ADVERSARIAL ROBUSTNESS BASED ON RANDOMIZED SMOOTHING IN QUANTUM MACHINE LEARNING

Anonymous authors
Paper under double-blind review

ABSTRACT

We present an end-to-end Quantum Machine Learning algorithm that encodes a classical input into a Quantum Computing state and provides a certified radius for a base classifier, with robustness guarantees based on randomized smoothing - current state-of-the-art defense against adversarial attacks. Classically, the number of samples, also the number of queries to the classifier, scale with $O(1/\epsilon^2)$ where ϵ is the desired error bound in the expected value of the probability measure ρ defined over the randomized smoothing neighborhood. Our algorithm is designed to solve the same problem for a Quantum Computing classifier. We prove that the number of queries to the classifier is $O(1/\epsilon)$ for the same confidence and error bound. We also present the unitary circuit corresponding to the quantum randomized smoothing algorithm, and the state preparation methods and circuits for smoothing distributions used to defend against common adversaries - modeled using l_0, l_1, l_2 norms, and other metrics. The results of the comparison between the classical and simulation of the quantum algorithm are also discussed.

1 INTRODUCTION

Machine Learning (ML) models have become ubiquitous over the last decade. There has also been a massive interest in Quantum Computing (QC) and Quantum Machine Learning (QML) (Aaronson, 2015), with algorithms like Shor's factoring (Shor, 1999), HHL (Harrow et al., 2009) algorithm for solving a linear system of equations, etc providing an exponential speedup over their classical counterparts. There exists another class of QC algorithms with a polynomial speedup - Grover's algorithm (Grover, 1996) for random database search, Quantum Amplitude Estimation (QAE) (Brassard et al., 2002) for counting problem, Bernstein-Vazirani (Bernstein & Vazirani, 1997) algorithm for parity problem, etc.

Shortcomings of classical ML algorithms against malicious actors is a widely-studied sub-domain (Goodfellow et al., 2014; Madry et al., 2017) of ML. Common attack vectors include data poisoning, backdoor attacks, and adversarial attacks. Adversarial attacks can easily trick well-trained classifiers into misclassifying an input perturbed by a small, usually imperceptible, noise. QML algorithms are prone to the same problems (Weber et al., 2021; Liao et al., 2021; Guan et al., 2020; Lu et al., 2020; Ren et al., 2022).

Popular methods for adversarial defense (Madry et al., 2017) find it difficult to train large robust classifiers, which are essential to solving real-world problems at the scale of ImageNet(Deng et al., 2009) or larger. These methods do not offer certifiable guarantees of robustness, even when they work well in practice.

Randomized smoothing is a state-of-the-art method that offers provable robustness against adversarial attacks, without any assumptions about the underlying classifier. The defense works by aggregating a classifier's output in a region around the input - henceforth called the smoothing neighborhood - and computing the average probability of a class ρ_c. It is prohibitively expensive to compute the exact value of ρ_c over the smoothing neighborhood since the number of points is exponential in the input dimension. In practice, Monte Carlo sampling algorithms are used to estimate ρ_c. Typically, randomized smoothing for adversarial robustness (Cohen et al., 2019; Yang et al., 2020; Lee et al., 2019) requires $N_{classical} \approx 10^5 - 10^6$ samples from the smoothing neighborhood.

In this paper, we discuss a purely QC approach to implementing randomized smoothing by using an orthogonal representation for the input space and use existing formalism for the Quantum Counting problem (Brassard et al., 2002). We create a superposition of the smoothing neighborhood of the input image and use our quantum circuit to output the average probability of prediction for a class ρ_c. This Quantum Adversarial Robustness (QuAdRo) algorithm provably needs $N_{quantum} = O(\sqrt{N}_{classical})$ calls to the base classifier oracle.

Our novel contribution is showing how QC subroutines like Quantum Amplitude Estimation (QAE) can be used in ML. Using the same theoretical framework as prior art (Cohen et al., 2019; Yang et al., 2020; Lee et al., 2019), we prove a quadratic speedup for the same guarantees of the robustness for QuAdRo. We also design qubit state encoding and state preparation circuits for l_0, l_1, and l_2 norm adversaries, and provide results from the simulation of the algorithm in Section 6.

2 Related Work

2.1 Randomized Smoothing

Randomized smoothing (Cohen et al., 2019; Yang et al., 2020) method has achieved provable robustness against adversarial attacks. Given an input, one can define a smoothing neighborhood based on the threat model of the adversary described by l_p norm and scale parameter λ. Such a robust model outputs the most likely class in the smoothing neighborhood returned by a base classifier, and this output is stable against l_p perturbations. Cohen et al. (2019) first proved tight robustness guarantees for l_2 norm adversary using Gaussian smoothing. Later, Yang et al. (2020) provided guarantees for a larger set of adversaries and smoothing distributions, except l_0 norm which was provided by Lee et al. (2019).

2.2 Qubit State Preparation

Qubits are logical units of information for Quantum Computers, equivalent to bits in classical computers. Any QC device is made up of qubits that have the following two properties - superposition and entanglement. Superposition refers to a qubit's ability to exist in multiple states at the same time, while entanglement refers to the ability of multiple qubits to exist in a shared state such that an operation on one qubit also affects the state of another qubit instantaneously, without any additional transfer of information.

There are numerous methods for encoding information into qubit states, often optimized for the target problem. For example, Novel Enhanced Quantum Representation (NEQR) (Zhang et al., 2013) and Flexible Representation of Quantum Images (FRQI) (Le et al., 2011) used for QC Image algorithms differ from variational heuristics for calculating molecular energies like Unitary Coupled-Cluster ansatz (Romero et al., 2018). Popular encoding methods used in QML applications use the amplitude of a quantum state $|i\rangle$ as the representation of input vector element $x[i]$ to be encoded. This representation is really efficient but has the drawback that quantum search, amplitude estimation, etc cannot be applied to such qubit states, due to lack of orthogonality.

A number of distributions can be prepared as a superposed qubit state(Rattew & Koczor, 2022) - Log concave distributions (Grover & Rudolph, 2002), Uniform distribution using Quantum Fourier Transform (Deutsch, 1985), etc. Distribution parameters can be modified, either during state preparation, or using circuits like QADD (Koch et al., 2022) mid-circuit. QFT and inverse QFT, in particular, are a common pair of pre- and post-processing circuits that concentrate information from a superposition via a Fourier transform.

2.3 Grover's Search Algorithm

Given a boolean objective function f defined over unstructured space S of size N such that $\exists x \in S : f(x) = 1$ and a QC Oracle O operator $O\,|x\rangle\,|y\rangle = |x\rangle\,|y \bigoplus f(x)\rangle$, Grover's Search algorithm (Grover, 1996) allows searching for x in $O(\sqrt{N})$ calls to O. Using state preparation subroutine U such that $U\,|0\rangle = |\psi\rangle$, the Grover diffusion operator

$$G = (2 * |\psi\rangle\,\langle\psi| - \mathbb{I})O \tag{1}$$

is applied repeatedly to concentrate the set of desired outcomes. The result of measurement of a concentrated state is a value $x \in S$ such that $f(x) = 1$ with an arbitrarily high probability. The original formulation requires $|\psi\rangle$ to be a uniform superposition of all values in the space, but later the same results were extended to arbitrary unitary operations U (Biron et al., 1999). No QC unstructured search algorithm can perform better than $O(\sqrt{N})$ (Boyer et al., 1998).

2.4 Quantum Amplitude Estimation

An alternate formulation of Grover's search algorithm into a counting problem allows us to use the same number of oracle queries to estimate the number of solutions in S for $f(x) = 1$ using Quantum Amplitude Estimation (QAE) (Brassard et al., 2002; Boyer et al., 1998). Variants of QAE (Grinko et al., 2021; Aaronson & Rall, 2020; Suzuki et al., 2020) tailored for low circuit depth and higher confidence use fewer oracle queries and/or fewer repetitions. The vanilla QAE in Figure 1 can be replaced with one of these variants without any loss of generality. A detailed comparison of these QAE variants can be found in Figure 3 in Grinko et al. (2021).

2.5 Quantum Machine Learning and Adversarial Robustness

There are multiple formulations of QML classifiers (Cong et al., 2019; Abohashima et al., 2020). For computer vision, experimental QML research focuses on MNIST task, using the output of intermediate hidden layers as features for final classification. In our algorithm, the oracle O in section 2.3 should be a QML classifier, such as Hasan (2021).

Previous works have shown that QML models are prone to adversarial attacks. For example, (Guan et al., 2020) checks the robustness of QML models against noise in the training data using classical methods by modeling the verification problem as a classical SDP, while (Weber et al., 2021) implements new protocols for QML to optimally certify robustness. The threat models and algorithm in these papers differ from the algorithm presented here, though a QAE based approach can provide a quadratic speedup in case of (Weber et al., 2021) as well.

3 Quantum Adversarial Robustness

As shown in Figure 1, the proposed algorithm for Quantum Adversarial Robustness (QuAdRo) uses state preparation $U^{p,\lambda}$ and Grover Diffusion Operator G_c described in subsection 3.1 and 3.2. Using QAE, we are able to measure $\overline{\rho_c}$ - an estimate of ρ_c - that base classifier answers correctly after smoothing, to a high precision with high confidence, as discussed in subsection 4.

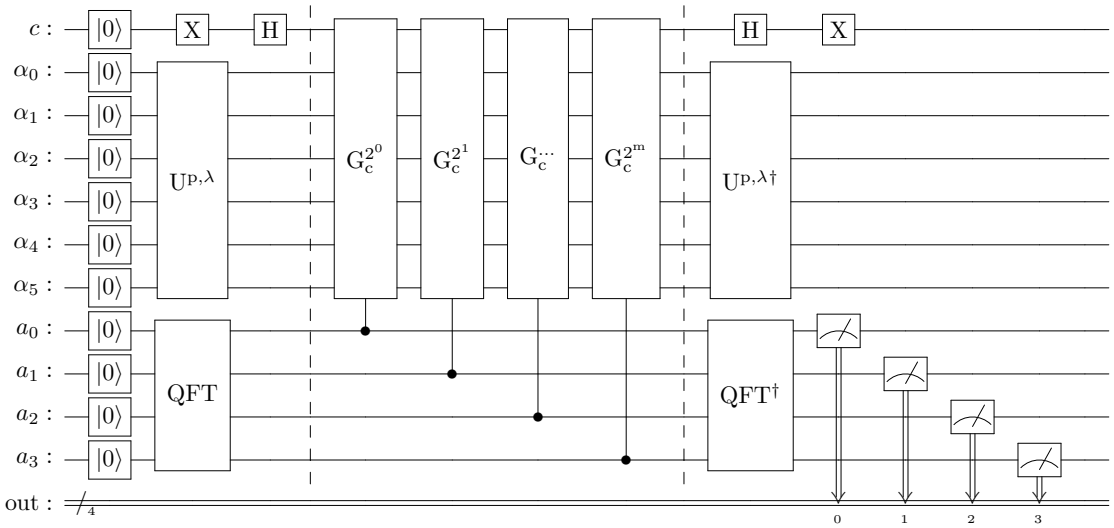

Figure 1: Quantum Estimation Circuit (QEC).

3.1 QUBIT STATE PREPARATION OPERATOR $U^{p,\lambda}$

Grover Diffusion Operator can only be used to search when distinct values in the input space are orthogonal. The smoothing operator maps a probability distribution defined over a region around the input such that the probability of a state i in the neighborhood is encoded in the amplitude for $|i\rangle$. For a detailed discussion about the design and functioning of the qubit encoding scheme, look at Appendix B.

Before state preparation, state qubits and ancilla are reset to $|0\rangle$. First, InputLoader (IL) loads the input I from classical bits into state qubits. Then, a superposition of the smoothing neighborhood of the input I based on a distribution $\phi(\lambda)$ is prepared based on the parameter λ, which is set to defend against a given l_p norm adversary. Each value $|i\rangle$ in the neighborhood is a valid input to the base classifier, a perturbed image for example.

For rest of the discussion, we define d as the size of the input vector I. For example, $d = height * width$ for an image, input is a bitmap $[\alpha_0, \alpha_1, \alpha_2, ...\alpha_{d-1}]$, which is encoded into qubit state $|\psi\rangle$ as follows

$$|\psi\rangle = |\alpha_0\alpha_1\alpha_2....\alpha_{d-1}\rangle = \bigotimes_{j=0}^{d-1} |\alpha_j\rangle \tag{2}$$

For all input values in Eq. 2, let v be the number qubits needed to represent any α_j. For an grayscale image, v is the resolution of a single pixel in the input, and if each pixel belongs to range 0-255 ie $[0, 2^8 - 1]$, then $v = 8$.

Let j^{th} value in input I, $\alpha_j = c_{j0}c_{j1}...c_{jv-1}$, i.e. $c_{ji} : i^{th}$ classical bit in α_j, and σ_i conventionally represents the ith Pauli operator where $\sigma_0 = \mathbb{I}$, $\sigma_1 = X$. Then, InputLoader IL_j that loads α_j into qubit state, can be defined as the tensor product of σ_i as follows

$$IL_j = \bigotimes_{i=0}^{v-1} \sigma_{c_{ji}} \tag{3}$$

and by definition $IL_j |0\rangle^v = |\alpha_j\rangle$. Since each value in the input vector I is loaded independently, complete InputLoader IL operator can be defined as

$$IL = \bigotimes_{j=0}^{d-1} IL_j \tag{4}$$

such that

$$IL |0\rangle^{dv} = |\alpha_0\alpha_1\alpha_2....\alpha_{d-1}\rangle = |\psi\rangle \tag{5}$$

Circuit IL comprises of a maximum of vd 1-qubit X gates. For any input I, if $\|I\|_1 = k$, then IL has exactly k X gates, corresponding to each $|1\rangle$ in $|\psi\rangle$.

Following IL, we define a smoothing circuit $S^{p,\lambda}$ with distribution ϕ defined by p, λ, such that

$$S_j^{p,\lambda} |\alpha_j\rangle = \Sigma_{k=0}^{2^v-1} \sqrt{\phi_j(\delta_k)} |\alpha_j + \delta_k\rangle \tag{6}$$

where $\alpha_j + \delta_k$ is a value in the neighborhood of α_j. In our scheme, the complete smoothing operator $S^{p,\lambda}$ maps each value α_j in the input vector independently to a probability weighted superposition of its neighborhood. For any i in the input space, probability weight of i in the smoothed qubit state is $\phi(i) = \Pi_j \phi_j(i_{jv}i_{jv+1}...i_{(j+1)v-1})$. Hence

$$S^{p,\lambda} |\psi\rangle = \bigotimes_{j=1}^{d} S_j^{p,\lambda} |\alpha_j\rangle = \Sigma_{i=0}^{2^{dv}-1} \sqrt{\phi(i)} |i\rangle \tag{7}$$

In this particular case, all ϕ_j are identical distributions centered around mean α_j. Detailed circuits for each p norm are discussed in section 5

The combination of IL and $S^{p,\lambda}$ is defined as the Qubit State Preparation Operator $U^{p,\lambda}$ (Figure 2a).

$$U^{p,\lambda} = \bigotimes_{j=0}^{d-1} U_j^{p,\lambda} = \bigotimes_{j=0}^{d-1} S_j^{p,\lambda} IL_j \tag{8}$$

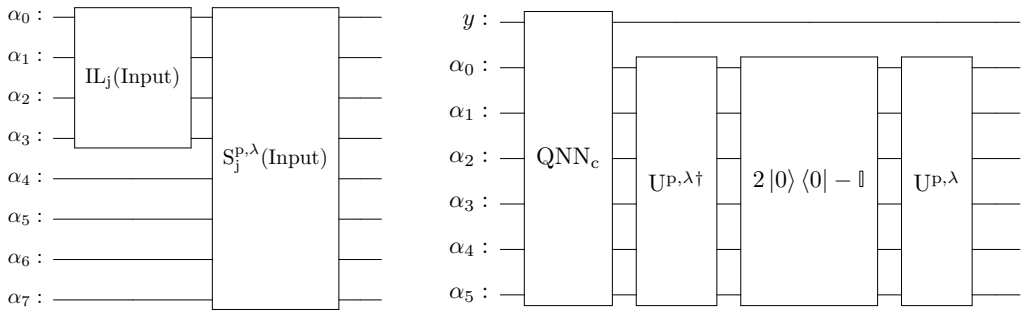

(a) $U^{p,\lambda}(Input)$ circuit. $S_j^{p,\lambda}$ can smooth the input image state to a higher resolution.

(b) Grover Diffusion Operator G_c circuit.

Figure 2: Building blocks of QEC.

3.2 Grover Diffusion Operator G_c

Based on section 2.3 and 2.4, we design a robust QML classifier from any base binary classifier f_c for class c provided by a unitary quantum parallel (Nielsen & Chuang, 2002) oracle QNN_c. If output of the QML classifier for input x is class c then $f_c(x) = 1$ else $f_c(x) = 0$.

$$QNN_c\,|x\rangle\,|y\rangle = |x\rangle\,\Big|y \bigoplus f_c(x)\Big\rangle \tag{9}$$

Setting $|y\rangle = |-\rangle$ (Deutsch & Jozsa, 1992).

$$QNN_c\,|x\rangle\,|-\rangle = (-1)^{f_c(x)}\,|x\rangle\,|-\rangle \tag{10}$$

When applying QNN_c to the smoothed superposition of inputs given by $U^{p,\lambda}\,|0\rangle^{dv}$ from section 3.1, ρ_c is the expected probability output by our algorithm i.e. input $|x+\delta\rangle$ in the smoothing neighborhood belongs to class c

$$\rho_c = \mathbb{E}(\mathbb{P}(f_c(x+\delta) = 1)) \tag{11}$$

The Grover Diffusion Operator derives from state preparation circuit $U^{p,\lambda}$ and QNN_c as shown in Figure 2b.

$$G_c = U^{p,\lambda}(2\,|0\rangle\,\langle 0| - \mathbb{I})U^{p,\lambda\dagger}\,QNN_c \tag{12}$$

3.3 Quantum Estimation Circuit

QEC uses QAE (Brassard et al., 2002) to solve the counting problem defined for QNN_c using $U^{p,\lambda}$ and G_c. The ancilla qubits $|a\rangle$ in Figure 1 are initialized in a uniform superposition in range [0, M] via QFT. Here, M is the total number of calls to the oracle in the circuit. Then, we apply controlled $G_c^{2^k}$ gates using $|a_k\rangle$ as control, which results in

$$QAE\,|\psi\rangle\,|a\rangle = G_c^a\,|psi\rangle \bigotimes |a\rangle \tag{13}$$

for all a. After applying $QFT^\dagger$, we measure the ancilla state to obtain θ such that the probability measure ρ_c is

$$\rho_c = sin^2(\theta/2) \tag{14}$$

4 Theoretical bounds

As noted earlier, the smoothing distribution ϕ_j for each value in the input vector I is independent, as per classical robustness criteria (Cohen et al., 2019; Yang et al., 2020). We claim that the proposed quantum algorithm QuAdRo finds ρ_c with a quadratic speedup compared to the best-known classical

Algorithm 1 Quantum Adversarial Robustness (QuAdRo)

procedure QUADRO(Input $|\psi\rangle$, Class c, $U^{p,\lambda}, G_c, M, \delta$)
  $N_{rep} = 12\log(\frac{1}{\delta})$ + 1
  $N_{QEC} = \frac{M}{N_{rep}}$
  for `i in Range(`N_{rep}`)` **do**
    $\theta_i = QEC1(|\psi\rangle, U^{p,\lambda}, G_c, N_{QEC}, \delta)$
    ρ_c_measured [i] = $sin^2(\frac{\theta_i}{2})$
  end for
  $\overline{\rho_c}$ = median(ρ_c_measured)
  lowerConfBound = $\overline{\rho_c} - \frac{7}{N_{QEC}}$
  if $lowerConfBound \geq \frac{1}{2}$ **then return** c, CertifiedRadius(lowerConfBound)
  else return -1, ABSTAIN
  end if
end procedure

algorithm with the same error bound. We show that the certified robustness problem can be reduced to a counting problem, and for a given error ϵ in the measurement of ρ_c for fixed confidence δ, the best known classical algorithm is $O(\frac{1}{\epsilon^2})$ and QuAdRo is $O(\frac{1}{\epsilon})$. This is also the best speedup possible for a counting problem using a quantum computer (Brassard et al., 2002).

4.1 QUANTUM COMPUTING BOUNDS

Given that m ancilla qubits are used for estimation, maximum number of oracle calls can be

$$M = \Sigma_{k=1}^{m} 2^k = 2^{m+1} - 1 \tag{15}$$

The probability that we measure θ correctly upto m bits is $\frac{8}{\pi^2}$, and the measured probability value $\overline{\rho_c}$ is such that

$$\|\overline{\rho_c} - \rho_c\| \leq \epsilon_0 = 2\pi\frac{\sqrt{\rho_c(1-\rho_c)}}{M} + \frac{\pi^2}{M^2} \tag{16}$$

More generally, for a single experiment with no repetitions, the error in measurement of ρ_c with a confidence $1-\delta$ (using theorem 12 from Brassard et al. (2002)) is

$$\|\overline{\rho_c} - \rho_c\| \leq \epsilon = \epsilon_0 + \frac{1}{\delta}(\frac{\sqrt{\rho_c(1-\rho_c)}}{2M} + \frac{\pi^2}{4M^2}) + \frac{1}{\delta^2}\frac{\pi^2}{4M^2} \tag{17}$$

Solving Eq. 17 for M in terms of ϵ, δ gives

$$M \approx \frac{\pi\sqrt{\rho_c(1-\rho_c)}}{\delta\epsilon} \sim O(\frac{1}{\delta\epsilon}) \tag{18}$$

The success probability $\frac{8}{\pi^2}$ in Eq. 16 can quickly be boosted to close to 100% by repeating the experiment multiple times and using the median estimate of ρ_c (Grinko et al., 2021). As a result,

$$M \sim O(\frac{1}{\epsilon}log(\frac{1}{\delta})) \tag{19}$$

Constants for calculating M in Eq. 19 are small, as shown in Algorithm 11 and Miyamoto (2022).

4.2 CLASSICAL COMPUTING BOUNDS

There are multiple intervals and bounds commonly used for statistical guarantees. They result in a similar order of magnitude in calculating M in terms of error ϵ in measuring ρ_c for a confidence $1-\delta$. If we use Chernoff Bounds

$$\mathbb{P}(\|X-\mu\| \leq \delta\mu) \geq 1 - 2e^{-\frac{\mu\delta^2}{3}} \tag{20}$$

ie for M repetitions of the function f classically,

$$\mathbb{P}(\|\overline{\rho_c} - \rho_c\| \leq \epsilon) \geq 1 - 2e^{-\frac{\epsilon^2}{3M\rho_c}} \tag{21}$$

For the same setting, using Clopper-Pearson interval is a more popular practice since it provides tighter bounds upto $O(M^{-\frac{3}{2}})$ (Thulin, 2014).

$$\epsilon \leq M^{-\frac{1}{2}} z_{\frac{\delta}{2}} \sqrt{\rho(1-\rho)} + M^{-1} \tag{22}$$

For confidence $1-\delta$, ignoring $\frac{1}{M}$ for large M,

$$\|\overline{\rho_c} - \rho_c\| \leq M^{-\frac{1}{2}} z_{\frac{\delta}{2}} \sqrt{\rho(1-\rho)} \tag{23}$$

Hence

$$M \approx \frac{z^2_{\frac{\delta}{2}} \rho(1-\rho)}{\epsilon^2} \sim O(\frac{1}{\epsilon^2}) \tag{24}$$

5 CIRCUITS FOR SMOOTHING DISTRIBUTIONS p, λ

Table 1 shows best performing smoothing distributions for each common adversary (Yang et al., 2020; Lee et al., 2019). As shown in Yang et al. (2020), same distribution can be used to counter multiple adversaries. These are a subset of all possible (ϕ, l_p) pairs, for which we present the state preparation circuits. Please refer to other works (Yang et al., 2020; Lee et al., 2019; Cohen et al., 2019) for detailed discussion and demonstration of optimal distributions. Density in Table 1 should match the probability distribution obtained from state preparation $U^{p,\lambda}$. Appendix B discusses this topic in detail.

Table 1: Adversary norm parameters and corresponding certified robustness radius.

Adversary l_p	Distribution ϕ	Density	Certified Radius
l_0	Tight Certificate	$x == \alpha \ ? \ \frac{\lambda}{2^v} : \frac{2^v - \lambda}{2^v(2^v-1)}$	$argmax_r \rho_r^{-1}(0.5) \geq p$
l_1	Uniform l_∞	$\propto \mathbb{I}(\|x\|_\infty \leq \lambda)$	$2\lambda(\rho_c - \frac{1}{2})$
l_2	Normal Distribution	$\propto e^{-\|\frac{x}{\lambda}\|_2^2/2}$	$\lambda GaussianCDF^{-1}(\rho; 0, 1)$

5.1 l_0 NORM

A state preparation circuit $U^{0,\lambda}$ requires $O(v^2)$ gates to smooth any pixel with amplitude $\sqrt{\frac{\lambda}{2^v}}$ for the base pixel state $|\alpha\rangle$ and an equal superposition of rest of the space, such that

$$S^{0,\lambda} |\alpha\rangle = \sqrt{\frac{\lambda}{2^v}} |\alpha\rangle + \Sigma_{x \neq \alpha} \sqrt{\frac{2^v - \lambda}{2^v(2^v-1)}} |x\rangle \tag{25}$$

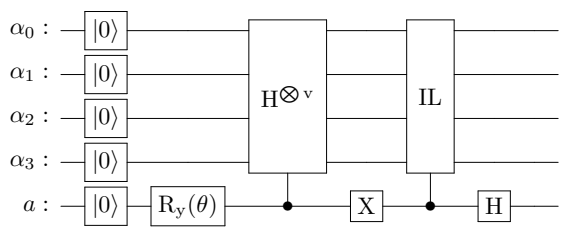

Figure 3: $U^{p,\lambda}$ circuit for l_0 norm adversary.

$U^{0,\lambda}$ in Figure 3 requires only $O(v)$ elementary gates ($\approx 2v$), where $cos(\frac{\theta}{2}) = \sqrt{\frac{\lambda-1}{2^v-1}}$. This circuit requires one additional ancilla qubit a that cannot be disentangled.

5.2 l_1 NORM

For l_1 norm adversary, we use a uniform distribution with width λ, centered at the pixel value α. To create a superposition state in a Uniform distribution, if $\lambda = 2^k, k \in \mathbb{Z}$, then the Hadamard operator H on k-qubits $H^{\bigotimes k}$ can be used. Circuit in Figure 3 can also be adapted to prepare a uniform distribution by setting $\theta = \frac{\pi}{2}$. More generally, $QFT(\lambda)$ can prepare the requisite uniform distribution for any λ. The prepared uniform distribution can then be shifted to its corresponding mean, resulting in $U^{p,\lambda}$ as shown in Figure 4.

5.3 l_2 NORM

Creating Gaussian distribution (and any other log-normal distribution) is a well-studied problem (Grover & Rudolph, 2002). The relevant circuit is described in Figure 8 in the Appendix B. The same effect can also be achieved using the SHIFT(α, w) operator, as shown in Figure 4.

5.4 OTHER ADVERSARIES

Any log-concave distribution can be created as a qubit state with a quantum circuit (Grover & Rudolph, 2002), and the SHIFT operator in Figure 7 to shift the distribution to a mean value. Yang et al. (2020) uses log-concave distributions for randomized smoothing, for instance Laplace or Uniform distribution for l_∞ norm adversary. Refer to Table A.1 in Yang et al. (2020) for a comprehensive list of viable distributions and corresponding certified radii.

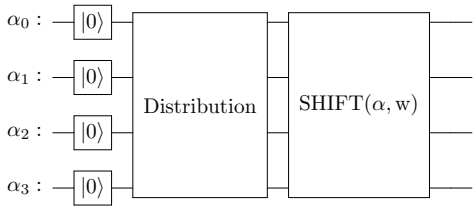

Figure 4: $U^{p,\lambda}$ circuit for l_p norm.

6 EXPERIMENTS

Due to a lack of access to QC hardware at the time of writing, all the presented results are obtained using simulation. For the base classifier, we train a 2-layer Deep Neural Network (DNN) model for recognizing handwritten digits, using the MNIST dataset from Torch v1.11.0. First, fully connected layer is followed by leakyRelu activation, normalization, and scaling to 4-bit integer representation. This discretized classifier has an accuracy $\approx 89.8\%$ on the MNIST test dataset. We found that Projected Gradient Descent (PGD) adversary (Madry et al., 2017) can successfully attack this classifier, and this attack can be defended against using Randomized Smoothing (Cohen et al., 2019). Detailed model architecture can be found in Appendix A. All training and simulation used an Nvidia RTX 2080 Ti GPU.

To compare QuAdRo with the classical randomized smoothing algorithm, we reduce the problem size to efficiently simulate the QC circuit. First, we use the feature vector after layer 1 of the classifier as input that the adversary can attack. Second, we use a small layer 2 - $d = 5$. Third, we discretize feature vectors to 4-bit integer values after training - $v = 4$. This limits the input space to 20 bits, ie 2^{20} possible inputs, which can be completely simulated classically. Additional 6-12 ancilla qubits are required for the QEC. Based on preliminary experiments, the simulation cost of QEC-based prediction is high so a classical predictor with $N_{predict} = 100$ is used in all experiments. Complete MNIST test set is used for experiments in Table 2, while simulation plots in Figure 5 use subsamples of size 1000 for 5a, b and 100 for 5c.

The experiments are designed around four parameters - λ adversary, N_{QEC} calls to the classifier oracle, 1 - δ confidence, and $2 * w + 1$ width of the truncated smoothing distribution. It is evident from Figure 5a that increasing λ reduces certified accuracy as expected while a smaller λ represents

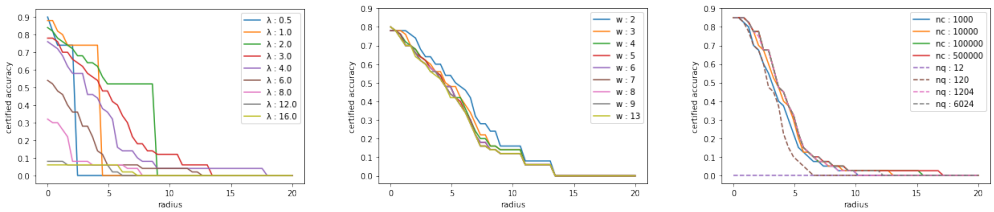

Figure 5: Certified accuracy vs radius for parameters - **a (left)** Adversary λ for Normal distribution; **b (middle)** Smoothing distribution truncation parameter w ; **c (right)** Oracle calls - nc : $N_{certify}$ for classical algorithm; nq : N_{QEC} for QuAdRo;

a weaker adversary. The effect of change in width w on $\overline{\rho_c}$ in Figure 5b is not very pronounced. We observed that $w < \lambda$ cannot model a λ adversary for radius $> w * d$. Since w linearly increases simulation cost, $\lambda \in [2, 4]$, $w \in [4, 5]$ are optimal for our experiments. Table 2 shows the results for our experiments.

Table 2: Results : Accuracy of robust classifiers for MNIST test dataset - abstain (Abs) when $\overline{\rho_c} < 0.5$. Acc (!Abs) refers to the accuracy for certified inputs where classifier does not abstain.

Classifier	**Acc** %	**Abs** %	**Acc (!Abs)** %	$\overline{\rho_c}$ **(y =** y_{pred}**)**	$\overline{\rho_c}$ **(y** $\neq y_{pred}$**)**
Base	89.73	-	-	-	-
Quadro N_{QEC} : 12	88.4	22.0	95.2	0.80	0.45
Quadro N_{QEC} : 120	88.4	17.7	94.1	0.83	0.39
Classical $N_{certify}$: 1000	87.7	17.8	94.0	0.82	0.40
Classical $N_{certify}$: 10000	88.0	17.8	94.0	0.83	0.40

Based on theoretical discussion (Thulin, 2014), given the number of calls to the classifier oracle $M = N_{certify}$, we can use $N_{certify}$ samples for classical certification while QuAdRo uses N_{REP} iterations of size N_{QEC}. Given $\epsilon_q \approx \frac{7}{N_{QEC}}$, $\epsilon_c \approx \sqrt{\frac{\rho_c(1-\rho_c)}{N_{certify}}}$, we should see a quantum advantage ($\epsilon_q < \epsilon_c$) for

$$N_{QEC} > 49 * \frac{12\, log\frac{1}{\delta} + 1}{z^2_{\frac{\delta}{2}} \rho_c(1-\rho_c)} \tag{26}$$

QuAdRo's accuracy outpaces the classical algorithm as N_{QEC}) i.e. $N_{certify}$ increases, specially for contentious examples since they have a low certified radius. For median $\overline{\rho_c} = 0.8$ from Table 2, the results for the two algorithms match around $N_{QEC} \approx 2000$ for $\delta = 0.001$, which is observed in Figure 5c. To match the certified radius for $\rho_c \approx 1.0$, N_{QEC} must scale by $\approx \frac{1}{1-\rho_c}$.

7 Conclusion

Section 4 establishes that the proposed QC algorithm QuAdRo offers a quadratic speedup over existing classical algorithms for certified adversarial robustness via randomized smoothing. The same guarantees were simulated and tested in Section 6. The results hold for any distribution and classifier that can be implemented on a Quantum Computing device, and the QC circuit to prepare popular smoothing distributions as qubit states are presented in Section 5 and Appendix B. We have not made any additional assumptions on the base classifier hence any guarantees that hold for prior art (Cohen et al., 2019; Yang et al., 2020) should hold for QuAdRo as well. Based on experiments, QuAdRo abstains less frequently and is more reliable compared the classical robust classifier for when $N_{certify}$ is increased.

Broader Impact and Limitations: The robustness guarantees in this paper apply only to the adversarial attacks on a classifier. The underlying bias (e.g training data, misrepresentation of objects/people, etc) in the base classifier will be inherited by the robust classifier as well. The algorithm presented here is also dependent on a functional QC hardware for the speedup, and application to meaningful image inputs will need >1000 qubits.

REFERENCES

Scott Aaronson. Read the fine print. *Nature Physics*, 11(4):291–293, 2015.

Scott Aaronson and Patrick Rall. Quantum approximate counting, simplified. In *Symposium on Simplicity in Algorithms*, pp. 24–32. SIAM, 2020.

Zainab Abohashima, Mohamed Elhosen, Essam H Houssein, and Waleed M Mohamed. Classification with quantum machine learning: A survey. *arXiv preprint arXiv:2006.12270*, 2020.

Ethan Bernstein and Umesh Vazirani. Quantum complexity theory. *SIAM Journal on computing*, 26 (5):1411–1473, 1997.

David Biron, Ofer Biham, Eli Biham, Markus Grassl, and Daniel A Lidar. Generalized grover search algorithm for arbitrary initial amplitude distribution. *Lecture notes in computer science*, pp. 140–147, 1999.

Michel Boyer, Gilles Brassard, Peter Høyer, and Alain Tapp. Tight bounds on quantum searching. *Fortschritte der Physik: Progress of Physics*, 46(4-5):493–505, 1998.

Gilles Brassard, Peter Hoyer, Michele Mosca, and Alain Tapp. Quantum amplitude amplification and estimation. *Contemporary Mathematics*, 305:53–74, 2002.

Jeremy Cohen, Elan Rosenfeld, and Zico Kolter. Certified adversarial robustness via randomized smoothing. In *International Conference on Machine Learning*, pp. 1310–1320. PMLR, 2019.

Iris Cong, Soonwon Choi, and Mikhail D Lukin. Quantum convolutional neural networks. *Nature Physics*, 15(12):1273–1278, 2019.

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.

David Deutsch. Quantum theory, the church–turing principle and the universal quantum computer. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, 400(1818): 97–117, 1985.

David Deutsch and Richard Jozsa. Rapid solution of problems by quantum computation. *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences*, 439(1907):553–558, 1992.

Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.

Dmitry Grinko, Julien Gacon, Christa Zoufal, and Stefan Woerner. Iterative quantum amplitude estimation. *npj Quantum Information*, 7(1):1–6, 2021.

Lov Grover and Terry Rudolph. Creating superpositions that correspond to efficiently integrable probability distributions. *arXiv preprint quant-ph/0208112*, 2002.

Lov K Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pp. 212–219, 1996.

Ji Guan, Wang Fang, and Mingsheng Ying. Robustness verification of quantum machine learning. *CoRR*, 2020.

Aram W Harrow, Avinatan Hassidim, and Seth Lloyd. Quantum algorithm for linear systems of equations. *Physical review letters*, 103(15):150502, 2009.

Areeq I Hasan. Igo-qnn: Quantum neural network architecture for inductive grover oracularization. *arXiv preprint arXiv:2105.11603*, 2021.

Daniel Koch, Michael Samodurov, Andrew Projansky, and Paul M Alsing. Gate-based circuit designs for quantum adder-inspired quantum random walks on superconducting qubits. *International Journal of Quantum Information*, pp. 2150043, 2022.

Phuc Q Le, Fangyan Dong, and Kaoru Hirota. A flexible representation of quantum images for polynomial preparation, image compression, and processing operations. *Quantum Information Processing*, 10(1):63–84, 2011.

Guang-He Lee, Yang Yuan, Shiyu Chang, and Tommi Jaakkola. Tight certificates of adversarial robustness for randomly smoothed classifiers. *Advances in Neural Information Processing Systems*, 32, 2019.

Haoran Liao, Ian Convy, William J Huggins, and K Birgitta Whaley. Robust in practice: Adversarial attacks on quantum machine learning. *Physical Review A*, 103(4):042427, 2021.

Sirui Lu, Lu-Ming Duan, and Dong-Ling Deng. Quantum adversarial machine learning. *Physical Review Research*, 2(3):033212, 2020.

Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.

Koichi Miyamoto. Bermudan option pricing by quantum amplitude estimation and chebyshev interpolation. *EPJ Quantum Technology*, 9(1):1–27, 2022.

Michael A Nielsen and Isaac Chuang. Quantum computation and quantum information, 2002.

Arthur G Rattew and Bálint Koczor. Preparing arbitrary continuous functions in quantum registers with logarithmic complexity. *arXiv preprint arXiv:2205.00519*, 2022.

Wenhui Ren, Weikang Li, Shibo Xu, Ke Wang, Wenjie Jiang, Feitong Jin, Xuhao Zhu, Jiachen Chen, Zixuan Song, Pengfei Zhang, et al. Experimental quantum adversarial learning with programmable superconducting qubits. *arXiv preprint arXiv:2204.01738*, 2022.

Jonathan Romero, Ryan Babbush, Jarrod R McClean, Cornelius Hempel, Peter J Love, and Alán Aspuru-Guzik. Strategies for quantum computing molecular energies using the unitary coupled cluster ansatz. *Quantum Science and Technology*, 4(1):014008, 2018.

Peter W Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM review*, 41(2):303–332, 1999.

Yohichi Suzuki, Shumpei Uno, Rudy Raymond, Tomoki Tanaka, Tamiya Onodera, and Naoki Yamamoto. Amplitude estimation without phase estimation. *Quantum Information Processing*, 19 (2):1–17, 2020.

Måns Thulin. The cost of using exact confidence intervals for a binomial proportion. *Electronic Journal of Statistics*, 8(1):817–840, 2014.

Maurice Weber, Nana Liu, Bo Li, Ce Zhang, and Zhikuan Zhao. Optimal provable robustness of quantum classification via quantum hypothesis testing. *npj Quantum Information*, 7(1):1–12, 2021.

Greg Yang, Tony Duan, J Edward Hu, Hadi Salman, Ilya Razenshteyn, and Jerry Li. Randomized smoothing of all shapes and sizes. In *International Conference on Machine Learning*, pp. 10693–10705. PMLR, 2020.

Yi Zhang, Kai Lu, Yinghui Gao, and Mo Wang. Neqr: a novel enhanced quantum representation of digital images. *Quantum information processing*, 12(8):2833–2860, 2013.

A SIMULATION : DISCUSSION, PLOTS, RESULTS

To be able to simulate QuAdRo on a classical computer, we reduced the problem size by following three changes - first, we used feature vector after layer 1 of the model as input that adversary can attack. Second, we use a small layer 2. Third, we discretize feature vectors to 4 bit int values after training. This allows us to limit the total input space to 20 bits, ie a space of 2^{20} possible inputs, which can be completely simulated.

Even for the smaller problem size, we show that

- The input space is local i.e. the nearby points in a neighborhood are likely to belong to the same class
- A non-trivial classical classifier C can be trained on the dataset, achieving $\approx 90\%$ accuracy
- An adversary A (PGD) can easily find inputs that are mis-classified by C
- Classifier based on randomized smoothing with l_p norm and parameter λ is effective against A's attack.

We created the quantum algorithm simulator in numpy, using output of the C over the complete input space of size 2^20 to calculate the superposition of states and exact ρ_c for comparison.

A.1 MODEL ARCHITECTURE

First, we train a 2-layer DNN model - trained from scratch : 2 layer DNN, Layer 1 of size 784*5, Layer 2 of size 5*10, using LeakyRelu activation with a scaling factor 0.03. Layer 2 inputs are normalized, and scaled by 2^4. Base classifier C inherits trained Layer2 from model, and inputs are the feature vector output by Layer 1 in the previously trained model, discretized to 4 bit integer representation in range $[0, 16]$. Leaky Relu activation with a scaling factor 0.03 is used after fc1, to make feature vectors partially invertible and can be used to visualize images for features outside the train and test datasets.

A.2 SIMULATION SETUP

We simulate QuAdRo for a 20 qubit input, with a total of 26-32 simulated qubits, since simulating a much larger system is not feasible on a classical computer. Due to lack of access to a quantum computer at the time of writing, the presented results were all simulated.

Given this setup, we implement both the classical and quantum computing algorithms for randomized smoothing to find the certified radius. We simulate Normal Distribution since it can defend against l_1, l_2, and l_∞ adversaries (Yang et al., 2020). We show that for same number of calls to the classifier, QuAdRo error bounds improves and surpasses the classical algorithm for same hyperparameters and input.

The simulator implementation is non-trivial. We recommend reading the code (in supplementary material) and comments in addition to Appendix A. Following is a list of noteworthy features of the simulator

- Calculates MNIST classifier output for the complete input space. This is used to simulate QNN_c oracle and calculate the exact value of rho_c and corresponding certified radius.
- Smoothing state psi is prepared using discretized Normal Distribution defined by

  ```
  f(i, mu, sd) = np.exp(-(i - mu)**2/(2*sd**2)
  ```

  for pixel value mu, $sd = \lambda$.
- The QFT operator on ancilla and the following Grover operator applied to psi have been simplified into an iterative operation on initial state. The iteration has N_{QEC} steps.

A.3 CODE SNIPPETS

We will release the code in supplementary material with its license. The code snippets here are only for reference, and we recommend using the code from supplementary material for simulation. The

following two functions can be used to simulate the Grover's Search operation by first preparing the requisite psi, then repeating the search operation, or using it in a subroutine for amplitude estimation.

```
def get_psi_np(vals, low, high sd, v):
    p = []
    psi = 0
    for i, val in enumerate(vals):
        x = np.zeros(num_points)
        for i in range(low, high):
            x[i] = np.exp(-(i - mu)**2/(2*sd**2))
        p.append(x/np.linalg.norm(x))
        if(i == 0):
            psi = p[i]
        else:
            psi = np.kron(psi, p[i])
    norm = np.linalg.norm(psi)
    psi = psi/norm
    return (p, psi)

def grover(classifier_tensor, psi, prev_state):
    a = np.dot(psi, prev_state)
    cur_state = 2*a*psi - prev_state
    cur_state = classifier_tensor*cur_state
    # Normalize to deal with floating point issues
    norm = np.linalg.norm(cur_state)
    return cur_state/norm
```

A.4 Other

One major difference in classical randomized smoothing and QuAdRo is the order of sampling and classifier operation. Classically, we sample from the neighborhood first, and then apply base classifier on each sample. In QuAdRo, we first apply the base classifier over the whole space, and then measure a posteriori. We expect to see a consistently better performance from QuAdRo for large N_{QEC} than suggested by the theoretical discussions.

B State Preparation Circuits

B.1 Qubit Encoding choice

Classical data can be encoded into a quantum state in numerous ways. Amplitude encoding is one of the most popular methods - a vector of real/complex values is encoded as the amplitude of the N states in a superposition such that the amplitude of $|i\rangle$ corresponds to the ith value in the input vector. Variations of this scheme have been used for different algorithms (Zhang et al., 2013; Le et al., 2011).

Oracle output for any input is pixel state dependent and any smoothing operation will destroy that information if pixel state was amplitude encoded - this is equivalent to running the algorithm for the mean of all images in the smoothing neighborhood for the given distribution.

B.2 l_0 Norm

The circuit in Figure 6 is optimized for low circuit depth and gate count. General state of this system will be referred to as $|\psi_n\rangle = |x\rangle\,|a\rangle$, where $|x\rangle$ represents the qubits that will eventually be in a superposition state of the smoothing distribution, while $|a\rangle$ is the additional qubit needed in the circuit. Subscript n represents n^{th} logical step in the circuit.

Initially, all qubits are set to the ground state $|0\rangle$, hence the state of the system is

$$|\psi_0\rangle = |0\rangle^v\,|0\rangle \tag{27}$$

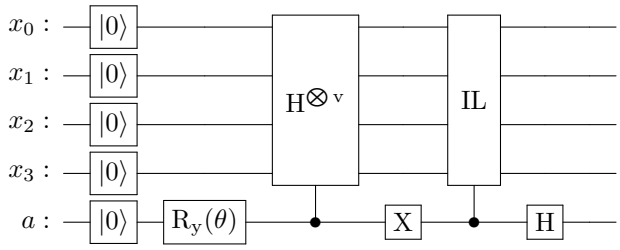

Figure 6: $U^{p,\lambda}$ circuit for l_0 norm adversary.

Assuming V elements in the input space and adversary λ such that $V > \lambda \geq 1$, the operator $R_y(\theta)$ rotates the ancilla state a, where $cos(\frac{\theta}{2}) = \sqrt{\frac{\lambda-1}{V-1}}$

$$R_y(\theta)\,|0\rangle = cos(\frac{\theta}{2})\,|0\rangle - sin(\frac{\theta}{2})\,|1\rangle \tag{28}$$

hence

$$|\psi_1\rangle = cos(\frac{\theta}{2})\,|0\rangle^v\,|0\rangle - sin(\frac{\theta}{2})\,|0\rangle^v\,|1\rangle \tag{29}$$

Applying v-dimensional Hadamard gate $H^{\otimes v}$ creates a uniform superposition of $V = 2^v$ states. Here, $H^{\otimes v}$ can be replaced with QFT if $V \neq 2^v$ for generality. State after the X gate is

$$|\psi_2\rangle = cos(\frac{\theta}{2})\,|0\rangle^v\,|1\rangle - \frac{1}{\sqrt{V}}sin(\frac{\theta}{2})(\Sigma_{i=0}^{V-1}\,|i\rangle)\,|0\rangle \tag{30}$$

After image loading,

$$|\psi_3\rangle = cos(\frac{\theta}{2})\,|\alpha\rangle\,|1\rangle - \frac{1}{\sqrt{V}}sin(\frac{\theta}{2})(\Sigma_{i=0}^{V-1}\,|i\rangle)\,|0\rangle \tag{31}$$

After Hadamard H on the ancilla

$$|\psi_4\rangle = (cos(\frac{\theta}{2})\,|\alpha\rangle - \frac{1}{\sqrt{V}}sin(\frac{\theta}{2})\Sigma_{i=0}^{V-1}\,|i\rangle)\,|0\rangle - (cos(\frac{\theta}{2})\,|\alpha\rangle + \frac{1}{\sqrt{V}}sin(\frac{\theta}{2})\Sigma_{i=0}^{V-1}\,|i\rangle)\,|1\rangle \tag{32}$$

Applying QEC on $|\psi_4\rangle$ works for either value of a, with initial probability

$$\mathbb{P}(|x\rangle = \alpha) = cos^2(\frac{\theta}{2}) + \frac{sin^2(\frac{\theta}{2})}{V-1} = \frac{\lambda}{V} \tag{33}$$

B.3 Alternate formulation

Alternatively, in a situation where repeating measurement of $|a\rangle$ is easier, we can use $cos(\frac{\theta}{2}) = \frac{\sqrt{V\lambda}-\sqrt{V+1-\lambda}}{V+1}$, and repeat the experiment until measurement $|a\rangle = 1$.

$$|x\rangle = (cos(\frac{\theta}{2}) + \frac{1}{\sqrt{V}}sin(\frac{\theta}{2}))\,|\alpha\rangle + \frac{1}{\sqrt{V}}sin(\frac{\theta}{2})\Sigma_{i\neq\alpha}\,|i\rangle \tag{34}$$

and probability of measuring ancilla

$$\mathbb{P}(|a\rangle = 1) = \frac{1}{2}(1 + 2(\frac{\theta}{2})sin(\frac{\theta}{2}) \tag{35}$$

Hence

$$\mathbb{P}(a=1) - \mathbb{P}(a=0) = \frac{2}{(V+1)}(1 - \frac{(V-1)\sqrt{\lambda(V+1-\lambda)}}{V+1}) \tag{36}$$

Since $\mathbb{P}(a = 1) > \frac{1}{2}$, estimated number of repetitions of the state preparation circuit is less than 2. Please note that state preparation is also a part of Grover's Diffusion Operator, and will concentrate ancilla state $|1\rangle$ after repeated applications. In general, measuring the output of QEC circuit multiple (<2) times until $a = 1$ yields the required state. We expect this alternate circuit to be more useful as a component in a variational circuit.

To get the correct result, we need to perform the experiment multiple times, until ancilla state is measured as $|1\rangle$. The expected number of trials $<= 2$. In addition, ancilla with the discarded state $|0\rangle$ has an error $\sim O(\frac{1}{V})$ in the amplitude of $|x\rangle$ which may be acceptable for some λ, V, and M.

B.4 SHIFT OPERATOR

Any prepared qubit state for a distribution can be shifted to a new mean using a quantum shift operator. The shift operator circuit $SHIFT\alpha$ based on QADD (Koch et al., 2022) shown in Figure 7 can take any superposition as input and shifts each $|i\rangle$ to $|i + \alpha\rangle$. We believe the same can be implemented using other quantum adders/circuits.

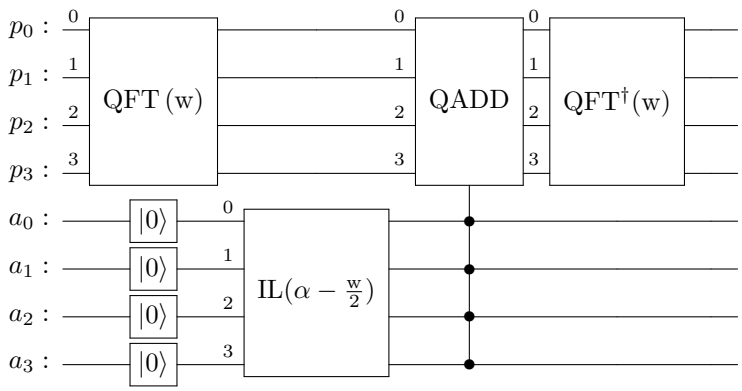

Figure 7: SHIFT(α, w) circuit to shift the input distribution $|p\rangle$ to mean α.

B.5 l_2 NORM

Basic Circuit for preparing l_2 norm distribution as a qubit state is shown in Figure 8.

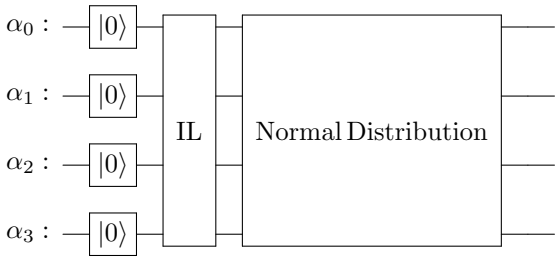

Figure 8: $U^{p,\lambda}(Input)$ implementation for l_2 norm (Grover & Rudolph, 2002).

The same effect can also be achieved using the SHIFT(α, w) operator, as shown in Figure 4.