) variance from the artificial noise added. An interesting aspect of the bound in Theorem 2 is the dependence on the number of shots N_s. Increasing N_s improves utility by reducing the shot error in gradient estimation, thereby tightening the bound. However, this comes with a trade-off: larger N_s diminishes the inherent randomness of quantum measurements, which contributes to privacy. Thus, more artificial noise is required to maintain the same privacy guarantee.

Theorem 2 (Bound on the Gradient Estimation Error for Q-ShiftDP)**.** *Let $\nabla C_S(\boldsymbol{\theta})$ be the true gradient over the full dataset S, and let $\tilde{\mathbf{g}}$ be the final noisy gradient estimate produced by Algorithm 1 (Q-ShiftDP) for a mini-batch of size B. The Mean Squared Error (MSE) of the estimator is bounded by:*

$$\mathbb{E}\left[\|\tilde{\mathbf{g}} - \nabla C_S(\boldsymbol{\theta})\|^2\right] \leq \frac{\Delta^2}{B}\left(1 + \frac{1}{2N_s}\right) + \frac{K\sigma^2\Delta^2}{B^2} \tag{11}$$

where Δ is the per-sample ℓ_2-sensitivity from Lemma 2, N_s is the number of measurement shots, K is the number of parameters, and σ is the artificial noise multiplier.

In Q-ShiftDP, the guaranteed lower bound on the single-shot variance σ^2_{shot} plays a crucial role in determining the required amount of artificial noise σ^2. However, without additional information about the positive and negative-shifted circuits, establishing a meaningful lower bound is challenging. In next section, we derive a non-trivial lower bound σ^2_{shot} in the presence of depolarizing noise in the quantum circuits.

3.2 Single-Shot Measurement Variance under Noisy Environment

Our analysis begins with a general quantum circuit that prepares the state $|\psi(x)\rangle$ from a given input x. Thus, the following theories and bounds for the general circuit apply directly to both positive- and negative-shifted circuits in Algorithm 1. Let O be an observable with a spectral decomposition $O = \sum_i \lambda_i P_i$, where λ_i are the distinct eigenvalues and P_i are the projectors to the corresponding eigenspaces. The expectation of O with respect to the output is $\langle O\rangle_x = \langle\psi(x)| O |\psi(x)\rangle$. This expectation is practically estimated through repeated measurements (shots). A single shot yields one of the eigenvalues λ_i, with probability

$$p_i(x) = \Pr(\hat{a} = \lambda_i) = \langle\psi(x)| P_i |\psi(x)\rangle \tag{12}$$

Performing N_s independent shots on identically prepared states produces i.i.d. outcomes $\hat{a} \in \{\lambda_i\}$ drawn from the distribution $\{p_i(x)\}$.

We denote the variance of a single-shot measurement by $\eta^2_{\text{shot}}(x)$. It is given by:

$$\begin{aligned}\eta^2_{\text{shot}}(x) &= \mathbb{E}[\hat{a}^2] - (\mathbb{E}[\hat{a}])^2 \\ &= \sum_i \lambda_i^2 p_i(x) - \left(\sum_i \lambda_i p_i(x)\right)^2\end{aligned} \tag{13}$$

By definition, σ^2_{shot} is a lower bound on the single-shot variance for both the positive- and negative-shifted circuits. This interpretation remains valid if we generalize the definition by assigning:

$$\sigma^2_{\text{shot}} := \inf_x \eta^2_{\text{shot}}(x) \tag{14}$$

We observe that if there exists some x and i such that $p_i(x) = 1$, the single-shot variance σ_{shot} becomes trivially zero. Consequently, without this inherent statistical noise, from Theorem 1, Q-ShiftDP loses its primary advantage, as the required artificial noise σ cannot be reduced. In practice, however, quantum noise in quantum circuits prevents this situation from occurring.

We now consider the effect of depolarizing noise, which is a common noise in quantum circuits. Under this noise model, the output state becomes mixed, which leads to a non-trivial lower bound on σ_{shot}.

To represent quantum states in a noisy environment, we employ the density matrix formalism. The ideal state is given by $\rho_{\text{ideal}}(x) = |\psi(x)\rangle\langle\psi(x)|$. Suppose this state passes through a global depolarizing channel $\mathcal{E}_\alpha$ of strength $\alpha \in [0, 1]$ before measurement. The resulting (noisy) state is

$$\rho_{\text{noisy}}(x) = \mathcal{E}_\alpha\big(\rho_{\text{ideal}}(x)\big) = (1-\alpha)\rho_{\text{ideal}}(x) + \alpha\frac{I}{d}, \tag{15}$$

where I is the identity operator and d is the Hilbert space dimension. The state $\frac{I}{d}$ is known as *uniform state*. When measured with respect to any observable, the outcomes are distributed uniformly, i.e., $p_i^{\text{uniform}} = \frac{1}{d}$ for all i. Thus, the single-shot variance under the uniform state is a constant:

$$\sigma^2_{\text{uniform}} = \left(\frac{1}{d}\sum_i \lambda_i^2\right) - \left(\frac{1}{d}\sum_i \lambda_i\right)^2 \tag{16}$$

Theorem 3. *The single-shot variance of a measurement of an observable O on this noisy state, $\eta^2_{\text{shot}}(x)$, is lower-bounded by:*

$$\eta^2_{\text{shot}}(x) \geq \alpha \cdot \sigma^2_{\text{uniform}} \tag{17}$$

Theorem 3 indicates that in the presence of noise, the single-shot variance for any input x is lower-bounded by $\alpha \cdot \sigma^2_{\text{uniform}}$. Intuitively, the depolarizing channel $\mathcal{E}_\alpha$ mixes the ideal state with the maximally mixed state, whose variance is maximized. As a result, the variance of the noisy state is always guaranteed to exceed a non-trivial threshold. Applying the noise channel $\mathcal{E}_\alpha$ to the outputs of the positive- and negative-shifted circuits in Q-ShiftDP, we conclude that $\sigma^2_{\text{shot}} \geq \alpha \cdot \sigma^2_{\text{uniform}}$, thus ensuring a meaningful lower bound. This, in turn, reduces the amount of artificial noise σ^2 that needs to be injected (see Corollary 1).

Corollary 1. *The presence of physical depolarizing noise on the quantum circuit (i.e., $\alpha > 0$) directly reduces the amount of artificial noise σ required by Q-ShiftDP algorithm.*

3.3 Estimation of Single-Shot Measurement Variance by Samples

The central guarantee of Theorem 1 relies on the single-shot variance, σ^2_{shot}, to reduce the required artificial noise. A global lower bound on this variance can

Algorithm 1: Differentially Private Parameter-Shift Rule (Q-ShiftDP)

Input: Data encoding $x \mapsto |\psi_0(x)\rangle$, Parametric unitary $\hat{U}(\boldsymbol{\theta}) = \prod_{k=1}^{K} e^{i\theta_k G_k}$ with frequencies Ω_k, Set of observables $\{\hat{O}_y\}$ with eigenvalues in $[\lambda_{\min}, \lambda_{\max}]$, Private dataset S, Learning rate $\mathbf{lr}$, NumSteps T, Batch size B, Noise multiplier σ.

Output: Trained parameters $\boldsymbol{\theta}_T$.

```
// Define sensitivity based on the theoretical maximum L2 norm of the quantum gradient
```

$\Delta \leftarrow \frac{\lambda_{\max}-\lambda_{\min}}{2}\sqrt{\sum_{k=1}^{K}\Omega_k^2}$

Initialize parameters $\boldsymbol{\theta}_0$ randomly

for $t \leftarrow 0$ **to** $T-1$ **do**

  Sample a mini-batch $B \subset S$

  Initialize an empty list for gradient estimates: $\mathcal{G} \leftarrow []$

  foreach *sample* $(x_j, y_j) \in B$ **do**

    Initialize vector: $\mathbf{g}^{(j)} \leftarrow \mathbf{0} \in \mathbb{R}^K$

    for $k \leftarrow 1$ **to** K **do**

      Define shift amount: $s \leftarrow \pi/(2\Omega_k)$

      $|\phi\rangle \leftarrow (\prod_{l=1}^{k-1}\hat{U}_l(\theta_l))|\psi_0(x_j)\rangle$ `// Define intermediate state`

      $\hat{A} \leftarrow (\prod_{l=k+1}^{K}\hat{U}_l(\theta_l))^\dagger \hat{O}_{y_j}(\prod_{l=k+1}^{K}\hat{U}_l(\theta_l))$ `// Define intermediate operator`

      Run positive-shifted circuit $\hat{A}\hat{U}_k(\theta_k+s)|\phi\rangle$ with N_s times, collect outcomes $\{\hat{a}_{j,k,i}\}_{i=1}^{N_s}$

      $r_+ \leftarrow \frac{1}{N_s}\sum_{i=1}^{N_s}\hat{a}_{j,k,i}$ `// Estimator for positive shift`

      Run negative-shifted circuit $\hat{A}\hat{U}_k(\theta_k-s)|\phi\rangle$ with N_s times, collect outcomes $\{\hat{b}_{j,k,i}\}_{i=1}^{N_s}$

      $r_- \leftarrow \frac{1}{N_s}\sum_{i=1}^{N_s}\hat{b}_{j,k,i}$ `// Estimator for negative shift`

      $\mathbf{g}_k^{(j)} \leftarrow \frac{\Omega_k}{2}(r_+ - r_-)$ `// Parameter-shift rule estimator`

    Append $\mathbf{g}^{(j)}$ to $\mathcal{G}$

  $\hat{\mathbf{g}}_B \leftarrow \sum_{\mathbf{g}^{(j)}\in\mathcal{G}}\mathbf{g}^{(j)}$

  Sample noise: $\mathbf{z} \sim \mathcal{N}(0, \sigma^2\Delta^2\mathbf{I}_K)$; $\tilde{\mathbf{g}} \leftarrow \frac{1}{B}(\hat{\mathbf{g}}_B + \mathbf{z})$; $\boldsymbol{\theta}_{t+1} \leftarrow \boldsymbol{\theta}_t - \mathbf{lr}\cdot\tilde{\mathbf{g}}$

return $\boldsymbol{\theta}_T$

be conservative. The actual variance for many samples may be much higher than this worst-case floor. Thus, we propose an adaptive algorithm that replaces the conservative bound with an empirical estimate of shot-noise variance from measurement outcomes. This approach allows us to tailor the amount of added privacy noise based on the overall randomness of each sample's measurements. To make this approach practical and simplify the analysis, we assume all generator frequencies are identical, $\Omega_k = \Omega \forall k$. This is a common scenario in variational circuits built from repeating layers of the same learnable gates (e.g., R_Y).

Given an input x_j and weight θ_k, let the true single-shot variances of the positive- and negative-shifted circuits be $\eta_{j,k,+}^2$ and $\eta_{j,k,-}^2$, respectively. We define the *total batch shot-noise variance*, $\boldsymbol{\eta}_B^2$, as the sum of all single-shot variances from both the positive and negative-shifted circuits for every sample and trainable weight in a batch B, i.e., $\boldsymbol{\eta}_B^2 = \sum_j\sum_k\sum_{\tau\in\pm}\eta_{j,k,\tau}^2$. We can achieve the same privacy guarantee by assigning a batch-specific noise level σ_B, rather than adding a global noise σ as in Algorithm 1, such that:

$$\sigma_B^2 \geq \max\left(0, C_{\mathrm{DP}} - \frac{\Omega^2\boldsymbol{\eta}_B^2}{4N_s\Delta^2}\right) \tag{18}$$

However, $\boldsymbol{\eta}_B^2$ is unknown a priori. Based on the outcomes of the positive- and negative-shifted circuits, we aim to construct an estimator $\hat{\boldsymbol{\eta}}_B^2$ for the total true variance $\boldsymbol{\eta}_B^2$. This estimator is then used to determine the batch-specific artificial noise σ_B^2. Ideally, the privacy guarantee level remains intact if the estimator is conservative, i.e., $\boldsymbol{\eta}_B^2 \geq \hat{\boldsymbol{\eta}}_B^2$. We also want the estimator to be as close as possible to the true variance, thereby reducing the amount of artificial noise added.

We now describe our variance estimation technique, which is based on Lemma 4. This lemma guarantees, with high probability, that $\boldsymbol{\eta}_B^2$ is lower-bounded by a function of sample variances and sample fourth central moments, both of which are computable from measurement outcomes. Given sets of outcomes for each sample j and weight k in the batch, we denote the computable sample variances as $\bar{\eta}_{j,k,\pm}^2$ and the sample fourth central moments as $\bar{\mu}_{4,j,k,\pm}$. Our variance

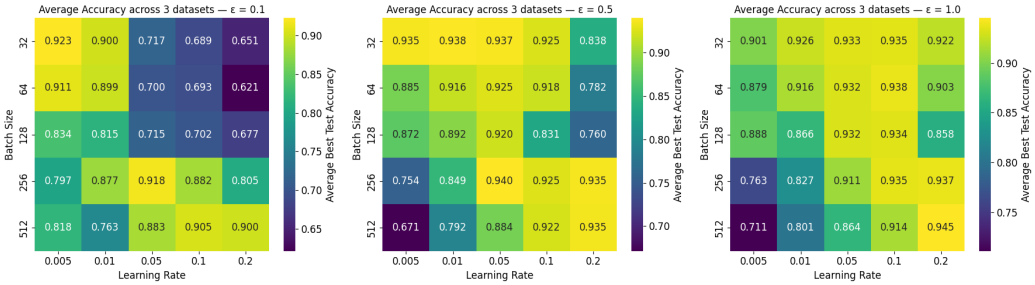

Figure 2: Average test accuracy on 3 datasets using Q-ShiftDP w.r.t batch size, learning rates and ε.

estimator is then formulated as:

$$\hat{\boldsymbol{\eta}}_B^2 = \bar{\boldsymbol{\eta}}_B^2 - z_\beta \sqrt{\sum_{j=1}^{B}\sum_{k=1}^{K}\sum_{\tau\in\pm}\frac{1}{N_s}\left(\bar{\mu}_{4,j,k,\tau} - (\bar{\eta}_{j,k,\tau}^2)^2\right)} \tag{19}$$

where $\bar{\boldsymbol{\eta}}_B^2 = \sum_{j=1}^{B}\sum_{k=1}^{K}\sum_{\tau\in\{+,-\}}\bar{\eta}_{j,k,\tau}^2$.

Lemma 4. *Let there be L independent groups, indexed by $i = 1, \ldots, L$. For each group i, let $y_{i,1}, \ldots, y_{i,n} \sim$ i.i.d. Y_i be n observations. Let η_i^2, $\bar{\eta}_i^2$, $\mu_{4,i}$ and $\bar{\mu}_{4,i}$ be the true variance, sample variance, true fourth central moment, and sample fourth central moment for the i-th group. Let $\eta_{batch}^2 = \sum_{i=1}^{L}\eta_i^2$ and $\bar{\eta}_{batch}^2 = \sum_{i=1}^{L}\bar{\eta}_i^2$. Assuming all $\mu_{4,i}$ are finite and given a significance level β with critical value z_β, then:*

$$P\left(\eta_{batch}^2 \geq \bar{\eta}_{batch}^2 - z_\beta\sqrt{\sum_{i=1}^{L}\frac{1}{n}\left(\bar{\mu}_{4,i} - (\bar{\eta}_i^2)^2\right)}\right) \to 1-\beta$$

as $L \to \infty, n \to \infty$.

By using this probabilistic estimator, we propose an adaptive algorithm, named Adaptive Q-ShiftDP (see Algorithm 2 in Appendix E), which tailors the amount of artificial noise for each mini-batch. The main idea is to calibrate the artificial noise per batch, σ_B^2, using a probabilistic estimator $\hat{\boldsymbol{\eta}}_B^2$. From Lemma 4, this estimator is conservative with high probability $\Pr(\boldsymbol{\eta}_B^2 \geq \hat{\boldsymbol{\eta}}_B^2) \to 1-\beta$. This reliance on a statistical estimate means the privacy guarantee itself is probabilistic. We can thus analyze the algorithm's behavior in two distinct *Modes* for each batch. The mechanism operates in the *Good Mode* when the estimate is a valid lower bound, ensuring the privacy guarantee holds (i.e., $\boldsymbol{\eta}_B^2 \geq \hat{\boldsymbol{\eta}}_B^2$). Conversely, in the *Bad Mode*, the estimator is an overestimate, thereby compromising the privacy guarantee (i.e., $\boldsymbol{\eta}_B^2 < \hat{\boldsymbol{\eta}}_B^2$). Lemma 5 provides a formal tool to analyze the overall privacy of such a probabilistic mechanism by converting the probability of success into a standard DP guarantee.

Lemma 5. *Let $\mathcal{M}$ be a randomized mechanism that operates in one of two modes. With probability γ, it operates in a "Good Mode," where its execution is equivalent to a mechanism $\mathcal{M}_{good}$ that is (ε, δ)-differentially private. With probability $1-\gamma$, it operates in a "Bad Mode," where no privacy guarantee holds. Then, the overall mechanism $\mathcal{M}$ is $(\varepsilon, \gamma\delta + (1-\gamma))$-DP.*

From Lemma 4, the probability of *Good Mode* is $1-\beta$, while that of *Bad Mode* is β. As discussed above, in *Good Mode*, the privacy level remains (ε, δ). Consequently, our variance estimation technique attains the DP guarantee formalized in Theorem 4. More details of Adaptive Q-ShiftDP are included in Appendix E.

Theorem 4. *For a mini-batch of size B, let $\hat{\boldsymbol{\eta}}_B^2$ be a probabilistic estimator for the true sum of all variances, $\boldsymbol{\eta}_B^2 = \sum_{j=1}^{B}\sum_{k=1}^{K}\sum_{\tau\in\pm}\eta_{j,k,\tau}^2$ such that it satisfies $\Pr(\boldsymbol{\eta}_B^2 \geq \hat{\boldsymbol{\eta}}_B^2) \to 1-\beta$ for a chosen significance level β. Adaptive Q-ShiftDP calculates a noise multiplier for each batch B, σ_B^2, satisfying:*

$$\sigma_B^2 \geq \max\left(0, C_{DP} - \frac{\Omega^2\hat{\boldsymbol{\eta}}_B^2}{4N_s\Delta^2}\right)$$

*where $C_{DP} = \left(c_2\frac{q\sqrt{T\log(1/\delta)}}{\varepsilon}\right)^2$ with c_2 is a constant. Adaptive Q-ShiftDP achieves $(\varepsilon, (1-\beta)\delta + \beta)$-**DP**.*

4 EXPERIMENTAL RESULTS

We evaluate the performance and privacy–utility trade-off of Q-ShiftDP on three representative datasets from recent QML benchmark suites: Bars & Stripes, Binary Blobs and Downscaled MNIST Bowles et al. (2024); Recio-Armengol et al. (2025). These benchmarks are carefully constructed, low-dimensional tasks that align with quantum circuit constraints and allow meaningful, reproducible evaluation of quantum models. Our classifier is a 4-qubit quantum neural network with a single Strongly Entangling Layer for Bars and Stripes dataset and 5 layers for other datasets. Details about datasets and models are given in Appendix D.1.

Performance of Q-ShiftDP under Varying Hyperparameter Settings. The performance of differentially private optimizers is highly sensitive to hyperparameter choices Bu et al. (2023). Unlike DP-SGD,

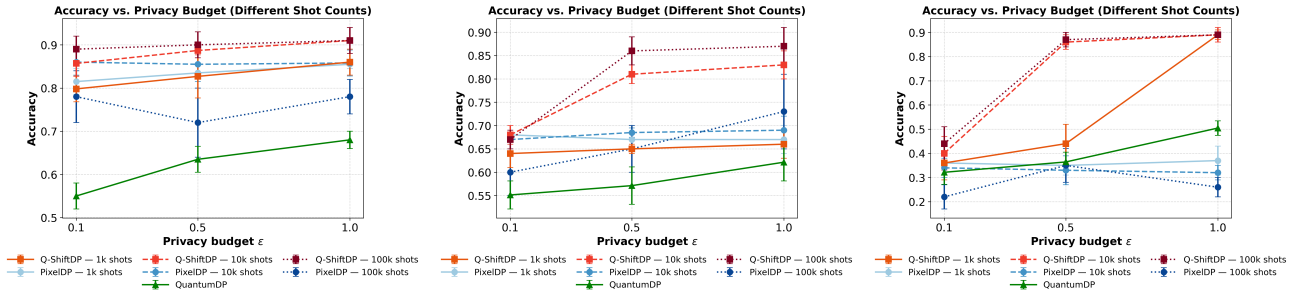

(a) Bars & Stripes (b) Downscaled MNIST (c) Binary Blobs

Figure 3: Comparison of model utility under equal privacy budgets: Q-ShiftDP vs. Pixel-level DP vs. QuantumDP in various datasets.

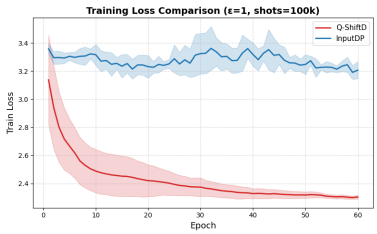

Figure 4: Training loss curves for Q-ShiftDP and Pixel-level DP under $\varepsilon = 1$ with 100k shots.

Q-ShiftDP leverages the parameter-shift rule and the bounded spectrum of quantum observables to derive an exact analytic sensitivity bound, removing the need to tune the clipping norm. Figure 2 reports the average test accuracy across all datasets for different combinations of learning rate, batch size, and privacy budget, with detailed results provided in Appendix D.2. We observe that large batch sizes paired with moderately high learning rates consistently yield the most stable performance. This result is consistent with previous DP-SGD literature De et al. (2022). Based on these findings, all subsequent experiments use a batch size of $B = 512$ and a learning rate of $\mathrm{lr} = 0.2$. In this experiment, we use analytic expectation values rather than limiting the number of measurement shots. The effect of finite shots on Q-ShiftDP's performance is further analyzed in Appendix D.3.

Comparison to Existing Privacy Baselines We compare Q-ShiftDP against PixelDP Lecuyer et al. (2019) and QuantumDP Hirche et al. (2023). While PixelDP enforces differential privacy by injecting classical noise directly into the input data, QuantumDP achieves DP by relying solely on intrinsic quantum noise (i.e., depolarizing channels). Importantly, we view Q-ShiftDP as a theoretically optimized adaptation of DP-SGD for QML. Thus, we do not include a direct empirical comparison to DP-SGD, as Q-ShiftDP theoretically dominates DP-SGD when applied to QML.

Figure 3 reports test accuracy as a function of the privacy budget ε for different shot counts across three benchmark datasets. For QuantumDP, results are reported using 100k shots. Across all datasets and experimental settings, Q-ShiftDP consistently outperforms both PixelDP and QuantumDP under the same privacy budget and shot count. The performance gap is particularly pronounced in the strong privacy regime ($\varepsilon = 0.1$), where input perturbation and quantum-only noise lead to substantial utility degradation. We note that training is highly unstable under very low privacy budget ($\varepsilon = 0.1$) or number of shots (1k shots). Figure 4 further compares the training dynamics of Q-ShiftDP and PixelDP. Under $\varepsilon = 1$ with 100k shots, Q-ShiftDP exhibits significantly more stable training behavior than PixelDP. Taken together, these results suggest that jointly accounting for both noise sources, as done in Q-ShiftDP, yields a more favorable privacy–utility trade-off in private quantum learning.

Additional Experiments. We provide further experimental results and analyses in the Appendix D.

5 CONCLUSION

In this work, we introduce Q-ShiftDP, a privacy mechanism designed for quantum machine learning. By proving that quantum gradients are inherently bounded and their intrinsic shot noise is a quantifiable privacy resource, our method eliminates the need for manual clipping and provably reduces the amount of artificial noise required. We established formal privacy and utility guarantee, further enhanced by an adaptive technique that tailors the noise to the specifics of each batch. Experimental results show that Q-ShiftDP outperforms the existing benchmark in multiple settings.

Acknowledgements

This work is partially supported by the National Science Foundation under grants IIS-2416606 and IIS-2123809.

References

Martin Abadi, Andy Chu, Ian Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, CCS '16, page 308–318, New York, NY, USA, 2016a. Association for Computing Machinery. ISBN 9781450341394. doi: 10.1145/2976749.2978318. URL https://doi.org/10.1145/2976749.2978318.

Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pages 308–318, 2016b.

Ya-Ru Bai, Yuan-Hong Tao, Shu-Hui Wu, Hui Zhang, and Shao-Ming Fei. Quantum differential privacy under noise channels. *Physica Scripta*, 99(3):035119, 2024.

Borja Balle, Gilles Barthe, and Marco Gaboardi. Privacy amplification by subsampling: Tight analyses via couplings and divergences. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL https://proceedings.neurips.cc/paper_files/paper/2018/file/3b5020bb891119b9f5130f1fea9bd773-Paper.pdf.

Franziska Boenisch, Christopher Mühl, Adam Dziedzic, Roy Rinberg, and Nicolas Papernot. Have it your way: individualized privacy assignment for dp-sgd. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, NIPS '23, Red Hook, NY, USA, 2023. Curran Associates Inc.

Joseph Bowles, Shahnawaz Ahmed, and Maria Schuld. Better than classical? the subtle art of benchmarking quantum machine learning models. *arXiv preprint arXiv:2403.07059*, 2024.

Zhiqi Bu, Yu-Xiang Wang, Sheng Zha, and George Karypis. Automatic clipping: Differentially private deep learning made easier and stronger. *Advances in Neural Information Processing Systems*, 36:41727–41764, 2023.

Vamshi Krishna Bukkarayasamudram, Pundru Chandra Shaker Reddy, K. Arun Kumar, R. M. Jagadish, Swati Sharma, Mudarakola Lakshmi Prasad, Yadala Sucharitha, Iftikhar Aslam Tayubi, and Gopal Kumar Thakur. Quantum computing revolution in healthcare: a systematic review of applications, issues and future directions. *Artificial Intelligence Review*, 58(12):389, Oct 2025. ISSN 1573-7462. doi: 10.1007/s10462-025-11381-w. URL https://doi.org/10.1007/s10462-025-11381-w.

X.-D. Cai, D. Wu, Z.-E. Su, M.-C. Chen, X.-L. Wang, Li Li, N.-L. Liu, C.-Y. Lu, and J.-W. Pan. Entanglement-based machine learning on a quantum computer. *Phys. Rev. Lett.*, 114: 110504, Mar 2015. doi: 10.1103/PhysRevLett.114.110504. URL https://link.aps.org/doi/10.1103/PhysRevLett.114.110504.

Kamalika Chaudhuri and Claire Monteleoni. Privacy-preserving logistic regression. *Advances in neural information processing systems*, 21, 2008.

Kamalika Chaudhuri, Claire Monteleoni, and Anand D Sarwate. Differentially private empirical risk minimization. *Journal of Machine Learning Research*, 12(3), 2011.

Chi-Sheng Chen and Samuel Yen-Chi Chen. Q-dpts: Quantum differentially private time series forecasting via variational quantum circuits, 2025. URL https://arxiv.org/abs/2508.05036.

Xiangyi Chen, Steven Z Wu, and Mingyi Hong. Understanding gradient clipping in private sgd: A geometric perspective. *Advances in Neural Information Processing Systems*, 33:13773–13782, 2020.

Soham De, Leonard Berrada, Jamie Hayes, Samuel L Smith, and Borja Balle. Unlocking high-accuracy differentially private image classification through scale. *arXiv preprint arXiv:2204.13650*, 2022.

Jinshuo Dong, Aaron Roth, and Weijie J Su. Gaussian differential privacy. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 84(1):3–37, 2022.

Yuxuan Du, Min-Hsiu Hsieh, Tongliang Liu, Dacheng Tao, and Nana Liu. Quantum noise protects quantum classifiers against adversaries. *Phys. Rev. Res.*, 3:023153, May 2021.

Cynthia Dwork. Differential privacy. In *International Colloquium on Automata, Languages and Programming*, 2006a. URL https://api.semanticscholar.org/CorpusID:2565493.

Cynthia Dwork. Differential privacy. In *International colloquium on automata, languages, and programming*, pages 1–12. Springer, 2006b.

Jiyan He, Xuechen Li, Da Yu, Huishuai Zhang, Janardhan Kulkarni, Yin Tat Lee, Arturs Backurs, Nenghai Yu, and Jiang Bian. Exploring the limits of differentially private deep learning with group-wise clipping. In *ICLR*, 2023.

Christoph Hirche, Cambyse Rouzé, and Daniel Stilck França. Quantum differential privacy: An information theory perspective. *IEEE Transactions on Information Theory*, 69(9):5771–5787, 2023. doi: 10.1109/TIT.2023.3272904.

Roger Iyengar, Joseph P Near, Dawn Song, Om Thakkar, Abhradeep Thakurta, and Lun Wang. Towards practical differentially private convex optimization. In *2019 IEEE symposium on security and privacy (SP)*, pages 299–316. IEEE, 2019.

Ismat Jarin and Birhanu Eshete. Dp-util: Comprehensive utility analysis of differential privacy in machine learning. In *Proceedings of the Twelfth ACM Conference on Data and Application Security and Privacy*, CODASPY '22, page 41–52, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450392204. doi: 10.1145/3508398.3511513. URL https://doi.org/10.1145/3508398.3511513.

Daniel Kifer, Adam Smith, and Abhradeep Thakurta. Private convex empirical risk minimization and high-dimensional regression. In *Conference on Learning Theory*, pages 25–1. JMLR Workshop and Conference Proceedings, 2012.

Mathias Lécuyer, Vaggelis Atlidakis, Roxana Geambasu, Daniel J. Hsu, and Suman Sekhar Jana. Certified robustness to adversarial examples with differential privacy. *2019 IEEE Symposium on Security and Privacy (SP)*, pages 656–672, 2018.

Mathias Lecuyer, Vaggelis Atlidakis, Roxana Geambasu, Daniel Hsu, and Suman Jana. Certified robustness to adversarial examples with differential privacy. In *2019 IEEE symposium on security and privacy (SP)*, pages 656–672. IEEE, 2019.

Jun Li, Xiaodong Yang, Xinhua Peng, and Chang-Pu Sun. Hybrid quantum-classical approach to quantum optimal control. *Phys. Rev. Lett.*, 118: 150503, Apr 2017. doi: 10.1103/PhysRevLett.118.150503. URL https://link.aps.org/doi/10.1103/PhysRevLett.118.150503.

Xuechen Li, Florian Tramer, Percy Liang, and Tatsunori Hashimoto. Large language models can be strong differentially private learners. *arXiv preprint arXiv:2110.05679*, 2021.

Seth Lloyd, Masoud Mohseni, and Patrick Rebentrost. Quantum algorithms for supervised and unsupervised machine learning, 2013. URL https://arxiv.org/abs/1307.0411.

Agostino Marengo and Vito Santamato. Quantum algorithms and complexity in healthcare applications: a systematic review with machine learning-optimized analysis. *Frontiers in Computer Science*, Volume 7 - 2025, 2025. ISSN 2624-9898. doi: 10.3389/fcomp.2025.1584114. URL https://www.frontiersin.org/journals/computer-science/articles/10.3389/fcomp.2025.1584114.

Ilya Mironov. Rényi differential privacy. In *2017 IEEE 30th computer security foundations symposium (CSF)*, pages 263–275. IEEE, 2017.

K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii. Quantum circuit learning. *Phys. Rev. A*, 98:032309, Sep 2018. doi: 10.1103/PhysRevA.98.032309. URL https://link.aps.org/doi/10.1103/PhysRevA.98.032309.

V. Moretti. *Fundamental Mathematical Structures of Quantum Theory: Spectral Theory, Foundational Issues, Symmetries, Algebraic Formulation.* Springer International Publishing, 2020. ISBN 9783030183486. URL https://books.google.com/books?id=Ekm_zQEACAAJ.

Seth Neel, Aaron Roth, Giuseppe Vietri, and Zhiwei Steven Wu. Differentially private objective perturbation: Beyond smoothness and convexity. *arXiv preprint arXiv:1909.01783*, 2019.

NhatHai Phan, Yue Wang, Xintao Wu, and Dejing Dou. Differential privacy preservation for deep auto-encoders: an application of human behavior prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30, 2016.

Anand Singh Rajawat, S. B. Goyal, Pradeep Bedi, Tony Jan, Md Whaiduzzaman, and Mukesh Prasad. Quantum machine learning for security assessment in the internet of medical things (iomt). *Future Internet*, 15(8), 2023. ISSN 1999-5903. doi: 10.3390/fi15080271. URL https://www.mdpi.com/1999-5903/15/8/271.

Erik Recio-Armengol, Shahnawaz Ahmed, and Joseph Bowles. Train on classical, deploy on quantum: scaling generative quantum machine learning to a thousand qubits. *arXiv preprint arXiv:2503.02934*, 2025.

Tom Sander, Pierre Stock, and Alexandre Sablayrolles. Tan without a burn: Scaling laws of dp-sgd. In *International Conference on Machine Learning*, pages 29937–29949. PMLR, 2023.

Maria Schuld and Nathan Killoran. Quantum machine learning in feature hilbert spaces. *Phys. Rev. Lett.*, 122:040504, Feb 2019. doi: 10.1103/PhysRevLett.122.040504. URL https://link.aps.org/doi/10.1103/PhysRevLett.122.040504.

Baobao Song, Shiva Raj Pokhrel, Athanasios V Vasilakos, Tianqing Zhu, and Gang Li. Towards a hybrid

quantum differential privacy. *IEEE Journal on Selected Areas in Communications*, 2025.

Shuang Song, Kamalika Chaudhuri, and Anand D Sarwate. Stochastic gradient descent with differentially private updates. In *2013 IEEE global conference on signal and information processing*, pages 245–248. IEEE, 2013.

Xinyu Tang, Ashwinee Panda, Milad Nasr, Saeed Mahloujifar, and Prateek Mittal. Private fine-tuning of large language models with zeroth-order optimization. *arXiv preprint arXiv:2401.04343*, 2024.

Om Thakkar, Galen Andrew, and H. B. McMahan. Differentially private learning with adaptive clipping. In *Neural Information Processing Systems*, 2019. URL `https://api.semanticscholar.org/CorpusID:150373827`.

Sohum Thakkar, Skander Kazdaghli, Natansh Mathur, Iordanis Kerenidis, André J. Ferreira-Martins, and Samurai Brito. Improved financial forecasting via quantum machine learning. *Quantum Machine Intelligence*, 6(1):27, May 2024. ISSN 2524-4914. doi: 10.1007/s42484-024-00157-0. URL `https://doi.org/10.1007/s42484-024-00157-0`.

William Watkins, Samuel Yen-Chi Chen, and Shinjae Yoo. Quantum machine learning with differential privacy. *Scientific Reports*, 13, 2021. URL `https://api.semanticscholar.org/CorpusID:232170478`.

William M Watkins, Samuel Yen-Chi Chen, and Shinjae Yoo. Quantum machine learning with differential privacy. *Scientific Reports*, 13(1):2453, 2023.

David Wierichs, Josh Izaac, Cody Wang, and Cedric Yen-Yu Lin. General parameter-shift rules for quantum gradients. *Quantum*, 6:677, March 2022. ISSN 2521-327X. doi: 10.22331/q-2022-03-30-677. URL `https://doi.org/10.22331/q-2022-03-30-677`.

Jun Zhang, Zhenjie Zhang, Xiaokui Xiao, Yin Yang, and Marianne Winslett. Functional mechanism: Regression analysis under differential privacy. *arXiv preprint arXiv:1208.0219*, 2012.

Li Zhou and Mingsheng Ying. Differential privacy in quantum computation. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pages 249–262, 2017. doi: 10.1109/CSF.2017.23.

Checklist

1. For all models and algorithms presented, check if you include:
   (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]
   (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes]
   (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Yes, link to anonymized source code is included in the appendix.]
2. For any theoretical claim, check if you include:
   (a) Statements of the full set of assumptions of all theoretical results. [Yes]
   (b) Complete proofs of all theoretical results. [Yes]
   (c) Clear explanations of any assumptions. [Yes]
3. For all figures and tables that present empirical results, check if you include:
   (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes]
   (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes]
   (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes, except error statistics are omitted due to the high computational cost of running repeated quantum simulations.]
   (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
   (a) Citations of the creator If your work uses existing assets. [Yes]
   (b) The license information of the assets, if applicable. [Not Applicable]
   (c) New assets either in the supplemental material or as a URL, if applicable. [Not Applicable]
   (d) Information about consent from data providers/curators. [Not Applicable]
   (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
   (a) The full text of instructions given to participants and screenshots. [Not Applicable]

(b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]

(c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

Supplementary Materials

A Additional Backgrounds

Quantum Systems in Noisy Environment. Realistic quantum systems are inherently open, interacting with their surrounding environment. Thus, it is subject to noise, decoherence, and operational imperfections. In this context, the description of quantum states is generalized from pure states to density matrices. A density matrix can represent mixed states, where the system is described as a statistical ensemble of pure states $\{p_i, |\psi_i\rangle\}$ as follow:

$$\rho = p_i \left|\psi_i\right\rangle \left\langle\psi_i\right|, \quad \sum_i p_i = 1 \tag{20}$$

The dynamics of such states are modeled by quantum channels, which are completely positive, trace-preserving (CPTP) maps acting on ρ. A general channel $\mathcal{E}$ follows a Kraus representation,

$$\mathcal{E}(\rho) = \sum_i K_i \rho K_i^\dagger, \quad \sum_i K_i^\dagger K_i = I \tag{21}$$

where $\{K_i\}$ are Kraus operators. Depolarizing channel is an example of quantum noisy channel. Specifically, in d-dimensional system, it replaces the input state with the maximally mixed state I/d with probability α, and leaves it unchanged with probability $1-\alpha$, i.e.,

$$\mathcal{E}_\alpha(\rho) = (1-\alpha)\rho + \alpha\frac{I}{d} \tag{22}$$

B Proofs

Lemma 1. *Given the objective function $C(\boldsymbol{\theta}) = \langle\psi_0|\hat{U}(\boldsymbol{\theta})^\dagger \hat{O} \hat{U}(\boldsymbol{\theta})|\psi_0\rangle$ where $\hat{O}$ is an observable with a minimum and maximum eigenvalues λ_{min} and λ_{max} respectively and $\hat{U}(\boldsymbol{\theta})$ is a product of unitary operators. The value of $C(\boldsymbol{\theta})$ is bounded by these same extremal eigenvalues for all possible parameter values $\boldsymbol{\theta}$.*

$$\lambda_{\min} \leq C(\boldsymbol{\theta}) \leq \lambda_{\max}$$

Proof. In a finite-dimensional space, $\hat{O}$ can be expressed in terms of its eigenvalues λ_i and corresponding orthonormal eigenvectors $|e_i\rangle$ as:

$$\hat{O} = \sum_i \lambda_i \left|e_i\right\rangle \left\langle e_i\right|$$

where $\hat{O}\left|e_i\right\rangle = \lambda_i \left|e_i\right\rangle$ and $\langle e_i|e_j\rangle = \delta_{ij}$. The set of eigenvectors $\{|e_i\rangle\}$ forms an orthonormal basis in the Hilbert space $\mathcal{H}$.

We call the state of the system after the unitary evolution as $|\psi(\boldsymbol{\theta})\rangle = \hat{U}(\boldsymbol{\theta})\left|\psi_0\right\rangle$. Then, the objective function can be rewritten as:

$$C(\boldsymbol{\theta}) = \langle\psi(\boldsymbol{\theta})|\hat{O}|\psi(\boldsymbol{\theta})\rangle$$

On the other hand, we can express $|\psi(\boldsymbol{\theta})\rangle$ in the eigenbasis of $\hat{O}$:

$$|\psi(\boldsymbol{\theta})\rangle = \sum_i c_i \left|e_i\right\rangle, \quad \text{where } c_i = \langle e_i|\psi(\boldsymbol{\theta})\rangle$$

By substitution, we have:

$$\begin{aligned} C(\boldsymbol{\theta}) &= \langle\psi(\boldsymbol{\theta})| \left(\sum_j \lambda_j |e_j\rangle\langle e_j|\right) |\psi(\boldsymbol{\theta})\rangle \\ &= \sum_j \lambda_j \langle\psi(\boldsymbol{\theta})\,|\,e_j\rangle\langle e_j\,|\,\psi(\boldsymbol{\theta})\rangle \\ &= \sum_j \lambda_j |\langle e_j|\psi(\boldsymbol{\theta})\rangle|^2 \\ &= \sum_j |c_j|^2\lambda_j \end{aligned}$$

This shows that $C(\boldsymbol{\theta})$ is a convex combination of the eigenvalues of $\hat{O}$ because $|c_j|^2 \geq 0$ and $\sum_j |c_j|^2 = 1$ (the normalization condition of quantum state $|\psi(\boldsymbol{\theta})\rangle$). As a result, we have the upperbound and lowerbound for $C(\boldsymbol{\theta})$ as following.

For the lower bound:

$$C(\boldsymbol{\theta}) = \sum_j |c_j|^2\lambda_j \geq \sum_j |c_j|^2\lambda_{\min} = \lambda_{\min}\left(\sum_j |c_j|^2\right) = \lambda_{\min}$$

For the upperbound:

$$C(\boldsymbol{\theta}) = \sum_j |c_j|^2\lambda_j \leq \sum_j |c_j|^2\lambda_{\max} = \lambda_{\max}\left(\sum_j |c_j|^2\right) = \lambda_{\max}$$

□

Lemma 2 (ℓ_2-Sensitivity of the Quantum Gradient)**.** *We define an objective function* $C(\boldsymbol{\theta}) = \left\langle \psi_0(x) \middle| \hat{U}(\boldsymbol{\theta})^\dagger \hat{O} \hat{U}(\boldsymbol{\theta}) \middle| \psi_0(x) \right\rangle$ *where the observable* $\hat{O}$ *has eigenvalues in range from* $\lambda_{\min}$ *to* $\lambda_{\max}$ *and parametric unitary* $\hat{U}(\boldsymbol{\theta}) = \prod_{k=1}^{K} e^{i\theta_k G_k}$ *with frequencies* Ω_k*. Suppose the gradient of* $C(\boldsymbol{\theta})$ *with respect to parameter* θ_k *is computed using the parameter-shift rule:*

$$\mathbf{g}_k = \tfrac{\Omega_k}{2}\,(r_+ - r_-), \tag{23}$$

where r_+ *and* r_- *are expectation values of the observable under a shifted amount* $s = \frac{\pi}{2\Omega_k}$*. The* ℓ_2***-norm*** *of this gradient is strictly bounded by the sensitivity* Δ*:*

$$\|\mathbf{g}_k\|_2 \leq \Delta \quad \textit{where} \quad \Delta = \frac{\lambda_{\max} - \lambda_{\min}}{2}\sqrt{\sum_{k=1}^{K}\Omega_k^2}$$

Proof. First, we establish a bound on each individual component of the gradient and then combine them to bound the L2 norm. The k-th component of the gradient is $\mathbf{g}_k^{(j)} = \frac{\Omega_k}{2}(r_+ - r_-)$. From Lemma 1, we have:

$$\lambda_{\min} \leq r_+ \leq \lambda_{\max} \quad \text{and} \quad \lambda_{\min} \leq r_- \leq \lambda_{\max}$$

As a result, we have:

$$|\mathbf{g}_k^{(j)}| = \left|\frac{\Omega_k}{2}(r_+ - r_-)\right| \leq \frac{|\Omega_k|}{2}(\lambda_{\max} - \lambda_{\min})$$

The squared L2 norm is the sum of the squares of the components. Thus, we have:

$$\begin{aligned}\|\mathbf{g}^{(j)}\|_2^2 &= \sum_{k=1}^{d}\left(\mathbf{g}_k^{(j)}\right)^2 \\ &\leq \sum_{k=1}^{d}\left(\frac{|\Omega_k|}{2}(\lambda_{\max}-\lambda_{\min})\right)^2 \\ &= \left(\frac{\lambda_{\max}-\lambda_{\min}}{2}\right)^2\sum_{k=1}^{d}\Omega_k^2\end{aligned}$$

Finally, we can conclude:

$$\|\mathbf{g}^{(j)}\|_2 \leq \frac{\lambda_{\max}-\lambda_{\min}}{2}\sqrt{\sum_{k=1}^{d}\Omega_k^2}$$

□

Lemma 3. *Suppose N_s is large enough so that the finite-shot estimators $\hat{r}_+$ and $\hat{r}_-$ are independent random variables drawn from Gaussian distributions with the variance at least $\sigma_{\text{shot}}^2/N_s$. Then, Algorithm 1 ensures $(\varepsilon_0, \delta_0)-\text{DP}$ per iteration with:*

$$\varepsilon_0 = \frac{\sqrt{2\ln\frac{1.25}{\delta_0}}}{\sqrt{\frac{2B\sigma_{\text{shot}}^2}{N_s(\lambda_{\max}-\lambda_{\min})^2}+\sigma^2}} \tag{24}$$

Proof. We consider the batch-averaged gradient:

$$\tilde{\mathbf{g}} = \frac{1}{B}\sum_{j=1}^{B}\mathbf{g}^{(j)} + \frac{1}{B}\mathbf{z}$$

where $\mathbf{g}^{(j)}$ is the per-example gradient and $z \sim \mathcal{N}(0, \sigma^2\Delta^2\mathbf{I}_K)$ is the Gaussian noise added for DP.

Because $\sum_{j=1}^{B}\mathbf{g}^{(j)}$ and $\mathbf{z}$ are independent, we have $\tilde{\mathbf{g}}$ is Gaussian with:

$$\text{Var}[\tilde{\mathbf{g}}] = \text{Var}[\frac{1}{B}\sum_{j=1}^{B}\mathbf{g}^{(j)}] + \text{Var}[\frac{1}{B}\mathbf{z}]$$

.

The estimated k-th component of the gradient is $\mathbf{g}_k^{(j)} = \frac{\Omega_k}{2}(r_+ - r_-)$. Its variance is:

$$\begin{aligned}\text{Var}[\mathbf{g}_k^{(j)}] &= \text{Var}\left[\frac{\Omega_k}{2}(r_+ - r_-)\right] = \left(\frac{\Omega_k}{2}\right)^2\text{Var}[r_+ - r_-] \\ &= \frac{\Omega_k^2}{4}\left(\text{Var}[r_+] + \text{Var}[r_-]\right) \geq \frac{\Omega_k^2}{4}(2\sigma_{\text{shot}}^2/N_s) = \frac{\Omega_k^2\sigma_{\text{shot}}^2}{2N_s}\end{aligned}$$

Then, the total variance of the statistical noise vector is the sum of the component variances:

$$\text{Var}[\mathbf{g}^{(j)}] = \sum_{k=1}^{K}\text{Var}[\mathbf{g}_k^{(j)}] \geq \frac{\sigma_{\text{shot}}^2}{2N_s}\sum_{k=1}^{K}\Omega_k^2$$

On the other hand, the variance of the explicitly added Gaussian noise is:

$$\text{Var}[\frac{1}{B}\mathbf{z}] = \frac{\sigma^2\Delta^2}{B^2}$$

where $\Delta = \frac{\lambda_{\max}-\lambda_{\min}}{2}\sqrt{\sum_{k=1}^{K}\Omega_k^2}$

Applying the Gaussian mechanism with sensitivity Δ/B then gives per-iteration (ε,δ)-DP Dwork (2006a):

$$\varepsilon = \frac{\Delta/B}{\sqrt{\frac{\sigma_{\text{shot}}^2}{2BN_s}\left(\sum_{k=1}^{K}\Omega_k^2\right)+\frac{\sigma^2\Delta^2}{B^2}}}\sqrt{2\ln\frac{1.25}{\delta}} = \frac{\sqrt{2\ln\frac{1.25}{\delta}}}{\sqrt{\frac{2B\sigma_{\text{shot}}^2}{N_s(\lambda_{\max}-\lambda_{\min})^2}+\sigma^2}}$$

□

Theorem 1. *Consider Algorithm 1 (Q-ShiftDP) applied to a private dataset S of size N, with mini-batch size B, sampling rate $q = B/N$, T training iterations, and per-sample ℓ_2-gradient sensitivity Δ (from Lemma 2).*

Let the desired total privacy guarantee be (ε,δ), for any $\delta \in (0,1)$ and for a sufficiently small ε (specifically, $\varepsilon < c_1 q^2 T$ for a universal constant c_1). The algorithm guarantees (ε,δ)-DP if the chosen explicit noise multiplier σ satisfies:

$$\sigma^2 \geq \max\left(0, C_{DP} - \frac{2B\sigma_{\text{shot}}^2}{N_s(\lambda_{\max}-\lambda_{\min})^2}\right) \tag{25}$$

where $C_{DP} = \left(c_2\frac{q\sqrt{T\log(1/\delta)}}{\varepsilon}\right)^2$ with c_2 is a constant.

Proof. Based on the works Balle et al. (2018); Abadi et al. (2016a), for an Subsampled Gaussian Mechanism (SGM) composed over T iterations with sampling rate q to be (ε,δ)-private, its per-iteration effective noise multiplier, σ_{total}, must satisfy:

$$\sigma_{\text{total}} \geq c_2\frac{q\sqrt{T\log(1/\delta)}}{\varepsilon}$$

for at least one universal constant c_2.

Each iteration of Q-ShiftDP is an instance of the SGM. Its total noise comes from two independent sources: the explicitly added noise (with multiplier σ) and the implicit quantum shot noise. Based on Lemma 3, the total effective noise multiplier, σ_{total}, is the sum of the squared individual multipliers:

$$\sigma_{\text{total}}^2 = \sigma^2 + \frac{2B\sigma_{\text{shot}}^2}{N_s(\lambda_{\max}-\lambda_{\min})^2}$$

As a result, to ensure the algorithm is (ε,δ)-private, we have:

$$\sigma^2 + \frac{2B\sigma_{\text{shot}}^2}{N_s(\lambda_{\max}-\lambda_{\min})^2} \geq \left(c_2\frac{q\sqrt{T\log(1/\delta)}}{\varepsilon}\right)^2$$

Because $\sigma \geq 0$, we have

$$\sigma^2 \geq \max\left(0,\ \left(c_2\frac{q\sqrt{T\log(1/\delta)}}{\varepsilon}\right)^2 - \frac{2B\sigma_{\text{shot}}^2}{N_s(\lambda_{\max}-\lambda_{\min})^2}\right)$$

. □

Theorem 2 (Bound on the Gradient Estimation Error for Q-ShiftDP)**.** *Let $\nabla C_S(\boldsymbol{\theta})$ be the true gradient over the full dataset S, and let $\tilde{\mathbf{g}}$ be the final noisy gradient estimate produced by Algorithm 1 (Q-ShiftDP) for a mini-batch of size B. The Mean Squared Error (MSE) of the estimator is bounded by:*

$$\mathbb{E}\left[\|\tilde{\mathbf{g}} - \nabla C_S(\boldsymbol{\theta})\|^2\right] \leq \frac{\Delta^2}{B}\left(1+\frac{1}{2N_s}\right) + \frac{K\sigma^2\Delta^2}{B^2} \tag{26}$$

where Δ is the per-sample ℓ_2-sensitivity from Lemma 2, N_s is the number of measurement shots, K is the number of parameters, and σ is the DP noise multiplier.

Proof. The total error can be decomposed into three independent, zero-mean sources: (1) artificial noise added for differential privacy, (2) variance from finite-shot estimation of gradients, and (3) variance from sampling a mini-batch instead of using the full dataset. We note that, in this proof, the symbol B is used to denote both the mini-batch itself and its size.

Let $\tilde{\mathbf{g}}$ be the final noisy gradient, $\hat{\mathbf{g}}_B$ be the shot-noise-only average gradient for the batch B, and $\nabla C_B(\boldsymbol{\theta})$ be the true (exact) average gradient for the batch B. We can decompose the error vector as:

$$\tilde{\mathbf{g}} - \nabla C_S(\boldsymbol{\theta}) = \underbrace{(\tilde{\mathbf{g}} - \hat{\mathbf{g}}_B)}_{\text{Artificial DP Noise}} + \underbrace{(\hat{\mathbf{g}}_B - \nabla C_B(\boldsymbol{\theta}))}_{\text{Estimator Variance}} + \underbrace{(\nabla C_B(\boldsymbol{\theta}) - \nabla C_S(\boldsymbol{\theta}))}_{\text{Sampling Variance}}$$

We have these three error terms are zero-mean and independent, so the expected squared norm of the sum is the sum of their expected squared norms (the cross-terms disappear).

$$\mathbb{E}\left[\|\tilde{\mathbf{g}} - \nabla C_S(\boldsymbol{\theta})\|^2\right] = \mathbb{E}\left[\|\tilde{\mathbf{g}} - \hat{\mathbf{g}}_B\|^2\right] + \mathbb{E}\left[\|\hat{\mathbf{g}}_B - \nabla C_B(\boldsymbol{\theta})\|^2\right] + \mathbb{E}\left[\|\nabla C_B(\boldsymbol{\theta}) - \nabla C_S(\boldsymbol{\theta})\|^2\right]$$

We bound each term separately.

Artificial DP Noise Error: The DP noise is $\tilde{\mathbf{g}} - \hat{\mathbf{g}}_B = \frac{1}{B}\mathbf{z}$, where $\mathbf{z} \sim \mathcal{N}(0, \sigma^2\Delta^2\mathbf{I}_K)$. Thus, we have:

$$\mathbb{E}\left[\|\tilde{\mathbf{g}} - \hat{\mathbf{g}}_B\|^2\right] = \frac{1}{B^2}\mathbb{E}\left[\|\mathbf{z}\|^2\right] = \frac{1}{B^2}(K \cdot \sigma^2\Delta^2) = \frac{K\sigma^2\Delta^2}{B^2}$$

Data Sampling Variance: This term measures the error from using a batch of size B instead of the full dataset of size N. We have:

$$\begin{aligned}
\mathbb{E}\left[\|\nabla C_B - \nabla C_S\|^2\right] &= \mathbb{E}\left[\left\|\frac{1}{B}\sum_{j=1}^{B}\nabla C_j - \nabla C_S\right\|^2\right] \\
&= \mathbb{E}\left[\left\|\frac{1}{B}\sum_{j=1}^{B}(\nabla C_j - \nabla C_S)\right\|^2\right] \\
&= \frac{1}{B^2}\sum_{j=1}^{B}\mathbb{E}\left[\|\nabla C_j - \nabla C_S\|^2\right] \\
&= \frac{1}{B}\mathbb{E}\left[|\nabla C_j - \nabla C_S|^2\right] \\
&\leq \frac{1}{B}\mathbb{E}\left[|\nabla C_j|^2\right] \text{(Because } \mathbb{E}[\nabla C_j] = \nabla C_S) \\
&\leq \frac{\Delta^2}{B} \text{(From Lemma 2)}
\end{aligned}$$

Estimator (Shot Noise) Variance: This term measures the error from estimating gradients with N_s shots. We have:

$$\mathbb{E}\left[\|\hat{\mathbf{g}}_B - \nabla C_B\|^2\right] = \mathbb{E}\left[\left\|\frac{1}{B}\sum_{j=1}^{B}(\mathbf{g}^{(j)} - \nabla C_j)\right\|^2\right] = \frac{1}{B^2}\sum_{j=1}^{B}\mathrm{Var}(\mathbf{g}^{(j)}) = \frac{1}{B}\mathrm{Var}(\mathbf{g})$$

where we used the i.i.d. nature of the estimators. The variance of a single estimated gradient vector $\hat{\mathbf{g}}$ is $\sum_{k=1}^{K}\mathrm{Var}(\hat{\mathbf{g}}_k)$. The variance of a single component is:

$$\mathrm{Var}(\mathbf{g}_k) = \mathrm{Var}\left(\frac{\Omega_k}{2}(r_+ - r_-)\right) = \frac{\Omega_k^2}{4}(\mathrm{Var}(r_+) + \mathrm{Var}(r_-))$$

The variance of the mean of N_s shots is $\mathrm{Var}(r_\pm) \leq \frac{\sigma_{\text{shot}}^2}{N_s}$.

$$\mathrm{Var}(\mathbf{g}_k) \leq \frac{\Omega_k^2}{4}\left(\frac{2\sigma_{\text{shot}}^2}{N_s}\right) = \frac{\Omega_k^2\sigma_{\text{shot}}^2}{2N_s}$$

Summing over all components:

$$\mathrm{Var}(\mathbf{g}) = \sum_{k=1}^{K} \mathrm{Var}(\mathbf{g}_k) \leq \frac{\sigma_{\text{shot}}^2}{2N_s} \sum_{k=1}^{K} \Omega_k^2$$

Using $\Delta = \frac{\lambda_{\max}-\lambda_{\min}}{2}\sqrt{\sum \Omega_k^2}$ and $\sigma_{\text{shot}}^2 \leq \frac{(\lambda_{\max}-\lambda_{\min})^2}{4}$, we get $\sum \Omega_k^2 = \frac{4\Delta^2}{(\lambda_{\max}-\lambda_{\min})^2}$. Substituting these in:

$$\mathrm{Var}(\hat{\mathbf{g}}) \leq \frac{(\lambda_{\max}-\lambda_{\min})^2/4}{2N_s} \cdot \frac{4\Delta^2}{(\lambda_{\max}-\lambda_{\min})^2} = \frac{\Delta^2}{2N_s}$$

So, the total estimator variance is bounded by $\frac{1}{B}\mathrm{Var}(\mathbf{g}) \leq \frac{\Delta^2}{2BN_s}$.

Combining the three error bounds gives the total MSE:

$$\mathbb{E}\left[\|\tilde{\mathbf{g}} - \nabla C_S(\boldsymbol{\theta})\|^2\right] \leq \frac{K\sigma^2\Delta^2}{B^2} + \frac{\Delta^2}{B} + \frac{\Delta^2}{2BN_s} = \frac{\Delta^2}{B}\left(1 + \frac{1}{2N_s}\right) + \frac{K\sigma^2\Delta^2}{B^2}$$

□

Theorem 3 (Lower Bound on Single-Shot Variance with Depolarizing Noise)**.** *The single-shot variance of a measurement of an observable O on this noisy state, $\eta_{\text{shot}}^2(x)$, is lower-bounded by:*

$$\eta_{\text{shot}}^2(x) \geq \alpha \cdot \sigma_{\text{uniform}}^2 \tag{27}$$

where $\sigma_{\text{uniform}}^2$ is the variance of the observable with respect to the maximally mixed state (I/d), a constant value given by:

$$\sigma_{\text{uniform}}^2 = \left(\frac{1}{d}\sum_i \lambda_i^2\right) - \left(\frac{1}{d}\sum_i \lambda_i\right)^2$$

Proof. Consider a quantum circuit that prepares an ideal state $\rho_{\text{ideal}}(x) = |\psi(x)\rangle\langle\psi(x)|$. Suppose this state passes through a global depolarizing channel $\mathcal{E}_\alpha$ of strength $\alpha \in [0,1]$ before measurement. The final (noisy) state is:

$$\rho_{\text{noisy}}(x) = \mathcal{E}_\alpha(\rho_{\text{ideal}}(x)) = (1-\alpha)\rho_{\text{ideal}}(x) + \alpha\frac{I}{d}$$

where I is the identity operator and d is the dimension of the Hilbert space.

First, we calculate the probability distribution of outcomes for the noisy state and then apply the definition of variance, $\eta_{\text{shot}}^2(x) = \mathbb{E}[\hat{a}^2] - (\mathbb{E}[\hat{a}])^2$.

Let $p_i^{\text{ideal}}(x) = \mathrm{Tr}(\rho_{\text{ideal}}(x)P_i)$ be the outcome probability for the ideal state. The probability of measuring outcome λ_i on the noisy state $\rho_{\text{noisy}}(x)$ is:

$$\begin{aligned}
p_i^{\text{noisy}}(x) &= \mathrm{Tr}(\rho_{\text{noisy}}(x)P_i) \\
&= \mathrm{Tr}\left(\left((1-\alpha)\rho_{\text{ideal}}(x) + \alpha\frac{I}{d}\right)P_i\right) \\
&= (1-\alpha)\mathrm{Tr}(\rho_{\text{ideal}}(x)P_i) + \frac{\alpha}{d}\mathrm{Tr}(IP_i) \\
&= (1-\alpha)p_i^{\text{ideal}}(x) + \frac{\alpha}{d}
\end{aligned}$$

where we have used $\mathrm{Tr}(P_i) = 1$ for a rank-1 projector. This shows that the noisy probability distribution is a convex combination of the ideal and uniform distributions.

Let $\mathbb{E}[\cdot]_{\text{noisy}}$, $\mathbb{E}[\cdot]_{\text{ideal}}$, and $\mathbb{E}[\cdot]_{\text{uniform}}$ denote expectations with respect to the noisy, ideal, and uniform distributions. The expectation values are linear combinations:

$$\mathbb{E}[\hat{a}]_{\text{noisy}} = \sum_i \lambda_i p_i^{\text{noisy}}(x) = (1-\alpha)\mathbb{E}[\hat{a}]_{\text{ideal}} + \alpha\mathbb{E}[\hat{a}]_{\text{uniform}}$$

$$\mathbb{E}[\hat{a}^2]_{\text{noisy}} = \sum_i \lambda_i^2 p_i^{\text{noisy}}(x) = (1-\alpha)\mathbb{E}[\hat{a}^2]_{\text{ideal}} + \alpha\mathbb{E}[\hat{a}^2]_{\text{uniform}}$$

Substituting these into the variance formula for the noisy state, $\eta^2_{\text{shot}}(x) = \mathbb{E}[\hat{a}^2]_{\text{noisy}} - (\mathbb{E}[\hat{a}]_{\text{noisy}})^2$, we have:

$$\begin{aligned}\eta^2_{\text{shot}}(x) &= \big((1-\alpha)\mathbb{E}[\hat{a}^2]_{\text{ideal}} + \alpha\mathbb{E}[\hat{a}^2]_{\text{uniform}}\big) - ((1-\alpha)\mathbb{E}[\hat{a}]_{\text{ideal}} + \alpha\mathbb{E}[\hat{a}]_{\text{uniform}})^2 \\ &= (1-\alpha)\mathbb{E}[\hat{a}^2]_{\text{ideal}} + \alpha\mathbb{E}[\hat{a}^2]_{\text{uniform}} \\ &\quad - \big((1-\alpha)^2(\mathbb{E}[\hat{a}]_{\text{ideal}})^2 + \alpha^2(\mathbb{E}[\hat{a}]_{\text{uniform}})^2 + 2\alpha(1-\alpha)\mathbb{E}[\hat{a}]_{\text{ideal}}\mathbb{E}[\hat{a}]_{\text{uniform}}\big)\end{aligned}$$

Recall that $\sigma^2 = \mathbb{E}[\hat{a}^2] - (\mathbb{E}[\hat{a}])^2$, so by replacing $\mathbb{E}[\hat{a}^2] = \sigma^2 + (\mathbb{E}[\hat{a}])^2$, we have:

$$\begin{aligned}\eta^2_{\text{shot}}(x) &= (1-\alpha)\sigma^2_{\text{ideal}}(x) + \alpha\sigma^2_{\text{uniform}} + (1-\alpha)(\mathbb{E}[\hat{a}]_{\text{ideal}})^2 - (1-\alpha)^2(\mathbb{E}[\hat{a}]_{\text{ideal}})^2 \\ &\quad + \alpha(\mathbb{E}[\hat{a}]_{\text{uniform}})^2 - \alpha^2(\mathbb{E}[\hat{a}]_{\text{uniform}})^2 - 2\alpha(1-\alpha)\mathbb{E}[\hat{a}]_{\text{ideal}}\mathbb{E}[\hat{a}]_{\text{uniform}} \\ &= (1-\alpha)\sigma^2_{\text{ideal}}(x) + \alpha\sigma^2_{\text{uniform}} + \alpha(1-\alpha)\left((\mathbb{E}[\hat{a}]_{\text{ideal}})^2 - 2\mathbb{E}[\hat{a}]_{\text{ideal}}\mathbb{E}[\hat{a}]_{\text{uniform}} + (\mathbb{E}[\hat{a}]_{\text{uniform}})^2\right) \\ &= (1-\alpha)\sigma^2_{\text{ideal}}(x) + \alpha\sigma^2_{\text{uniform}} + \alpha(1-\alpha)(\mathbb{E}[\hat{a}]_{\text{ideal}}(x) - \mathbb{E}[\hat{a}]_{\text{uniform}})^2\end{aligned}$$

Since all terms in the above equation are non-negative, we can drop the first and third terms to deliver the inequality:

$$\eta^2_{\text{shot}}(x) \geq \alpha \cdot \sigma^2_{\text{uniform}}$$

□

Theorem 4. *For a mini-batch of size B, let $\hat{\boldsymbol{\eta}}^2_B$ be a probabilistic estimator for the true sum of all variances, $\boldsymbol{\eta}^2_B = \sum_{j=1}^{B}\sum_{k=1}^{K}\sum_{\tau\in\pm}\eta^2_{j,k,\tau}$ such that it satisfies $\Pr(\boldsymbol{\eta}^2_B \geq \hat{\boldsymbol{\eta}}^2_B) \to 1-\beta$ for a chosen significance level β. Adaptive Q-ShiftDP calculates a noise multiplier for each batch B, σ^2_B, satisfying:*

$$\sigma^2_B \geq \max\left(0, C_{DP} - \frac{\Omega^2\hat{\boldsymbol{\eta}}^2_B}{4N_s\Delta^2}\right)$$

*where $C_{DP} = \left(c_2\frac{q\sqrt{T\log(1/\delta)}}{\varepsilon}\right)^2$ with c_2 is a constant. Adaptive Q-ShiftDP achieves $(\varepsilon, (1-\beta)\delta+\beta)$-**DP**.*

The proof of this result rely on supporting lemmas, which we state and prove below.

Lemma 4 (Multi-Group Variance Estimation)**.** *Let there be L independent groups, indexed by $i = 1, \ldots, L$. For each group i, let $y_{i,1}, \ldots, y_{i,n} \sim$ i.i.d. Y_i be n observations. Let η^2_i, $\bar{\eta}^2_i$, $\mu_{4,i}$ and $\bar{\mu}_{4,i}$ be the true variance, sample variance, true fourth central moment, and sample fourth central moment for the i-th group. Let $\eta^2_{batch} = \sum_{i=1}^{L}\eta^2_i$ and $\bar{\eta}^2_{batch} = \sum_{i=1}^{L}\bar{\eta}^2_i$. Assuming all $\mu_{4,i}$ are finite and given a significance level β with critical value z_β, then:*

$$P\left(\eta^2_{batch} \geq \bar{\eta}^2_{batch} - z_\beta\sqrt{\sum_{i=1}^{L}\frac{1}{n}\left(\bar{\mu}_{4,i} - (\bar{\eta}^2_i)^2\right)}\right) \to 1-\beta$$

as $L \to \infty, n \to \infty$.

Proof. First, we prove a sublemma for single-group case as follow:

Lemma 4.1 (Single-Group Variance Estimation)**.** *Let $y_1, \ldots, y_n$ be independent and identically distributed (i.i.d.) samples drawn from a distribution with true variance η^2 and a finite true fourth central moment μ_4. Let $\bar{\eta}^2$ and $\bar{\mu}_4$ be the sample variance and sample fourth central moment, respectively, computed from these observations. For a chosen significance level β, let z_β be the corresponding critical value of the standard normal distribution.*

Then, the following probabilistic lower bound holds:

$$\Pr\left(\eta^2 \geq \bar{\eta}^2 - z_\beta\sqrt{\frac{1}{n}\left(\bar{\mu}_4 - (\bar{\eta}^2)^2\right)}\right) \to 1-\beta$$

as the number of samples $n \to \infty$.

Proof. The proof is based on the asymptotic distribution of the sample variance. Let the true variance be η^2 and the true fourth central moment be μ_4. The corresponding sample estimators are $\bar{\eta}^2$ and $\bar{\mu}_4$.

By the Central Limit Theorem and Slutsky's Theorem, the standardized statistic T_n is known to converge in distribution to a standard normal random variable $Z \sim \mathcal{N}(0,1)$:

$$T_n = \frac{\sqrt{n}(\bar{\eta}^2 - \eta^2)}{\sqrt{\bar{\mu}_4 - (\bar{\eta}^2)^2}} \xrightarrow{d} Z$$

By definition of the one-sided critical value z_β, we know that for a standard normal variable Z, $\Pr(Z \leq z_\beta) = 1-\beta$. Since T_n converges in distribution to Z, we can state the asymptotic probability:

$$\Pr(T_n \leq z_\beta) \to 1 - \beta \quad \text{as } n \to \infty$$

Substituting the definition of T_n, we have:

$$\Pr\left(\frac{\sqrt{n}(\bar{\eta}^2 - \eta^2)}{\sqrt{\bar{\mu}_4 - (\bar{\eta}^2)^2}} \leq z_\beta\right) \to 1 - \beta$$

We now algebraically manipulate the inequality inside the probability to isolate the true parameter η^2. Each of the following steps is equivalent.

$$\begin{aligned}
\frac{\sqrt{n}(\bar{\eta}^2 - \eta^2)}{\sqrt{\bar{\mu}_4 - (\bar{\eta}^2)^2}} &\leq z_\beta \\
\bar{\eta}_n^2 - \eta^2 &\leq \frac{z_\beta}{\sqrt{n}}\sqrt{\bar{\mu}_4 - (\bar{\eta}^2)^2} \\
-\eta^2 &\leq -\bar{\eta}^2 + \frac{z_\beta}{\sqrt{n}}\sqrt{\bar{\mu}_4 - (\bar{\eta}^2)^2} \\
\eta^2 &\geq \bar{\eta}^2 - \frac{z_\beta}{\sqrt{n}}\sqrt{\bar{\mu}_4 - (\bar{\eta}^2)^2}
\end{aligned}$$

We can rewrite the final term by bringing the $1/\sqrt{n}$ factor inside the square root:

$$\eta^2 \geq \bar{\eta}^2 - z_\beta\sqrt{\frac{1}{n}\left(\bar{\mu}_4 - (\bar{\eta}^2)^2\right)}$$

Since this inequality is algebraically equivalent to our starting point ($T_n \leq z_\beta$), the probability of it being true must have the same limit. Therefore, we have proven the statement in the sublemma. □

The proof extends the single-group case by applying the Central Limit Theorem to the sum of the sample variances, $\bar{\eta}^2_{\text{batch}} = \sum_{i=1}^{L} \bar{\eta}_i^2$.

From the original lemma, we know the asymptotic variance of each individual sample variance $\bar{\eta}_i^2$ as $n \to \infty$ is given by:

$$\text{Var}(\bar{\eta}_i^2) \approx \frac{1}{n}(\mu_{4,i} - (\eta_i^2)^2)$$

We consider $\bar{\eta}^2_{\text{batch}}$ as a sum of L independent random variables $(\bar{\eta}_1^2, \ldots, \bar{\eta}_L^2)$. By the Central Limit Theorem, for a large number of groups L, the distribution of $\bar{\eta}^2_{\text{batch}}$ approaches a normal distribution. The mean of this sum is the sum of the means, $E[\bar{\eta}^2_{\text{batch}}] = \sum_{i=1}^{L} E[\bar{\eta}_i^2] = \sum_{i=1}^{L} \eta_i^2 = \eta^2_{\text{batch}}$. The variance of the sum is the sum of the variances, $\text{Var}(\bar{\eta}^2_{\text{batch}}) = \sum_{i=1}^{L} \text{Var}(\bar{\eta}_i^2)$.

Thus, the standardized sum converges in distribution to a standard normal variable $Z \sim \mathcal{N}(0,1)$. By Slutsky's Theorem, we can replace the true parameters in the variance term with their consistent sample estimators without changing the asymptotic distribution. This gives the test statistic T_L:

$$T_L = \frac{\bar{\eta}^2_{\text{batch}} - \eta^2_{\text{batch}}}{\sqrt{\sum_{i=1}^{L} \frac{1}{n}\left(\bar{\mu}_{4,i} - (\bar{\eta}_i^2)^2\right)}} \xrightarrow{d} Z$$

By definition of the critical value z_β, we have $\Pr(T_L \leq z_\beta) \to 1-\beta$ as $L \to \infty$. We now isolate η^2_{batch} within the probability statement.

$$\frac{\bar{\eta}^2_{\text{batch}} - \eta^2_{\text{batch}}}{\sqrt{\sum_{i=1}^{L} \frac{1}{n}\left(\bar{\mu}_{4,i} - (\bar{\eta}_i^2)^2\right)}} \leq z_\beta$$

$$\bar{\eta}^2_{\text{batch}} - \eta^2_{\text{batch}} \leq z_\beta \sqrt{\sum_{i=1}^{L} \frac{1}{n}\left(\bar{\mu}_{4,i} - (\bar{\eta}_i^2)^2\right)}$$

$$-\eta^2_{\text{batch}} \leq -\bar{\eta}^2_{\text{batch}} + z_\beta \sqrt{\sum_{i=1}^{L} \frac{1}{n}\left(\bar{\mu}_{4,i} - (\bar{\eta}_i^2)^2\right)}$$

$$\eta^2_{\text{batch}} \geq \bar{\eta}^2_{\text{batch}} - z_\beta \sqrt{\sum_{i=1}^{L} \frac{1}{n}\left(\bar{\mu}_{4,i} - (\bar{\eta}_i^2)^2\right)}$$

Since the final inequality is algebraically equivalent to $T_L \leq z_\beta$, its probability must also converge to $1-\beta$, which completes the proof. □

Lemma 5. *Let $\mathcal{M}$ be a randomized mechanism that operates in one of two modes. With probability γ, it operates in a "Good Mode," where its execution is equivalent to a mechanism $\mathcal{M}_{good}$ that is (ε, δ)-differentially private. With probability $1-\gamma$, it operates in a "Bad Mode," where no privacy guarantee holds. Then, the overall mechanism $\mathcal{M}$ is $(\varepsilon, \gamma\delta + (1-\gamma))$-differentially private.*

Proof. Let x and y be any two adjacent databases, and let S be any set of possible outcomes. The mechanism's choice of mode is determined by its internal randomness, which is resampled for each run. Therefore, the events of being in a good mode are independent for the runs on x and y.

Let $E_{good,x}$ be the event that the mechanism operates in the Good Mode when run on database x. Let $E_{good,y}$ be the event that the mechanism operates in the Good Mode when run on database y. We are given $\Pr(E_{good,x}) = \Pr(E_{good,y}) = \gamma$.

By the law of total probability, we can write the probability of an output for database x:

$$\Pr[\mathcal{M}(x) \in S] = \Pr[\mathcal{M}(x) \in S | E_{good,x}] \cdot \gamma + \Pr[\mathcal{M}(x) \in S | \neg E_{good,x}] \cdot (1-\gamma) \tag{28}$$

We now bound the two terms on the right-hand side. The mechanism that runs conditioned on the Good Event is, by assumption, the same (ε, δ)-DP mechanism, $\mathcal{M}_{good}$, regardless of the input. Thus, its behavior on x (conditioned on $E_{good,x}$) can be compared to its behavior on y (conditioned on $E_{good,y}$):

$$\Pr[\mathcal{M}(x) \in S | E_{good,x}] \leq e^\varepsilon \Pr[\mathcal{M}(y) \in S | E_{good,y}] + \delta$$

For the Bad Event, we can only bound the probability by 1:

$$\Pr[\mathcal{M}(x) \in S | \neg E_{good,x}] \leq 1$$

Substituting these bounds into Equation 28:

$$\begin{aligned} \Pr[\mathcal{M}(x) \in S] &\leq (e^\varepsilon \Pr[\mathcal{M}(y) \in S | E_{good,y}] + \delta) \cdot \gamma + 1 \cdot (1-\gamma) \\ &\leq \gamma \cdot e^\varepsilon \Pr[\mathcal{M}(y) \in S | E_{good,y}] + \gamma\delta + (1-\gamma) \end{aligned} \tag{29}$$

To complete the proof, we relate the conditional probability for database y back to the total probability. By the law of total probability for the run on y:

$$\Pr[\mathcal{M}(y) \in S] = \Pr[\mathcal{M}(y) \in S | E_{good,y}] \cdot \gamma + \Pr[\mathcal{M}(y) \in S | \neg E_{good,y}] \cdot (1-\gamma)$$

Since probabilities are non-negative, it follows that:

$$\Pr[\mathcal{M}(y) \in S] \geq \Pr[\mathcal{M}(y) \in S | E_{good,y}] \cdot \gamma$$

$$\implies \Pr[\mathcal{M}(y) \in S | E_{good,y}] \leq \frac{\Pr[\mathcal{M}(y) \in S]}{\gamma}$$

Finally, we substitute this back into our main inequality (Equation 29):

$$\begin{aligned}\Pr[\mathcal{M}(x) \in S] &\leq \gamma \cdot e^{\varepsilon} \left(\frac{\Pr[\mathcal{M}(y) \in S]}{\gamma} \right) + \gamma\delta + (1-\gamma) \\ &\leq e^{\varepsilon} \Pr[\mathcal{M}(y) \in S] + \gamma\delta + (1-\gamma)\end{aligned}$$

This expression matches the definition of $(\varepsilon, \gamma\delta + (1-\gamma))$-differential privacy. □

We are now ready to prove the main theorem for the advanced adaptive mechanism, Adaptive Q-ShiftDP. The privacy of the algorithm for a given batch depends on an aggregated estimator, $\hat{\boldsymbol{\eta}}_B^2$, which provides a probabilistic lower bound for the true sum of all shot-noise variances $\boldsymbol{\eta}_B^2$ in that batch.

The algorithm is in a "Good Mode" if this estimator $\hat{\boldsymbol{\eta}}_B^2$ is a correct lower bound on $\boldsymbol{\eta}_B^2$. Lemma 4 directly provides the probability for this single event:

$$\Pr_{\text{good}} = \Pr(\boldsymbol{\eta}_B^2 \geq \hat{\boldsymbol{\eta}}_B^2) \to 1 - \beta$$

Consequently, the algorithm is in a "Bad Mode" with the remaining probability, $\Pr_{\text{bad}} = 1 - \Pr_{\text{good}} \to \beta$.

By applying Lemma 5 with these probabilities, we conclude that the adaptive mechanism is $(\varepsilon, \Pr_{\text{good}} \cdot \delta + \Pr_{\text{bad}})$-DP. This yields the final privacy guarantee of:

$$(\varepsilon, (1-\beta)\delta + \beta)\text{-DP}$$

C Related Works

In this section, we first review the literature on differential privacy in classical machine learning. Then, we discuss existing works that extend this concept to quantum machine learning.

C.1 Differential Privacy in Classical Machine Learning

The goal of DP in Classical ML is to modify the ML training process so the final model is differentially private, i.e, the model would not be sufficiently different, no matter whether a particular instance was or was not included in the training data. Early works focused on modifying already trained model weights, a strategy often called output perturbation. Originating from Dwork's definition of differential privacy Dwork (2006b), these methods add noise proportional to the sensitivity of the final model parameters. However, computing tight sensitivity bounds is tractable only for relatively simple models such as linear regression due to the complex dependence of weights on the data Zhang et al. (2012). For example, Chaudhuri and Monteleoni (2008) proposed an algorithm that injects calibrated noise into trained logistic regression models to achieve differential privacy.

Another line of work achieves DP-Training by perturbing the loss function itself with noise, which is then optimized normally using SGD or other optimizers Chaudhuri et al. (2011); Kifer et al. (2012); Phan et al. (2016). Both noising weights and objective perturbation generally require strong convexity assumptions and convergence to a global optimum to guarantee privacy. Under these assumptions they can provide very strong privacy ($\varepsilon \approx 1$) with minimal accuracy loss. Subsequent studies have attempted to relax the requirements of global optimality Iyengar et al. (2019) or convexity Neel et al. (2019), but their privacy guarantees still depend on the loss function being bounded and Lipschitz continuity in the model weights. These methods remain computationally expensive and are therefore practical only for relatively small datasets.

Because most deep models are non-convex and are not trained to the global optimality due to time constraints and the computing cost, such loss perturbation techniques are rarely practical for modern deep learning. Gradient perturbation based methods are currently the most practical approach for ensuring rigorous privacy guarantees in non-convex settings such as large scale deep neural networks. The most widely adopted variant is differentially private stochastic gradient descent (DP-SGD) Song et al. (2013); Abadi et al. (2016b), which has become the standard method for differentially private training. Briefly, DP-SGD clips each per example gradient to a fixed

norm, adds Gaussian noise to the averaged gradients, and tracks the cumulative privacy loss using accounting techniques. Recent advances in DP-SGD focus on tightening privacy accounting bounds Mironov (2017); Dong et al. (2022), developing adaptive clipping techniques Chen et al. (2020); He et al. (2023) to improve the noise–utility trade-off, and introducing memory-efficient algorithms Li et al. (2021); Tang et al. (2024) to mitigate the large memory footprint caused by gradient accumulation.

C.2 Differential Privacy in Quantum Setting

With the growing popularity of quantum computing, the concept of differential privacy has naturally been extended to quantum machine learning (QML). The first formal notion of quantum differential privacy (QDP) was introduced by Zhou and Ying (2017), who defined it as a direct quantum analogue of classical differential privacy. Building on this foundation, Du et al. (2021) explored its practical relevance to QML. Specifically, they illustrate that intrinsic quantum noise can be harnessed to achieve QDP in quantum classifiers. For example, they study the depolarizing noise channel as a privacy-preserving mechanism, and then derive a precise relationship between the noise strength and the resulting (ε, δ)-QDP guarantee. Notably, they also established that this mechanism not only preserves the differential privacy but simultaneously improves the model's adversarial robustness.

Subsequently, Hirche et al. (2023) advanced the field by developing a comprehensive information-theoretic framework for QDP. In this work, the authors leverage tools such that quantum relative entropy and channel divergence to establish a rigorous and general foundation for analyzing privacy guarantees in quantum settings. This framework not only formalizes the definition of QDP but also connects it with fundamental quantities in quantum information theory. Building on these foundations, more recent studies Bai et al. (2024); Watkins et al. (2023); Song et al. (2025) have turned their attention to the practical implications of QDP under realistic quantum noise. In particular, these works investigate how different types of quantum noise channels, such as depolarizing, bit-flip, and phase-flip channels, contribute to the privacy budget.

Most existing works in quantum differential privacy (QDP) focus on injecting quantum noise into the forward pass to preserve privacy, which is conceptually equivalent to input perturbation in classical DP Lécuyer et al. (2018). However, in classical neural networks, perturbing the forward pass (i.e., inputs or outputs) often results in higher utility degradation than perturbing model's gradients Jarin and Eshete (2022). The same limitation may arise in QDP. Recently, a few studies Watkins et al. (2021); Chen and Chen (2025) have explored adding noise to quantum gradients. But, these approaches remain largely ad hoc, as they directly apply DP-SGD to quantum gradients without accounting for their unique properties, such as boundedness and inherent stochasticity. This can lead to injecting an excessive amount of artificial DP noise. Therefore, it is crucial to design methods that exploit the structural characteristics of quantum gradients to achieve stronger privacy–utility trade-offs.

D Experimental Settings

D.1 Datasets and Quantum Models.

We assess the performance and privacy–utility trade-off of Q-ShiftDP on three representative datasets from recent QML benchmark suites Bowles et al. (2024); Recio-Armengol et al. (2025) — Bars & Stripes, Binary Blobs, and Downscaled MNIST, which are specifically designed to offer a controlled and reliable test bed for evaluating quantum models. The Bars and Stripes dataset is a simple, translation-invariant benchmark consisting of 2D pixel grids representing either horizontal bars or vertical stripes. The Binary Blobs dataset is a bitstring version of the common Gaussian blobs dataset. Each sample is a 16-bit string generated by randomly selecting one of eight predefined bit patterns and independently flipping each bit with a 5% probability. This process creates a distribution with eight clearly separated, visually distinct classes. Finally, Downscaled MNIST provides a reduced-dimension digit classification task derived from the original MNIST dataset, preserving class-discriminative information while keeping the input size tractable for quantum circuits. Our classifier is a fully quantum 4-qubit neural network (QNN) implemented in PennyLane and trained using PyTorch. For Bars & Stripes, we employ a single Strongly Entangling Layer, while for Binary Blobs and Downscaled MNIST, we use a five-layer Strongly Entangling Layers ansatz. The first two basis probabilities are used for binary classification, with their log-probabilities passed to the NLL loss. This is equivalent to using two binary observables O_0 and O_1, one for each class, which simplifies implementation.

D.2 Impact of Hyperparameters on the Performance of Q-ShiftDP

Differentially private stochastic optimizers are known to be highly sensitive to hyperparameter selection Bu et al. (2023); Sander et al. (2023); even slight adjustments to the clipping norm or batch size can lead to substantial performance changes. In contrast to DP-SGD which must clip unbounded per-example gradients to guarantee privacy, Q-ShiftDP exploits the parameter-shift rule and the bounded spectrum of quantum observables to compute an exact analytic sensitivity bound (see Lemma 2), removing the need to manually tune a clipping norm. To investigate how batch size and learning rate affect Q-ShiftDP's performance, we perform a grid search over batch sizes $\{32, 64, 128, 256, 512\}$, learning rates $\{0.2, 0.1, 0.05, 0.01, 0.005\}$, and privacy budgets $\varepsilon \in \{0.1, 0.5, 1.0\}$. For simplicity, this analysis assumes analytic expectation values rather than a finite number of measurement shots.

Figure 5: Test accuracy on Bars and Stripes dataset using QNN with gradient privated under Q-ShiftDP by different batch size and learning rates.

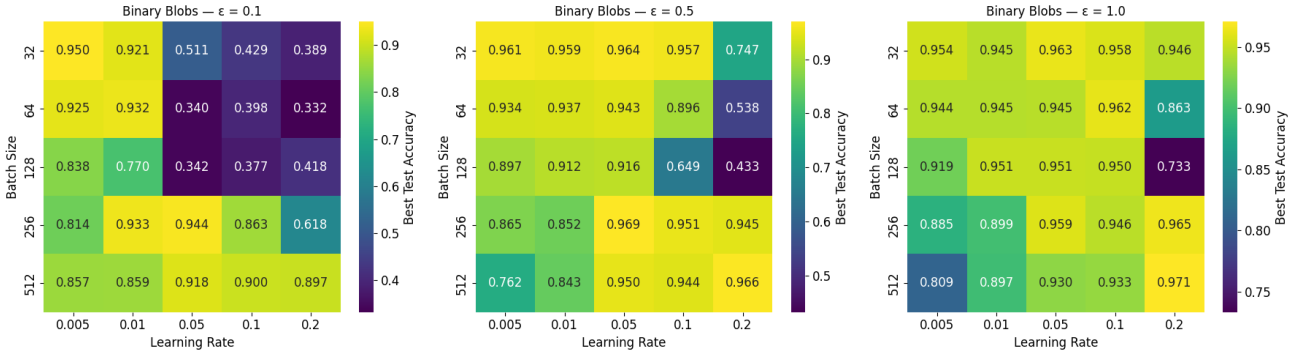

Figure 6: Test accuracy on Binary Blobs dataset of QNN with gradient privated using Q-ShiftDP by different batch size and learning rates

Figure 5, 6 and 7 shows that privacy budget, batch size, and learning rate strongly influence DP training on all datasets. Higher privacy budgets (larger ε) lead to better accuracy, while very small values (≈ 0.1) hurt performance. Large batche size usually perform best when privacy is strong ($\varepsilon = 0.1$ or 0.5), especially when paired with a large enough learning rate. This is consistent with prior DP-SGD literature showing that larger batch sizes yield more stable training across most privacy budgets De et al. (2022). For the remaining experiments, we use a batch size of $B = 512$ and a learning rate of $\mathrm{lr} = 0.2$, as this combination remains stable across all tested privacy budgets and datasets.

D.3 Impact of number of *shots* on performance of Q-ShiftDP

Table 1 shows the impact of number of *shots* on the performace of Q-ShiftDP. We observe that increasing the number of shots consistently improves model accuracy across all privacy budgets. With only 1,000 shots the gradient estimates are noisy, leading to reduced performance, while using 100,000 or infinitely many shots brings

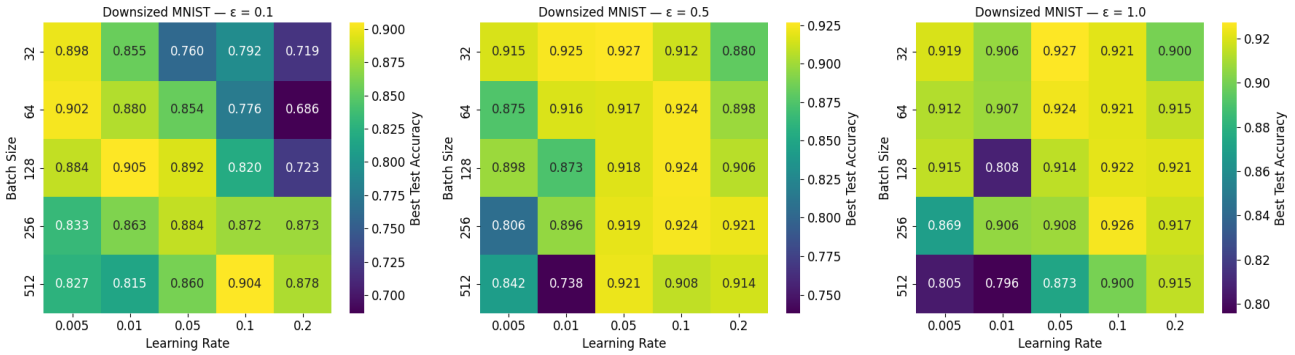

Figure 7: Test accuracy on Downsized MNIST dataset of QNN with gradient privated using Q-ShiftDP by different batch size and learning rates

ε	1,000 shots	10,000 shots	100,000 shots	∞ shots
1	0.86	0.91	0.91	0.950
0.5	0.83	0.89	0.90	0.940
0.1	0.80	0.86	0.89	0.925

Table 1: Accuracy of the QNN model on Bars and Stripes dataset for different privacy budgets ε and numbers of shots. The ∞ column represents the ideal case with no sampling noise.

the accuracy closer to the ideal non-sampling case. This effect is especially important when the privacy budget is tight (small ε), where the added stochasticity from few shots compounds with the differential privacy noise and harms learning.

D.4 Performance of Adaptive Q-ShiftDP

In this section, we provide experimental validation for our adaptive algorithm, **Adaptive Q-ShiftDP**. The central premise behind Adaptive Q-ShiftDP is that the non-adaptive Q-ShiftDP algorithm relies on a conservative, worst-case lower bound for the shot-noise variance (σ^2_{shot}), which can be trivially zero or based on a small depolarizing noise factor α. We hypothesize that the empirically measured variance is, in practice, significantly higher. To demonstrate the practical benefits of exploiting this, our evaluation is structured in three parts. First, we visualize the distribution of empirically measured shot-noise variances to confirm that they are consistently greater than the theoretical lower bound. Then, we quantify the resulting reduction in the required artificial noise (σ_B^2) when using our adaptive estimator. Finally, we compare the end-to-end performance of Adaptive Q-ShiftDP against the non-adaptive Q-ShiftDP in various depolarizing noise settings, showing that this reduction in noise translates directly to improved utility.

To begin, we set up experiments to measure the distribution of empirically observed shot-noise variances. We consider a PQC with $n = 4$ qubits and 12 parameterized gates. For each of 1000 random samples, consisting of input x and parameters θ, we simulate the exact single-shot variance according to Eq. 13. The simulations are performed with two observables: a binary observable with 2 eigenvalues, which is also used throughout our experiments on classification tasks, and a non-degenerate observable with 2^n eigenvalues, which is commonly used in regression tasks. Figure 8 shows the statistics of practical single-shot variances for the binary observable, where we observe that the practical variances remain far from the global variance lower bound for all values of α. The same trend is observed in Figure 9, which reports the statistics for the non-degenerate observable. These results highlight the necessity of estimating practical variances to reduce the amount of required artificial noise.

Next, we quantify the practical benefit of the adaptive mechanism by measuring the reduction in artificial noise it achieves. For each batch, we define the Artificial Noise Reduction(%) as the fraction of the baseline artificial noise that Adaptive Q-ShiftDP is able to eliminate by leveraging the inherent shot noise:

$$\text{Artificial Noise Reduction}(\%) = \frac{C_{\text{DP}} - \sigma_B^2}{C_{\text{DP}}} \times 100$$

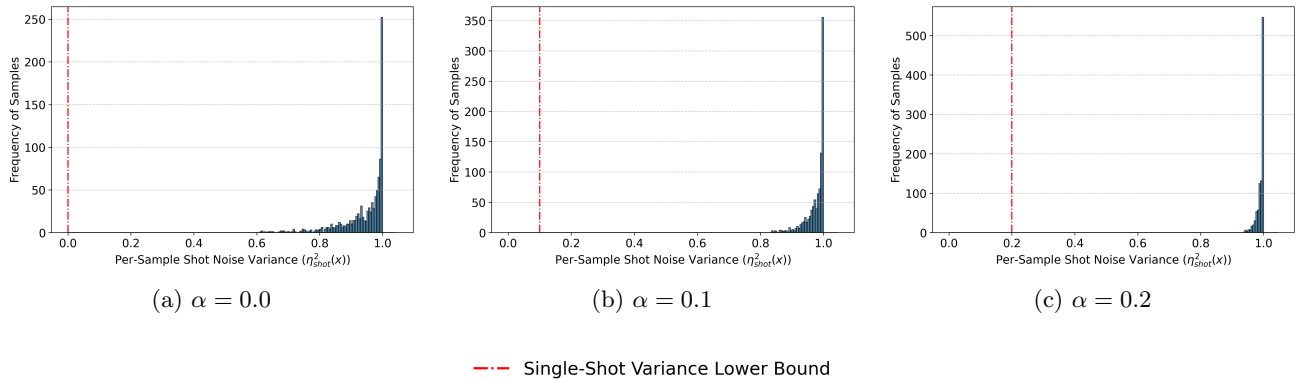

(a) $\alpha = 0.0$ (b) $\alpha = 0.1$ (c) $\alpha = 0.2$

Figure 8: Statistics of Practical Single-Shot Variances in Parameterized Quantum Circuits with a Binary Observable for Depolarizing Noise Levels $\alpha \in [0.0, 0.1, 0.2]$.

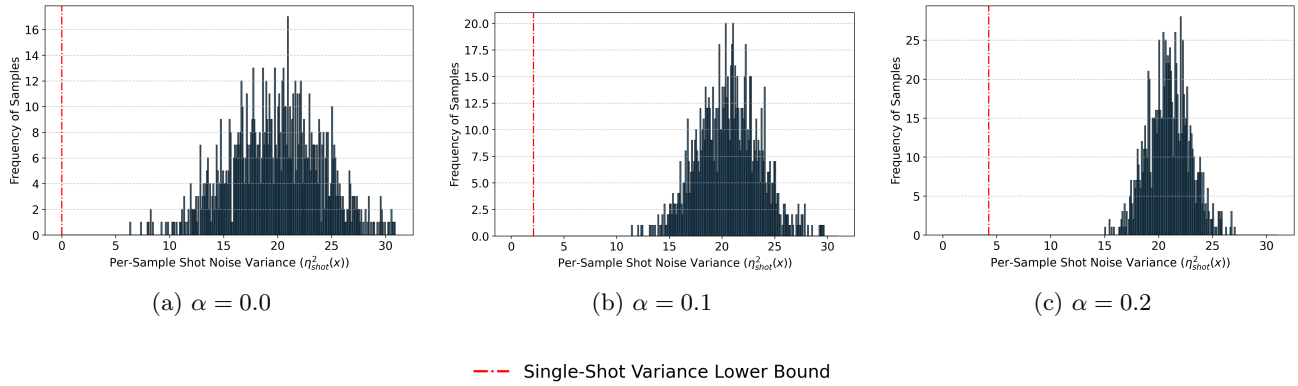

(a) $\alpha = 0.0$ (b) $\alpha = 0.1$ (c) $\alpha = 0.2$

Figure 9: Statistics of Practical Single-Shot Variances in Parameterized Quantum Circuits with a Non-Degenerate Observable for Depolarizing Noise Levels $\alpha \in [0.0, 0.1, 0.2]$.

Figure 10 plots the average reduction across various numbers of shots ($N_s \in \{10^2, 10^3, 10^4\}$) and batch sizes ($B \in \{64, 128, 256, 512\}$). First, the reduction is most dramatic when the number of shots is low. With only 100 shots, the inherent randomness from measurement is substantial, allowing the algorithm to reduce the required artificial noise by up to 16%. As the number of shots increases to 10^3 and 10^4, the gradient estimates become more precise (less shot noise), diminishing the contribution from inherent randomness. Consequently, the noise reduction drops to around 4% and less than 1%, respectively. Second, the relative impact of shot noise is greater for smaller batches. For instance, with 100 shots, a batch of 64 achieves a 16% reduction, whereas a larger batch of 512 sees a smaller reduction of less than 12%.

Finally, we demonstrate the practical efficiency of Adaptive Q-ShiftDP by comparing its performance with Q-ShiftDP under varying environmental conditions, represented by the strength of depolarizing noise $\alpha \in \{0, 0.1, 0.2\}$. For a fair comparison, both methods are trained with $\varepsilon = 1$, $\delta = \frac{1}{|S|}$ with $|S|$ is the dataset's size, and a small $\beta = 10^{-5} \ll \delta$, ensuring an equivalent level of differential privacy. The QML model used in this experiment consists of a single layer and is trained on the Bars and Stripes dataset ($|S| = 1000$) with 1000 shots and a batch size of 512. Figure 11 presents the comparison results. Specifically, it shows that across all three noise settings, Adaptive Q-ShiftDP achieves approximately 10% higher accuracy than Q-ShiftDP. These results highlight the effectiveness of reducing artificial noise in improving the practical performance of differentially private QML models.

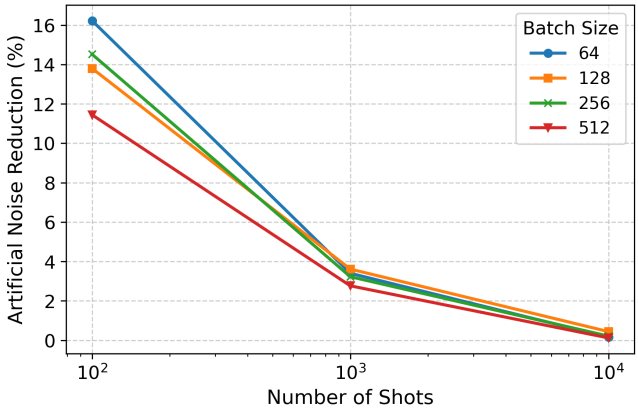

Figure 10: Percentage reduction in required artificial noise achieved by Adaptive Q-ShiftDP as a function of the number of shots (N_s) and batch size (B).

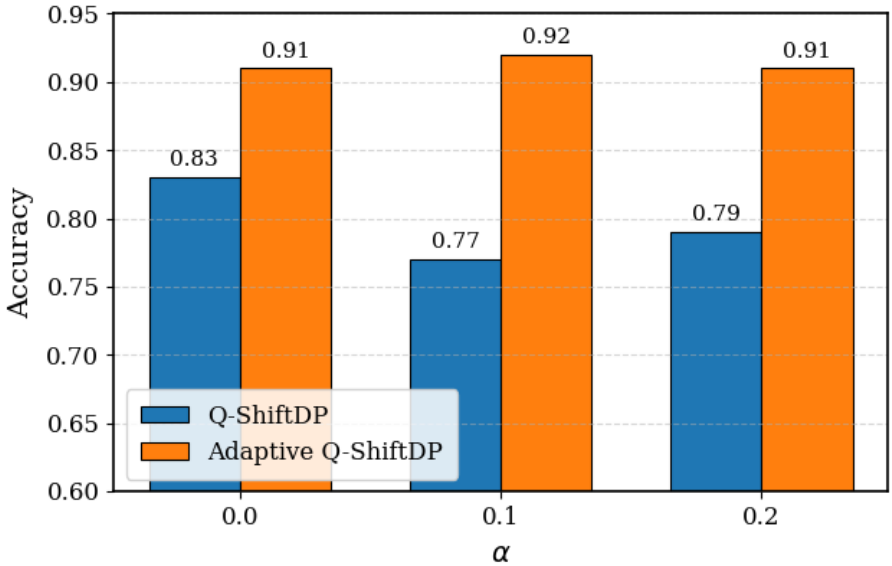

Figure 11: Accuracy comparison between Q-ShiftDP and Adaptive Q-ShiftDP under varying α values.

E Estimation of Single-Shot Measurement Variance by Samples (more details)

This section provides a detailed walkthrough of our adaptive algorithm, Adaptive Q-ShiftDP, presented in Algorithm 2. Adaptive Q-ShiftDP improves upon Q-ShiftDP by replacing the fixed, global artificial noise multiplier, σ^2, with a per-batch multiplier, σ_B^2, that adapts to the inherent randomness of the specific samples in each mini-batch (All changes from Q-ShiftDP are highlighted in red). This adaptation is driven by the per-batch variance estimator, $\hat{\boldsymbol{\eta}}_B^2$. This estimator, which provides a probabilistic lower bound on the true sum of variances in the batch, is constructed using Lemma 4. It, in turn, is built from two key statistics computed for each sample j and weight k in the batch: the sample variances, $\bar{\eta}^2_{j,k,\pm}$, and the sample fourth central moments, $\bar{\mu}_{4,j,k,\pm}$.

These statistics are calculated directly from the measurement outcomes. For the positive-shifted circuit of sample j with outcomes $\{\hat{a}_{j,k,1}, \ldots, \hat{a}_{j,k,N_s}\}$, we first compute the sample mean, $\bar{a}_{j,k} = \frac{1}{N_s}\sum_i \hat{a}_{j,k,i}$. The estimators are then formulated as follows:

- **Sample Variance ($\bar{\eta}^2_{j,k,+}$):**

$$\bar{\eta}^2_{j,k,+} = \frac{1}{N_s - 1}\sum_{i=1}^{N_s}(\hat{a}_{j,k,i} - \bar{a}_{j,k})^2 \tag{30}$$

- **Sample Fourth Central Moment ($\bar{\mu}_{4,j,k,+}$):**

$$\bar{\mu}_{4,j,k,+} = \frac{1}{N_s}\sum_{i=1}^{N_s}(\hat{a}_{j,k,i} - \bar{a}_{j,k})^4 \tag{31}$$

The same formulations are used for the negative-shifted circuit's outcomes, $\{\hat{b}_{j,k,i}\}_{i=1}^{N_s}$, to compute $\bar{\eta}^2_{j,k,-}$ and $\bar{\mu}_{4,j,k,-}$.

In the main manuscript, we state that given a DP budget (ε, δ) and assuming identical gate frequencies Ω, Q-ShiftDP can achieve (ε, δ)-DP guarantee by assigning a batch-specific noise level σ_B such that:

$$\sigma_B^2 \geq \max\left(0, C_{\mathrm{DP}} - \frac{\Omega^2 \boldsymbol{\eta}_B^2}{4N_s\Delta^2}\right) \tag{32}$$

where $\boldsymbol{\eta}_B^2$ is the true per-batch variance of the batch B. We explain this claim in more detail here. First, we remind that the privacy guarantee of the algorithm stems from the total effective noise variance per iteration (i.e., per batch), which must meet the threshold C_{DP} (see proof of Theorem 1). This total noise is the sum of two independent sources: the variance from the artificial noise we inject ($\sigma_B^2\Delta/B$) and the variance from the inherent quantum shot noise per batch. For a single gradient component $\mathbf{g}_k^{(j)}$, we have:

$$\mathrm{Var}[\mathbf{g}_k^{(j)}] = \mathrm{Var}\left[\frac{\Omega}{2}(r_{j,k,+} - r_{j,k,-})\right] = \frac{\Omega^2}{4N_s}\left(\eta^2_{j,k,+} + \eta^2_{j,k,-}\right)$$

Thus, the shot-noise variance for the entire batch is the sum of the variances of all $B \times K$ components:

$$\begin{aligned}\mathrm{Var}\left[\sum_{j=1}^{B}\mathbf{g}^{(j)}\right] &= \sum_{j=1}^{B}\sum_{k=1}^{K}\mathrm{Var}[\mathbf{g}_k^{(j)}] \\ &= \frac{\Omega^2}{4N_s}\sum_{j=1}^{B}\sum_{k=1}^{K}\left(\eta^2_{j,k,+} + \eta^2_{j,k,-}\right) \\ &= \frac{\Omega^2}{4N_s}\boldsymbol{\eta}_B^2\end{aligned}$$

Following the proof of Lemma 3, we have the total gradient variance per batch as:

$$\mathrm{Var}[\tilde{\mathbf{g}}] = \frac{\Omega^2}{4N_sB^2}\boldsymbol{\eta}_B^2 + \frac{\sigma_B^2\Delta^2}{B^2}$$

With the sensitivity Δ/B Dwork (2006a), we have the total noise multiplier per batch as:

$$\sigma^2_{\mathrm{total}} = \sigma_B^2 + \frac{\Omega^2\boldsymbol{\eta}_B^2}{4N_s\Delta^2}$$

Following the proof of Theorem 1,, (ε, δ)-DP is guaranteed when $\sigma^2_{\mathrm{total}} \geq C_{DP}$ or equivalently $\sigma_B^2 \geq \max\left(0, C_{\mathrm{DP}} - \frac{\Omega^2\boldsymbol{\eta}_B^2}{4N_s\Delta^2}\right)$.

Another important point is that $\boldsymbol{\eta}_B^2$ is unknown *a priori*. Therefore, in Adaptive Q-ShiftDP, we estimate it using $\hat{\boldsymbol{\eta}}_B^2$. By setting

$$\sigma_B^2 \leftarrow \max\left(0,\, C_{\mathrm{DP}} - \frac{\Omega^2\hat{\boldsymbol{\eta}}_B^2}{4N_s\Delta^2}\right),$$

Adaptive Q-ShiftDP ensures $(\varepsilon,\, (1-\beta)\delta + \beta)$-differential privacy with β is the significance level, as stated in Theorem 4.

Algorithm 2: Adaptive Differentially Private Parameter-Shift Rule (Adaptive Q-ShiftDP)

Input: Data encoding $x \mapsto |\psi_0(x)\rangle$, Parametric unitary $\hat{U}(\boldsymbol{\theta}) = \prod_{k=1}^{K} e^{i\theta_k G_k}$ with identical frequency Ω, Set of observables $\{\hat{O}_y\}$ with eigenvalues in $[\lambda_{\min}, \lambda_{\max}]$, Private dataset S, Learning rate $\mathbf{lr}$, NumSteps T, Batch size B, Ideal DP budget (ε, δ), significance level β.

Output: Trained parameters $\boldsymbol{\theta}_T$.

```
// Define sensitivity based on the theoretical maximum L2 norm of the quantum gradient
```

$\Delta \leftarrow \frac{\lambda_{\max}-\lambda_{\min}}{2}\sqrt{\sum_{k=1}^{K}\Omega_k^2}$

Initialize parameters $\boldsymbol{\theta}_0$ randomly

for $t \leftarrow 0$ **to** $T-1$ **do**

  Sample a mini-batch $B \subset S$

  Initialize an empty list for gradient estimates: $\mathcal{G} \leftarrow []$

  foreach *sample* $(x_j, y_j) \in B$ **do**

    Initialize vector: $\mathbf{g}^{(j)} \leftarrow \mathbf{0} \in \mathbb{R}^K$

    for $k \leftarrow 1$ **to** K **do**

      Define shift amount: $s \leftarrow \pi/(2\Omega_k)$

      $|\phi\rangle \leftarrow (\prod_{l=1}^{k-1}\hat{U}_l(\theta_l))|\psi_0(x_j)\rangle$ `// Define intermediate state`

      $\hat{A} \leftarrow (\prod_{l=k+1}^{K}\hat{U}_l(\theta_l))^\dagger \hat{O}_{y_j}(\prod_{l=k+1}^{K}\hat{U}_l(\theta_l))$ `// Define intermediate operator`

      Run positive-shifted circuit $\hat{A}\hat{U}_k(\theta_k+s)|\phi\rangle$ N_s times, collect outcomes $\{\hat{a}_{j,k,i}\}_{i=1}^{N_s}$

      $r_+ \leftarrow \frac{1}{N_s}\sum_{i=1}^{N_s}\hat{a}_{j,k,i}$ `// Estimator for positive shift`

      Run negative-shifted circuit $\hat{A}\hat{U}_k(\theta_k-s)|\phi\rangle$ N_s times, collect outcomes $\{\hat{b}_{j,k,i}\}_{i=1}^{N_s}$

      $r_- \leftarrow \frac{1}{N_s}\sum_{i=1}^{N_s}\hat{b}_{j,k,i}$ `// Estimator for negative shift`

      Estimate sample variances $\bar{\eta}^2_{j,k,\pm}$ and sample fourth central moments $\bar{\mu}_{4,j,k,\pm}$. from outcomes using Eq 30 and 31.

      $\mathbf{g}_k^{(j)} \leftarrow \frac{\Omega_k}{2}(r_+ - r_-)$ `// Parameter-shift rule estimator`

    Append $\mathbf{g}^{(j)}$ to $\mathcal{G}$

  $\hat{\mathbf{g}}_B \leftarrow \sum_{\mathbf{g}^{(j)}\in\mathcal{G}}\mathbf{g}^{(j)}$

  `// --- Adaptive Step (Per-Batch) ---`

  $\bar{\boldsymbol{\eta}}_B^2 \leftarrow \sum_{j=1}^{B}\sum_{k=1}^{K}\sum_{\tau\in\{+,-\}}\bar{\eta}^2_{i,\tau}$

  $\hat{\boldsymbol{\eta}}_B^2 \leftarrow \bar{\boldsymbol{\eta}}_B^2 - z_\beta\sqrt{\sum_{j=1}^{B}\sum_{k=1}^{K}\sum_{\tau\in\pm}\frac{1}{N_s}\left(\bar{\mu}_{4,j,k,\tau} - (\bar{\eta}^2_{j,k,\tau})^2\right)}$

  $C_{\mathrm{DP}} \leftarrow \left(c_2\frac{q\sqrt{T\log(1/\delta)}}{\varepsilon}\right)^2$

  $\sigma_B^2 \leftarrow \max\left(0, C_{\mathrm{DP}} - \frac{\Omega^2\hat{\boldsymbol{\eta}}_B^2}{4N_s\Delta^2}\right)$

  Sample noise: $\mathbf{z} \sim \mathcal{N}(0, \sigma_B^2\Delta^2\mathbf{I}_K)$

  $\tilde{\mathbf{g}} \leftarrow \frac{1}{B}(\hat{\mathbf{g}}_B + \mathbf{z})$

  $\boldsymbol{\theta}_{t+1} \leftarrow \boldsymbol{\theta}_t - \mathbf{lr}\cdot\tilde{\mathbf{g}}$

return $\boldsymbol{\theta}_T$