# Improved maximum-likelihood quantum amplitude estimation

Adam Callison<sup>1</sup>, Dan E. Browne<sup>1</sup>

 $^1$  Department of Physics and Astronomy, University College London, London, WC1E 6BT, UK

E-mail: a.callison@ucl.ac.uk

Abstract. Quantum amplitude estimation is a key subroutine in a number of powerful quantum algorithms, including quantum-enhanced Monte Carlo simulation and quantum machine learning. Maximum-likelihood quantum amplitude estimation (MLQAE) is one of a number of recent approaches that employ much simpler quantum circuits than the original algorithm based on quantum phase estimation. In this article, we deepen the analysis of MLQAE to put the algorithm in a more prescriptive form, including scenarios where quantum circuit depth is limited. In the process, we observe and explain particular ranges of 'exceptional' values of the target amplitude for which the algorithm fails to achieve the desired precision. We then propose and numerically validate a heuristic modification to the algorithm to overcome this problem, bringing the algorithm even closer to being useful as a practical subroutine on near- and midterm quantum hardware.

#### 1. Introduction

Quantum computing promises significant computational speedups compared to classical computing in a number of important tasks throughout science and industry [1, 2], and has consequently been an active and important topic of research for several decades [3-5]. Among the wide variety of quantum algorithms that have been developed exists a core set of general purpose subroutines with broad applicability, and a key example is the quantum amplitude estimation algorithm. Quantum amplitude estimation can be used to estimate the probability that the output of another algorithm satisfies a particular target property, and can do this with a quadratic speedup in terms of the desired precision of the estimate, compared to classical sampling. This can be employed as a subroutine in a number of important algorithms, including quantum-enhanced Monte Carlo simulation [6-8] and algorithms in quantum machine learning [9-14].

The original quantum amplitude estimation due to Brassard et al. [15] essentially combines Grover's algorithm [16] and quantum phase estimation [17], and as a consequence involves quantum circuits with complicated controlled operations that may be challenging to implement in near- and mid-term quantum hardware. More recently, a number of much simpler quantum amplitude estimation algorithms have been developed that employ much simpler circuits while achieving the same speedup [18–26], including the maximum likelihood method (MLQAE) of Suzuki et al. [27] and its extensions [25, 28–33].

MLQAE employs a combination of Grover-like circuits and classical likelihoodmaximisation-based post-processing in order to find a high-quality amplitude estimate, and as such can be viewed as a hybrid quantum-classical algorithm [34]. Unlike the original quantum amplitude estimation algorithm and many of the other more recent proposals, the theoretical scaling analysis of maximum-likelihood quantum amplitude estimation relies on certain regularity and asymptotic assumptions related to the Bernstein-von Mises theorem [35], and as a consequence the quadratic speedup is not completely analytically rigorous. Nevertheless, numerical experiments in prior work show that the algorithm appears to achieve the quadratic speedup in practice with low constant overhead using relatively simple quantum circuits.

In this work, we extend the analysis of MLQAE to give the algorithm a more prescriptive structure: given an achievable circuit depth (expressed abstractly as the achievable number of applications of a Grover-like operator), we show how to approximately target a particular desired precision with some target probability. In the process of validating this extension through numerical simulations, we observe and explain in detail particular ranges of 'exceptional' values of the target amplitude for which the algorithm fails to achieve the desired precision. We then propose a heuristically-motivated modification to the algorithm to prevent this issue, and show through numerical simulation that our approach works. Overall, through the insights in this work, we improve the understanding of MLQAE, and move it toward usefulness as a practical algorithm in the near- and mid-term.

The article is structured as follows: in section 2, we review the relevant background for quantum amplitude estimation, including the problem definition, a discussion of some of its applications, and a detailed description of maximumlikelihood quantum amplitude estimation and its information-theoretic analysis. We begin our contribution in section 3 by extending the analysis and showing how to target a desired precision with a specified probability (up to certain regularity assumptions related to the Bernstein-von Mises theorem), including in depth-limited scenarios. In doing so, we identify particular 'exceptional' values of the target amplitude for which the algorithm fails to achieve the desired precision, and we explore and explain this phenomenon in detail in section 4. In section 5, we develop and numerically validate a heuristic method to mitigate the problematic exceptional values. Finally, in section 6, we summarise the article and suggest some directions for future work.

#### 2. Background

## 2.1. Problem definition

The formal setting in which quantum amplitude estimation (QAE) is typically considered assumes that one has access to some unitary quantum algorithm  $\hat{A}$  on n qubits and an oracle  $\hat{O}$  that partitions the computational basis  $\{|j\rangle_n\}_{j=1}^{2^n}$  into two disjoint subsets, the 'good' subset G and the 'bad' subset  $B = \{|j\rangle_n\}_{j=1}^{2^n} - G$ , via the action

$$\hat{O}|j\rangle = \begin{cases} |j\rangle, & \text{for } j \in G \\ -|j\rangle, & \text{for } j \in B. \end{cases}$$
(1)



Figure 1. Circuit for conventional quantum amplitude estimation

The state prepared by applying the algorithm  $\hat{A}$  to some reference state, which can be taken without loss of generality to be the all-zero state  $|0\rangle^{\otimes n}$ , can be written as

$$\left|A\right\rangle_{n} \equiv \hat{A}\left|0\right\rangle^{\otimes n} \tag{2}$$

$$= \sqrt{a} \left| A_G \right\rangle_n + \sqrt{1 - a} \left| A_B \right\rangle_n, \qquad (3)$$

where the (normalised) states  $|A_G\rangle_n$  and  $|A_B\rangle_n$  have support only on the 'good' and 'bad' subsets respectively, with phases chosen such that a is a positive real number. With these definitions, the goal of QAE is to find an estimate  $\tilde{a}$  of the probability a(referred to in this setting as an amplitude) that is correct up to a particular precision  $\epsilon$  with high probability. The precision requirement is sometimes considered to be *relative* 

$$a(1-\epsilon) \le \tilde{a} \le a(1+\epsilon),\tag{4}$$

and is sometimes considered to be *additive* 

$$a - \epsilon \le \tilde{a} \le a + \epsilon. \tag{5}$$

## 2.2. Conventional QAE

The first algorithm for quantum amplitude estimation is due to Brassard et al. [15], and draws on ideas from Grover's algorithm [16] and the Quantum Phase Estimation (QPE) algorithm [17]. The core of the algorithm is the Grover-like iteration operator

$$\hat{Q} = \hat{A} \left( \mathbb{1} - 2 \left| 0 \right\rangle \left\langle 0 \right|_n \right) \hat{A}^{\dagger} \hat{O}, \tag{6}$$

which can be shown to have a pair of eigenvalues  $\lambda_{\pm} = e^{\pm 2i\theta_a}$ , where  $\theta_a$  is defined by

$$a \equiv \sin^2 \theta_a \tag{7}$$

$$0 \le \theta_a \le \frac{\pi}{2}.\tag{8}$$

The corresponding eigenvectors  $|\lambda_{\pm}\rangle_n$  span the subspace span $(\{|A_G\rangle_n, |A_B\rangle_n\})$  and have no support outside this subspace.

The original QAE algorithm uses the circuit shown in Fig. 1; it requires not only the *n*-qubit (target) register needed for applying the algorithm  $\hat{A}$ , but an additional *q*-qubit control register, used for controlling applications of the Grover iteration  $\hat{Q}$ . The first application of the algorithm  $\hat{A}$  on the target register prepares the state  $|A\rangle_n$ , which has no support outside span( $\{|A_G\rangle_n, |A_B\rangle_n\}$ ), and thus no overlap with any eigenstates of the Grover iteration  $\hat{Q}$  other than  $|\lambda_{\pm}\rangle_n$ . The remainder of the circuit simply applies standard QPE to the Grover iteration  $\hat{Q}$ ; that is, turning the results  $m_0, m_1, \ldots, m_q - 1$  of the final layer of measurements into an integer  $k = 2^0 m_0 + 2^1 m_1 + \ldots, 2^{q-1} m_{q-1}$  and converting it into an angle will, with high probability, result in either  $\frac{\pi k}{2q} = \theta_{\tilde{a}}$  or  $\frac{\pi k}{2q} = \pi - \theta_{\tilde{a}}$  where  $|\theta_{\tilde{a}} - \theta_a| = O\left(\frac{1}{2q}\right)$ . Thus, with high probability, the amplitude estimate

$$\tilde{a} = \sin^2 \left( \min\left[ \left\{ \frac{\pi k}{2^q}, \pi - \frac{\pi k}{2^q} \right\} \right] \right) \tag{9}$$

will be within an error  $\epsilon \sim O\left(\frac{1}{2^q}\right)$  of the true amplitude *a*.

An application of the Grover iterator  $\hat{Q}$  requires two calls to the algorithm  $\hat{A}$ (more correctly, one call to  $\hat{A}$  and one to its inverse  $\hat{A}^{\dagger}$ ). Thus, an application of a *power* l of  $\hat{Q}$  ( $\hat{Q}^{l}$ ) requires 2l applications of the algorithm  $\hat{A}$ . Inspecting the target register in Fig. 1, it can be seen that the total number of calls to  $\hat{A}$  is

$$N_{\rm calls} = 1 + 2\sum_{j=0}^{q-1} 2^j \tag{10}$$

$$= 2 \cdot 2^q - 1 \tag{11}$$

$$= O\left(\frac{1}{\epsilon}\right). \tag{12}$$

This constitutes a quadratic speedup over the case of simply sampling from the output of  $\hat{A}$ , for which the number of calls scales as  $N_{\text{calls}} = O\left(\frac{1}{\epsilon^2}\right)$ .

While other forms of QPE exist that avoid the use of a quantum Fourier transform (QFT), they typically still require the unitary being studied (in this case the Grover iterator  $\hat{Q}$ ) to be controlled [17, 36–40], or the ability to implement Haar-random unitaries [41], both of which can be costly in terms of quantum resources [27].

# 2.3. Applications of quantum amplitude estimation

In this subsection, we briefly review a number of applications of QAE as a subroutine in quantum algorithms.

2.3.1. Monte Carlo simulation In [6], Montanaro showed that QAE can be used as a subroutine in otherwise classical Monte Carlo algorithms in order to achieve a quadratic quantum speedup, compared to the best-known fully classical approaches. As a simple example, consider the problem of estimating the mean  $\mu$  of a randomized (classical) algorithm A to within some additive error  $\epsilon$  with high probability. If an upper bound  $\sigma^2$  is known on the variance of the random output from the algorithm A, then it is a well-known result from classical statistics that such an estimate  $\tilde{\mu}$  for the mean  $\mu$  requires a number of Monte Carlo samples that scales as

$$N_{\rm samples}^{\rm (Monte Carlo)} = O\left(\frac{\sigma^2}{\epsilon^2}\right). \tag{13}$$

Montanaro showed that, by essentially replacing the classical algorithm  $\hat{A}$  with a quantum algorithm  $\hat{A}$  that outputs a quantum state which, if it were to be measured in the computational basis, would have the same probability distribution as the original algorithm  $\hat{A}$ , then QAE can be applied to produce an estimate  $\tilde{\mu}$  of the mean  $\mu$  more efficiently. The number of calls to the quantum algorithm  $\hat{A}$  in this case scales (up to polylogarithmic factors) as

$$N_{\text{calls}}^{(\text{QAE Monte Carlo})} = O\left(\frac{\sigma}{\epsilon}\right),\tag{14}$$

essentially a quadratic speedup over purely classical Monte Carlo. As an important extension, he also pointed out that if the quantum algorithm  $\hat{A}$  already has a quantum speedup compared to its classical counterpart A, that is if  $\hat{A}$  produces 'quantum samples' more efficiently than A produces samples, then this speedup will *concatenate* with the QAE Monte Carlo speedup.

Among a number of other more sophisticated examples, Montanaro showed how to use QAE to speed up Markov chain Monte Carlo algorithms for computing approximate partition functions; in this case, discrete-time quantum walks [42] are used to gain a quadratic speedup in the mixing time of the Markov chain, which then concatenates with the QAE-related speedup, producing an overall speedup compared to the best known classical algorithms. Other applications of quantum-enhanced Monte Carlo simulation are reported elsewhere [7, 8].

2.3.2. Machine learning Within the flourishing field of quantum machine learning, two key research directions are that of quantum clustering algorithms and quantum classification algorithms [9, 13]. In these algorithms, a core step is the estimation of the inner product  $\langle \boldsymbol{x} | \boldsymbol{y} \rangle$  between two vectors  $\boldsymbol{x}$  and  $\boldsymbol{y}$  for which unitary circuits  $U_{\boldsymbol{x}}$  and  $U_{\boldsymbol{y}}$  are known that can load them into quantum states via

$$|\boldsymbol{x}, \boldsymbol{y}\rangle = U_{\boldsymbol{x}, \boldsymbol{y}} |\boldsymbol{0}\rangle. \tag{15}$$

By choosing the product of unitaries  $U_{\boldsymbol{y}}^{\dagger}U_{\boldsymbol{x}}$  as the quantum algorithm  $\hat{A}$  and the state  $|\mathbf{0}\rangle$  as the (unique) 'good' state (as well as the initial state), QAE can be performed to produce an estimate of the inner product

$$\langle \boldsymbol{x} | \boldsymbol{y} \rangle = \langle \boldsymbol{0} | U_{\boldsymbol{y}}^{\dagger} U_{\boldsymbol{x}} | \boldsymbol{0} \rangle, \qquad (16)$$

as required. Other applications of QAE within quantum machine learning are reported elsehwere [10-12, 14].

2.3.3. Finance In the field of finance, an important task is to calculate a quantity known as the Economic Capital Requirement (ECR), which is essentially the amount of capital an investor must keep in order to protect against losses. This in turn requires an accurate understanding of the risk associated with a particular portfolio, and an often used measure of this risk is the "Value at Risk" (VaR) metric, defined for a loss distribution  $\mathcal{L}$  as the smallest total loss with probability P greater than or equal to the parameter  $\alpha$ ; that is

$$\operatorname{VaR}_{\alpha}\left[\mathcal{L}\right] = \inf_{x \ge 0} \left[ x | P(\mathcal{L} \le x) \ge \alpha \right]$$
(17)

where  $\alpha$  is often chosen to be around 99.9% [43, 44]. In [44], Egger et al. showed that QAE can be used to gain a quantum speed-up for the calculation of VaR (and other

related measures of risk). To do this, Egger et al. propose a quantum algorithm  $A_{\text{VaR}}$  that can be described at a high-level as being a combination of three parts; that is,  $\hat{A}_{\text{VaR}} = CSU$ . The first operator U loads the correlated uncertainty for a portfolio of k assets into one register of qubits. The second operator S then computes the total loss over the portfolio and stores the result, as a superposition, in another qubit register. The final operator C applies a threshold to the total loss and flips a single target qubit if the total loss is less than or equal to a given x value, and QAE is used to accurately determine the amplitude/probability of this target qubit having been flipped. A classical outer loop is then used to perform a bisection search to find the smallest x that is flipped with probability greater than or equal  $\alpha$ .

# 2.4. QAE without QPE

The conventional QAE of [15] (see subsection 2.2) includes a number of complicated controlled operations, as well as a quantum Fourier transform as the final step (see Fig. 1). As a result, the entire procedure may be difficulty to implement, particularly on near- and mid-term quantum hardware. The subsequent identification of a number of important applications for QAE (see subsection 2.3.1) led to a renewed interest in the algorithm, and a number of new approaches to the problem were developed that avoid the need for the quantum Fourier transform and the associated controlled- $\hat{Q}$ operations.

One of the first examples of QAE without QPE was proposed in [18]. In that work, Aaronson and Rall present a new, iterative algorithm for QAE that makes use of the Grover-like iteration operator  $\hat{Q}$  but which does not use the quantum Fourier transform and does not require the  $\hat{Q}$  operations to be controlled. Furthermore, Aaronson and Rall gave a rigorous analysis, proving that the same quadratic speedup as in conventional QAE is achieved, albeit with a large constant factor [23].

Around the same time, an entirely different approach was proposed in [27] that, like the algorithm of [18], uses simple circuits consisting of various numbers of iterations of the Grover-like iteration operator Q, but exploits the measurement outcomes via maximum-likelihood classical post-processing in order to make efficient use of the information gained in each measurement. This maximum-likelihood QAE (MLQAE) also achieves the quadratic speedup of the conventional QAE in practice, and seemingly with a smaller constant factor than the algorithm of [18]. However, the analysis makes assumptions that are only approximately true, and thus the formal scaling results are themselves only approximate and without the rigor of previous QAE algorithms. Many extensions to this algorithm have since been proposed, including a version intended for low-circuit-depth scenarios [25] (a different algorithm with more rigorous analysis, based on the Chinese Remainder Theorem, was also presented in this work) and versions which modify the algorithm and analysis to account for noise on the quantum device [25, 28–33, 45]. Despite the formal drawbacks, MLQAE has the potential to be a promising algorithm in practice, and is the focus of the present work. We will detail the MLQAE algorithm in the next subsection.

There have also been other approaches to QAE that are iterative in style, similar to the algorithm of [18], which have reduced the constant factor on the scaling while retaining the rigorous guarantees [19–21, 26, 46, 47] (the algorithms of [27] and [21] are compared experimentally in [48]), including one that is also sensitive to the sign of the amplitude [24]. Other algorithms which have the added advantage of only requiring a single state preparation are presented in [22, 23].



Figure 2. Example circuit for maximum-likelihood quantum amplitude estimation.

#### 2.5. Maximum-likelihood QAE

It has been shown that MLQAE [27] performs as well or better than many of the other QAE algorithms [21]. Due to this, MLQAE has the potential to be a promising QAE algorithm in practice, and so it is worthwhile to find ways to better understand and improve the MLQAE algorithm.

MLQAE involves running a series of quantum circuits of the form shown in Fig. 2, much simpler than the circuit in Fig. 1, and feeding the measurement outcomes into some classical post-processing that includes numerically maximising a likelihood function. Like the conventional QAE algorithm of [15], the MLQAE algorithm begins by applying the algorithm  $\hat{A}$  to the *n*-qubit register (now the only the register) to prepare the state  $|A\rangle$ , and follows by applying some number of iterations of the Grover-like iterator  $\hat{Q}$ . However, here, these applications of  $\hat{Q}$  are not controlled on another register.

It can be shown that performing  $d_j$  (where the index j is introduced for later convenience) applications of the Grover iteration operator  $\hat{Q}$  (which we will refer to as the 'Grover-depth') to the state  $|A\rangle = \sin(\theta_a) |A_G\rangle + \cos(\theta_a) |A_G\rangle$  produces the state

$$\left|\psi_{d_{j}}\right\rangle \equiv \hat{Q}^{d_{j}}\left|A\right\rangle \tag{18}$$

$$= \sin\left[\left(2d_j + 1\right)\theta_a\right] |A_G\rangle + \cos\left[\left(2d_j + 1\right)\theta_a\right] |A_B\rangle.$$
(19)

A computational-basis measurement then produces a 'good' state from G with probability

$$p_{d_i}(\theta_a) = \sin^2 \left[ (2d_j + 1) \,\theta_a \right]. \tag{20}$$

If the state  $|\psi_{d_j}\rangle$  is prepared and measured some number  $N_{\text{shot}}$  times and the number of time a 'good' state is found is denoted as  $h_{d_j}$ , then the likelihood  $L_{d_j}(\theta_a = \theta; h_{d_j})$ of this value of  $h_{d_j}$  as a function of the possible values  $\theta$  for the angle  $\theta_a$  is given by

$$L_{d_j}(\theta_a = \theta; h_{d_j}) = \left[ p_{d_j}(\theta) \right]^{h_{d_j}} \left[ 1 - p_{d_j}(\theta) \right]^{N_{\text{shot}} - h_{d_j}}.$$
 (21)

In the MLQAE algorithm, a number of shots  $N_{\text{shot}}$  and a Grover-depth schedule  $D = \{d_j\}_{j=0}^{q-1}$  are chosen (where the number of depths q can be considered arbitrary at this point), and for each depth  $d_j$  a circuit like the one in Fig. 2 is used to prepare and measure the  $|\psi_{d_j}\rangle$  state  $N_{\text{shot}}$  times, and the measurement record  $\boldsymbol{h} = (h_{d_0}, h_{d_1}, \ldots, h_{q-1})$  is obtained. The individual likelihood functions [Eq. (21)] can then be combined to produce an overall likelihood function for the entire measurement record  $\boldsymbol{h}$ 

$$L(\theta_a = \theta; \mathbf{h}) = \prod_{j=0}^{q-1} L_{d_j}(\theta_a = \theta; h_{d_j})$$
(22)



**Figure 3.** (a) Illustrations of the likelihood  $L_{d_j}(\theta_a = \theta; h_{d_j})$  that the angle  $\theta_a$  is equal to a value  $\theta$ , given that a Grover-like QAE circuit with  $d_j$  Grover iterations  $(d_j \text{ larger for lower plots})$  produced  $h_{d_j}$  'good' states from  $N_{\text{shot}}$  shots. The actual angle  $\theta_a$  that generated the outcomes is indicated by the dashed line. (b) The overall likelihood  $L(\theta_a = \theta; h)$  produced by multiplying together the likelihoods for each  $d_j$ .

Typical forms for the individual likelihood functions  $L_{d_j}(\theta_a = \theta; h_{d_j})$  are illustrated in Fig. 3(a)  $(d_j \text{ larger for lower plots})$ . It can be seen that smaller Groverdepths  $d_j$  produce a likelihood function with few broad peaks, while larger Groverdepths produce a likelihood function with many sharp peaks. The typical form of the overall likelihood function, produced by multiplying together the likelihoods for each  $d_j$ , is illustrated in Fig. 3(b); the combined effect is a single, sharp peak which, if located by a numerical maximisation, will produce an accurate estimate of  $\theta_{\tilde{a}}$  of the angle  $\theta_a$  (which can be turned into an accurate estimate  $\tilde{a}$  for the amplitude a).

2.5.1. Information theory analysis of MLQAE It can be shown (see appendix Appendix A) that the Fisher information carried by the measurement record h regarding the amplitude a is

$$I(a) = \frac{N_{\text{shot}}}{a(1-a)} \sum_{j=0}^{q-1} (2d_j + 1)^2$$
(23)

$$=\frac{N_{\rm shot}}{a(1-a)}(S_D^{(2)})^2,$$
(24)

where in the second line the definition

$$S_D^{(2)} \equiv \sqrt{\sum_{j=0}^{q-1} (2d_j + 1)^2}$$
(25)

has been made.

Assuming the likelihood function satisfies certain regularity conditions detailed in the Bernstein–von Mises theorem [35], which will be explained in more detail in section 4, then as  $N_{\text{shot}}$  increases the estimate  $\tilde{a}$  becomes approximately distributed according to a normal distribution  $\mathcal{N}\left[a, I^{-1}(a)\right]$  with mean equal to the true amplitude a and standard deviation related to the Fisher information as  $I^{-\frac{1}{2}}(a)$ . That is, the average additive error in the estimate  $\tilde{a}$  will be approximately

$$\epsilon_{\text{avg}} \equiv \sqrt{\mathbb{E}\left[(\tilde{a}-a)^2\right]} \\\approx \sqrt{\frac{a(1-a)}{N_{\text{shot}}}} \cdot \frac{1}{S_D^{(2)}}.$$
(26)

The number of calls to the algorithm  $\hat{A}$ , including the call at the beginning of each circuit and the two calls required for each application of the Grover iteration operator  $\hat{Q}$ , is

$$N_{\text{calls}} = N_{\text{shot}} \sum_{j=0}^{q-1} (2d_j + 1)$$
(27)

$$= N_{\rm shot} S_D^{(1)} \tag{28}$$

where in the second line the definition

$$S_D^{(1)} \equiv \sum_{j=0}^{q-1} (2d_j + 1)$$
(29)

has been made. Combining Eq. (26) and Eq. (28) gives

$$\epsilon_{\rm avg} \approx \sqrt{\frac{a(1-a)}{N_{\rm calls}}} \cdot \frac{\sqrt{S_D^{(1)}}}{S_D^{(2)}}.$$
(30)

The primary Grover-depth schedule considered by Suzuki et al. in [27] is the *exponential* schedule

$$D_{\text{EXP}} = \{d_0 = 0\} \cup \{d_j = 2^{j-1}\}_{j=1}^{q-1}.$$
(31)

With this choice, and writing the maximum Grover-depth as  $d \equiv d_{q-2}$ , the quantities  $S_D^{(1,2)}$  can be evaluated (except for the trivial case of d = 0) as

$$S_{D_{\rm EXP}}^{(1)} = 4d + \log_2 d \tag{32}$$

$$=O(d) \tag{33}$$

$$S_{D_{\text{EXP}}}^{(2)} = \sqrt{\frac{16d^2}{3} + 8d + \log_2 d - \frac{10}{3}}$$
(34)

$$=O(d).$$
(35)

Combining the scalings in Eq. (33) and Eq. (35) with Eq. (30) gives

$$\epsilon_{\rm avg} \sqrt{N_{\rm calls}} \approx O\left(\frac{1}{\sqrt{d}}\right).$$
 (36)

Thus, the quadratic speedup of  $\epsilon_{\text{avg}} \approx O\left(N_{\text{calls}}^{-1}\right)$  over simply sampling from the output of the algorithm  $\hat{A}$  can be achieved by allowing (via the choice of the number qof depths) the maximum Grover-depth d to scale linearly with the desired average accuracy as  $d = O\left(\epsilon_{\text{avg}}^{-1}\right)$ . The numerical maximisation of the likelihood function can also be performed in time  $O\left(\epsilon_{\text{avg}}^{-1}\right)$  by simply calculating the likelihood function (or, in practice, its logarithm) at  $O\left(\epsilon_{\text{avg}}^{-1}\right)$  regularly-spaced values of  $\theta$  and choosing the maximum, though this may not be the fastest method in practice.

It is worth noting that in the case where the maximum Grover-depth d scales with the desired average accuracy as  $d = O(\epsilon_{\text{avg}}^{-1})$ , the average error  $\epsilon_{\text{avg}}$  and the target amplitude a would appear in Eq. (30) at the same order, suggesting that a version of this algorithm exists that achieves *relative* error. More explicitly, we find from Eq. (30) that

$$\frac{\epsilon_{\text{avg}}^2}{a(1-a)} N_{\text{calls}} \approx \frac{S_D^{(1)}}{\left(S_D^{(2)}\right)^2}$$
$$= O\left(\frac{1}{d}\right)$$
$$= O\left(\epsilon_{\text{avg}}\right)$$
$$\frac{\epsilon_{\text{avg}}}{a(1-a)} N_{\text{calls}} \approx O(1).$$

Making use of this scaling in practice will require a more sophisticated approach to the numerical likelihood maximisation, such as a grid-search that iteratively focuses more finely on smaller regions of the  $\theta$  domain, to ensure that it can be performed in time  $O\left(\frac{a(1-a)}{\epsilon_{\text{avg}}}\right)$  rather than in time  $O\left(\epsilon_{\text{avg}}^{-1}\right)$ . While this is an interesting observation, we will not consider relative error further in this work.

#### 3. Targeting a desired precision

Suzuki et al.'s analysis [27] shows how the maximum Grover-depth d can be scaled to achieve a quadratic speedup compared to classical sampling, but leaves the number of shots  $N_{\rm shot}$  performed at each Grover-depth as a free parameter. In this section, we extend the analysis to show how, for a given maximum Grover-depth d, the number of shots  $N_{\rm shot}$  can be chosen to approximately target a particular precision  $\epsilon$  with probability at least  $1 - \delta$ .

As mentioned in subsection 2.5, the Bernstein-von Mises theorem implies, assuming the likelihood function satisfies its regularity conditions, that for sufficiently large number of shots  $N_{\text{shot}}$ , the estimate  $\tilde{a}$  becomes approximately distributed according to a normal distribution  $\mathcal{N}\left[a, I^{-1}(a)\right]$  with mean equal to the true amplitude a and standard deviation related to the Fisher information as  $I^{-\frac{1}{2}}(a)$ . The probability that the estimate  $\tilde{a}$  is within a precision  $\epsilon$  of the true amplitude a is thus approximately

$$P(|\tilde{a} - a| \le \epsilon) \approx \frac{1}{I^{-\frac{1}{2}}(a)\sqrt{2\pi}} \int_{a-\epsilon}^{a+\epsilon} dx \exp\left[-\frac{(x-a)^2}{2I^{-1}(a)}\right].$$
 (37)

Writing this probability as  $1 - \delta \equiv P(|\tilde{a} - a| \leq \epsilon)$ , it can be shown that

$$I(a) \approx \frac{2\mathrm{erfinv}^2 \left(1 - \delta\right)}{\epsilon^2},\tag{38}$$

where erfinv is the inverse error function. Combining Eq. 38 and Eq. 24 then gives an expression for the required number of shots

$$N_{\rm shot} \approx a(1-a) \frac{2 {\rm erfinv}^2 (1-\delta)}{\left(S_D^{(2)}\right)^2 \epsilon^2}.$$
(39)

Given that the goal is to find an estimate for the amplitude a, which is not known a priori, we must also maximise  $N_{\text{shot}}$  over all values of the amplitude a. This maximum occurs at a = 0.5 and gives

$$N_{\rm shot} \approx \frac{\operatorname{erfinv}^2 \left(1 - \delta\right)}{2 \left(S_D^{(2)}\right)^2 \epsilon^2}.$$
(40)

Hereafter, we will focus on the regime in which the targeted precision  $\epsilon$  and acceptable failure probability  $\delta$  are small compared to  $1/\left(S_D^{(2)}\right)^2 = O(1/d^2)$  to avoid contradicting the the assumption that the number of shots is sufficiently large for the Bernstein-von Mises theorem to apply. For small  $\delta$ , the erfinv<sup>2</sup>  $(1 - \delta)$  factor can be approximated as  $\ln\left(\frac{\sqrt{\pi}}{\delta}\right)$  [49]. Returning to the exponential schedule  $D_{\text{EXP}}$  from Eq. (31) and relating  $N_{\text{shot}}$  to  $N_{\text{calls}}$  by Eq. (28), this gives a scaling of

$$\epsilon \sqrt{N_{\text{calls}}} \approx O\left(\sqrt{\frac{\ln\left(\frac{1}{\delta}\right)}{d}}\right),$$
(41)

in agreement with Eq. (36).

#### 3.1. Depth-limited scenarios

Relevant to implementations of MLQAE on near- and mid-term hardware, it is important to consider scenarios in which the maximum Grover-depth d cannot be made to scale linearly with the desired precision  $\epsilon$ . In prior work [25], Giurgica-Tiron et al. introduce a parameter  $0 \leq \beta \leq 1$ , and allow the maximum Grover-depth d to scale only as  $O(\epsilon^{-(1-\beta)})$ . Then, a polynomial Grover-depth schedule is chosen as

$$D_{\text{POLY},\beta} = \left\{ d_j = \text{Round}\left(j^{\frac{1-\beta}{2\beta}}\right) \right\}_{j=1}^q,$$
(42)

where  $q = \max\left(\frac{1}{\epsilon^{2\beta}}, \ln\left(\frac{1}{\epsilon}\right)\right)$  so that the schedule approaches an exponential schedule for small  $\beta$  and the maximum Grover-depth  $d \equiv d_q$  becomes approximately  $\frac{1}{\epsilon^{1-\beta}}$ otherwise, as required. Through analysis of the Fisher information content of measurements, Giurgica-Tiron et al. show that that the total number of calls  $N_{\text{calls}}$ to the algorithm  $\hat{A}$  scales with the precision  $\epsilon$  as

$$N_{\text{calls}} \approx O\left(\frac{1}{\epsilon^{1+\beta}}\right).$$
 (43)

From this result, it can be seen that this scheme allows for trade-off between circuit depth and quantum speedup compared to classically sampling from simple measurements of the algorithm  $\hat{A}$ .

However, we observe here that the same trade-off can be achieved using the standard exponential schedule  $D_{\text{EXP}}$  (or a similar one, as will be explained below), by simply choosing the maximum Grover-depth d to scale appropriately as  $O(\epsilon^{-(1-\beta)})$ . In this case, it can be seen from Eq. (41) (or even Eq. (36)) that the number of calls  $N_{\text{calls}}$  scales in the same way as in Eq. (43). If the Grover-depth scaling of  $d = O(\epsilon^{-(1-\beta)})$  is not compatible with  $D_{\text{EXP}}$  as defined in Eq. (31), then once a maximum Grover-depth d is identified for a particular run of MLQAE, a similar schedule can be chosen as

$$D_{\text{EXP},\nu} = \{d_0 = 0\} \cup \{d_j = \text{Round}\left(\nu^{j-1}\right)\}_{j=1}^{q-1},$$
(44)

where the base  $\nu$  and number of depths q are chosen together such that  $\nu$  is the closest value to 2 that supports  $\nu^{q-2} = d$ . Explicit pseudo-code for generating the schedule is given in Algorithm 1. The difference  $|\nu - 2|$  becomes smaller over larger scales of

## Algorithm 1 Generate depth-limited exponential schedule

```
Input: d: Maximum Grover-depth.
Output: D_{\text{EXP},\nu}: Compatible exponential Grover-depth schedule, in
ascending order.
 1: set p_{\max,\text{ceil}} \coloneqq \text{Ceil}[\log_2 d], p_{\max,\text{floor}} \coloneqq \text{Floor}[\log_2 d] \text{ #two options for}
     the maximum power
 2: if p_{\rm max,floor}=0: #discard a maximum power of 0
 3:
            set p_{\max} \coloneqq p_{\max,\text{ceil}}
             set \nu \coloneqq d^{1/p_{\max}}
 4:
 5: else:
            set base_upper := d^{1/p_{\max,\text{ceil}}}, base_lower := d^{1/p_{\max,\text{floor}}}
 6:
            if |base\_upper-2| < |base\_lower-2|:#choose the base closest
 7:
     to 2
 8:
                   set p_{\max} \coloneqq p_{\max,\text{ceil}}, \nu \coloneqq \texttt{base\_upper}
 9:
             else:
                   set p_{\max} \coloneqq p_{\max, \text{floor}}, \nu \coloneqq \texttt{base\_lower}
10:
11: set D_{\text{EXP},\nu} = \{d_0 = 0\} \cup \{d_j = \text{Round}(\nu^j)\}_{i=0}^{p_{\text{max}}}
12: return D_{\text{EXP},\nu}
```

maximum Grover-depth d (see appendix Appendix B), meaning that the approximate scaling results for the standard exponential schedule  $D_{\text{EXP}}$  shown in subsubsection 2.5.1 and in this section still hold approximately.

An important case to consider is that of  $\beta = 1$ ; that is, when the maximum Grover-depth d is fixed and not allowed to scale with the desired precision  $\epsilon$ . This case is relevant to near- and mid-term hardware for which the useful circuit depth, and thus the achievable Grover-depth d for MLQAE on a particular algorithm  $\hat{A}$ , will have some upper-limit on any particular device due to interactions with the environment. Indeed, even specific devices employing error-correction will have some limit on the achievable Grover-depth d, as the available number of qubits will limit the size of error-correcting codes [50]. In this case, as the maximum Grover-depth d is constant, it is clear from the approximate scaling in Eq. (41) that the required number of calls  $N_{\text{calls}}$  will scale inversely-proportionally to the square  $\epsilon^2$  of the desired precision, as in the case of classical sampling; that is, no scaling advantage is achieved. However, it is also the case in this case that the achievable Grover-depth d determines the constant factor on this scaling. Combining Eq. (28) and Eq. (40) gives the required number of calls  $N_{\text{calls}}$  to be

$$N_{\text{calls}} \approx \frac{S_{D_{\text{EXP},\nu}}^{(1)}}{\left(S_{D_{\text{EXP},\nu}}^{(2)}\right)^2} \cdot \frac{\text{erfinv}^2 \left(1-\delta\right)}{2\epsilon^2}.$$
(45)

The speed-up factor  $R_d \equiv N_{\text{calls}}^{(\text{classical})}/N_{\text{calls}}$  achieved by using a maximum Groverdepth *d* compared to classical sampling case (for which d = 0 and  $S_D^{(1)} = S_D^{(2)} = 1$ ), can be quantified approximately as

$$R_d \approx \frac{\left(S_{D_{\text{EXP},\nu}}^{(2)}\right)^2}{S_{D_{\text{EXP},\nu}}^{(1)}}.$$
(46)

Taking for simplicity the case where the maximum Grover-depth d is a power of 2, this can be rewritten as

$$R_d \approx \frac{\frac{16d^2}{3} + 8d + \log_2 d - \frac{10}{3}}{4d + \log_2 d} \tag{47}$$

$$=O(d).$$
(48)

The meaning of Eq. (48) is that, assuming the Bernstein-von Mises theorem holds, then using a quantum processor that supports a maximum Grover-depth d (for the particular algorithm  $\hat{A}$ ) to run MLQAE should produce an estimate  $\tilde{a}$  for the amplitude a within precision  $\epsilon$  with probability approximately greater than  $1 - \delta$  with a O(d) factor speed-up compared to simply classically sampling from the algorithm  $\hat{A}$ . If we return to the case where the maximum Grover-depth d can scale with the targeted precision  $\epsilon$  as  $d = O(\epsilon^{-1})$ , then a quadratic speedup over the classical scaling of  $N_{\text{calls}}^{(\text{classical})} = O(\epsilon^{-2})$  is achieved, as expected.

In practice, even when limiting the circuit-depth to keep the noise under control, low levels of noise may still affect the results and place a limit on the achievable precision  $\epsilon$ . Methods to mitigate this issue by modifying the the likelihood functions  $L_{d_j}(\theta = \theta_a; h_{d_j})$  to include the effects of noise are discussed in [25, 27–29, 32], while other QAE-specific noise-mitigation methods are discussed in [30, 31, 33, 45]. In this work, we do not consider the effects of noise beyond the notion of a depth limitation.

#### 3.2. Numerical validation

To test the analysis presented in this section, we have performed numerical simulations of the MLQAE algorithm. Given the simple description of the action of the operators  $\hat{Q}^{d_j}$  in terms of the angle  $\theta_a$  given in Eq. 19, it is not necessary to perform full quantum-circuit simulations in order to simulate MLQAE for a particular pre-chosen value of the amplitude *a*. Instead, for each Grover-depth  $d_j$  in the chosen schedule *D*, we calculate the 'good' state probability  $p_{dj}$  ( $\theta_a$ ), according to Eq. (20), draw a number



Figure 4. The error in an estimate  $\tilde{a}$ , produced from a single run of the MLQAE algorithm, of the amplitude a for 1 000 000 equally-spaced values of the amplitude a covering the full range from a = 0 to a = 1, using maximum Grover-depth d, targeting a precision  $\epsilon$  with failure probability at most  $\delta$  and a number of shots  $N_{\rm shot}$  calculated according to Eq. (40). The black, horizontal dashed line is the target precision  $\epsilon$ , while the vertical dashed lines indicate the exceptional values according to Eq. (49), where the blue (orange) lines are fore even (odd) j. The parameters for the two plots are (a) d = 16,  $\epsilon = 10^{-3}$ ,  $\delta = 0.01$  (resulting in  $N_{\rm shot} = 1111$ ) and (b) d = 50,  $\epsilon = 10^{-4}$ ,  $\delta = 0.01$  (resulting in  $N_{\rm shot} = 11688$ ).

of 'good' states  $h_{d_j}$  from the binomial distribution  $\mathcal{B}\left[N_{\text{shot}}, p_{d_j}\left(\theta_a\right)\right]$ , and collect these together into the measurement record  $\boldsymbol{h}$ . From this measurement record, we construct the likelihood function  $L(\theta_a = \theta; \boldsymbol{h})$  according to Eq. (22) and numerically maximise its logarithm by simply calculating it at  $3/\epsilon$  regularly-spaced values of  $\theta$  (where the 3 has been chosen arbitrarily) and finding the maximum.

Fig. 4 shows, for two sets of parameters, the error in the estimate  $\tilde{a}$ , produced from a single run of the MLQAE algorithm, of the amplitude *a* for 1 000 000 equallyspaced values of the amplitude *a* covering the full range from a = 0 to a = 1, targeting a precision  $\epsilon$  with failure probability at most  $\delta$ , and using using maximum Groverdepth *d* and a number of shots  $N_{\text{shot}}$  calculated according to Eq. (40). The parameters for the two plots are (a) d = 16,  $\epsilon = 10^{-3}$ ,  $\delta = 0.01$  (resulting in  $N_{\text{shot}} = 1111$ ) and (b) d = 50,  $\epsilon = 10^{-4}$ ,  $\delta = 0.01$  (resulting in  $N_{\text{shot}} = 11688$ ).

It can be seen in both plots of Fig. 4 that, for most typical values of the amplitude a, MLQAE produces an estimate with an error that is either close to or much smaller than the target precision  $\epsilon$ . However, it is also clear that some small, regularly-spaced regions of the amplitude a exist in which the estimates  $\tilde{a}$  produced are significantly worse. These exceptional values occur in the vicinity of a-values for which the circuit with the maximum Grover-depth d produces 'good' states with probability 1 or 0; that is, the exceptional values occur in the vicinity of

$$a_k^{(d \to p_d(\theta_a)=0,1)} = \sin\left[\frac{k\pi}{2(2d+1)}\right]^2 \text{ for } k = 0, 1, \dots, 2d, 2d+1.$$
(49)

The reason for this, as well as precisely what is meant by "in the vicinity of", will be explained in detail in the next section, but for now, this observation allows us to examine the exceptional values more closely.

In Fig. 5, each data point represents the minimum precision  $\epsilon_{\text{achieved}}$  achieved with probability at least  $1 - \delta$ , estimated from 10 000 runs of MLQAE, for various values of the number of shots  $N_{\text{shot}}$  and for various values of the amplitude a. The  $N_{\rm shot}$  values used are powers of 2, as well as an extra data point for the  $N_{\rm shot}$  calculated according to Eq. (40) for a particular target precision  $\epsilon$ , which are also indicated in each plot by the vertical dotted and horizontal dashed lines respectively. The left column uses the parameters d = 16,  $\epsilon = 10^{-3}$ ,  $\delta = 0.01$ , while the right column uses the parameters d = 50,  $\epsilon = 10^{-4}$ ,  $\delta = 0.01$ . The top (bottom) row shows some *typical* (*exceptional*) values of the amplitude a. For each plot in Fig. 5, we performed the numerical likelihood maximisation with a grid of  $3/\epsilon_{\rm min}$  equally-spaced points, where  $\epsilon_{\rm min}$  is the precision expected for the largest value of  $N_{\rm shot}$  plotted. It can be seen that for the typical values of the amplitude a, the target precision  $\epsilon$  is achieved at values of  $N_{\rm shot}$  that are lower than or close to the value predicted by Eq. (40). For the exceptional values of the amplitude a, however, it can be seen that values of  $N_{\rm shot}$ that are between 2 and 4 times larger than the value predicted by Eq. (40) are required to achieve the target precision  $\epsilon$ . This is consistent with the plots in Fig. 4, and the reasons for these exceptional values will be explored in the next section.

A similar observation was made in [29] in the case of MLQAE in a noisy setting, where a multiparameter likelihood maximisation is performed that includes the noise rate as a nuissance parameter. In that work, 'anomalous target' values of the amplitude a are observed for which the algorithm fails to achieve the desired precision; however, those anomalous target values are identified as being caused by a failure of the multiparameter likelihood maximisation procedure to accurately estimate the noise rate, which is not relevant in the noiseless case where the likelihood-maximisation is performed only in one dimension. Thus, the exceptional values seen here are a distinct phenomenon to the anomalous target values seen in [29].

#### 4. Exceptional values

In subsection 3.2, we showed numerically that choosing the number of shots  $N_{\rm shot}$  according to Eq. (40) approximately achieves the target precision  $\epsilon$  with probability at least  $1 - \delta$  for *typical* values of the amplitude a, but that it fails for particular exceptional values in the vicinity of  $a = a_k^{(d \to p_d(\theta_a)=0,1)}$  for which the maximum Grover-depth d produces 'good' states with probability 1 or 0. To see why this problem occurs, it is helpful to examine the likelihood-maximisation process at a conceptual level.

To this end, Fig. 6 shows a qualitative illustration of the emergence of exceptional in the vicinity of  $a = a_k^{(d \to p_d(\theta_a)=0,1)}$  [see Eq. (49)]. The top row illustrates the case of typical values of the amplitude a (and of the corresponding angle  $\theta_a$ ). Fig. 6(a) illustrates the probability  $p_d(\theta_a)$  of measuring a 'good' state after running the circuit with the maximum Grover-depth, d. The true value of the angle  $\theta_a$  is indicated by the black cross and the vertical dashed line, while other angles that produce the same probability are indicated by grey crosses. As the typical values occur far from  $p_d(\theta_a) = 0$  and  $p_d(\theta_a) = 1$ , the corresponding Grover-depth d likelihood function  $L_d$ typically looks as in Fig. 6(b), where the likelihood peaks correspond to the crosses from plot (a). The Grover-depth d measurement data does not contain the information to determine which of the peaks correspond to the true angle  $\theta_a$ , but the information gained from the lower Grover-depth circuits is typically sufficient to narrow  $\theta$  down to a small range, indicated by the shaded area in plot (b), that contains only one of these peaks.

However, the bottom row illustrates the case of an exceptional value, which in



Figure 5. Each data point represents the minimum precision  $\epsilon_{achieved}$  achieved with probability at least  $1-\delta$ , estimated from 10000 runs of MLQAE, for various values of the number of shots  $N_{\rm shot}$  and for various values of the amplitude a. The  $N_{\rm shot}$  values used are powers of 2, as well as an extra data point for the  $N_{\rm shot}$  calculated according to Eq. (40) for a particular target precision  $\epsilon$ , which are also indicated by the vertical dotted and horizontal dashed lines respectively. The left column ((a) and (c)) uses the parameters d = 16,  $\epsilon = 10^{-3}$ ,  $\delta = 0.01$ , while the right column ((b) and (d)) uses the parameters d = 50,  $\epsilon = 10^{-4}$ ,  $\delta = 0.01$ . The top row ((a) and (b)) shows typical values of the amplitude a, chosen as  $a = 0.1, 0.2, \ldots, 0.8, 0.9$  as well as the results for choosing a randomly from the full range for each of the  $10\,000$  runs. The bottom row ((c) and (d))shows exceptional values of the amplitude a, chosen as  $a = a_k^{(d \to p_d(\theta_a) = 0, 1)} + \epsilon$ see [Eq. (49)], with (c) k = 14, 15, 16, 17, 18 and (d)  $k = 45, 46, \dots, 54, 55$  in order to be reasonably close to a = 0.5. For each plot, we performed the numerical likelihood maximisation with a grid of  $3/\epsilon_{\min}$  equally-spaced points, where  $\epsilon_{\min}$ is the precision expected for the largest value of  $N_{\rm shot}$  plotted.

this example is close to  $p_d(\theta_a) = 1$ . As some pairs of crosses occur close together in Fig. 6(c), some pairs of peaks occur close together in the corresponding likelihood function in Fig. 6(d), including two peaks that both appear in the shaded region, meaning that it is not possible to determine which of the two peaks correspond to the true angle  $\theta_a$ . If the distance  $|\tilde{a} - \tilde{a}_{\text{False}}|$  between the correct estimate  $\tilde{a}$  and the nearby 'false' estimate  $\tilde{a}_{\text{False}}$  is much smaller than the target precision, that is if  $|\tilde{a} - \tilde{a}_{\text{False}}| \ll \epsilon$ , then either estimate is likely to work well, and this ambiguity will not cause an issue. In the opposite extreme, if  $|\tilde{a} - \tilde{a}_{\text{False}}| \gtrsim 2\epsilon$ , then the false estimate  $\tilde{a}_{\text{False}}$  is unlikely to be in the ambiguous region (illustrated by the shaded regions in Fig. 6), as  $2\epsilon$  is approximately the precision that would be achieved if the data from the maximum Grover-depth d runs were not included in the analysis.



**Figure 6.** A qualitative illustration of the emergence of exceptional values of the amplitude *a* in the vicinity of  $a = a_k^{(d \to p_d(\theta_a)=0,1)}$  [see Eq. (49)]. The top row illustrates the case of typical values of the amplitude *a* (and of the corresponding angle  $\theta_a$ ). Plot (a) illustrates the probability  $p_d$  of measuring a 'good' state after running the circuit with the maximum Grover-depth, d. The true value of the angle  $\theta_a$  is indicated by the black cross and the vertical dashed line, while other angles that produce the same probability are indicated by grey crosses. As the typical values occur far from  $p_d = 0$  and  $p_d = 1$ , the corresponding Grover-depth d likelihood function  $L_d$  typically looks as in plot (b), where the likelihood peaks correspond to the crosses from plot (a). The Grover-depth d measurement data does not contain the information to determine which of the peaks correspond to the true angle  $\theta_a$ , but the information gained from the lower Grover-depth circuits is typically sufficient to narrow  $\theta$  down to a small range, indicated by the shaded area in plot (b), that contains only one of these peaks. The bottom row ((c) and (d)) illustrates the case of exceptional values, in this example close to  $p_d = 1$ . As some pairs of crosses occur close together in plot (c), some pairs of peaks occur close together in the corresponding likelihood function in plot (d), including two peaks that both appear in the shaded region, meaning that it is not possible to determine which of the two peaks correspond to the true angle  $\theta_a$ .

The exceptional values are thus expected to occur roughly in the intermediate regime, where  $\epsilon \leq |\tilde{a} - \tilde{a}_{\text{False}}| \leq 2\epsilon$ . In Fig. 7, this corresponds to the expectation that the exceptional values should be approximately contained in the region of medium darkness. This intuition is approximately confirmed, although the exceptional values do extend further out in to the lightest shaded region.

This can also be understood in terms of the Bernstein-von Mises theorem. The starting point of the derivation of the required number of shots  $N_{\text{shot}}$  in Eq. (40) was the assumption that the regularity conditions set out in the Bernstein-von Mises theorem hold. The theorem is neatly explained and summarised by Giurgica-Tiron et al. in [25], and the conditions can be expressed most simply as the following three requirements:

(i) The likelihood function  $L(\theta_a = \theta; h)$  and its logarithm should be smooth over the



**Figure 7.** The red line shows the precision  $\epsilon_{\text{achieved}}$  achieved with probability  $1 - \delta = 99\%$  (estimated from 10 000 MLQAE runs), when targeting a precision  $\epsilon = 10^{-4}$ , indicated by the horizontal dashed line, using an exponential Grover-depth schedule with maximum Grover-depth d = 50, at 1000 different values of the amplitude a in the vicinity of  $a_k^{(d \to p_d(\theta_a)=0,1)}$  with k = 50, indicated by the vertical dotted line, and the shaded regions are centred on this value. The numerical likelihood maximisation has been performed by calculating the likelihood function at  $\sim 30/\epsilon$  equally spaced values of  $\theta$ , a much finer search than is needed in practice in order to reduce the size of the fluctuations in the plot. The darkest shaded regions have width  $2\epsilon$  and  $4\epsilon$  respectively.

relevant interval of  $\theta$ .

- (ii) The Fisher information I(a) should be smooth over the relevant interval of a (and thus of  $\theta$ ). This implies that the logarithm of the likelihood function should be at least thrice differentiable
- (iii) The first derivative of the prior distribution should exist and be continuous.

In the present case, the prior distribution is uniform, and so condition iii is trivially satisfied. However, it can be seen from Eq. (21), Eq. (22) and Fig. 3 that the likelihood function  $L(\theta_a = \theta; \mathbf{h})$  goes to 0 at the amplitude values  $a_k^{(d_j \to P_{\text{good}} = 0,1)}$  if the measurement record  $h_{d_j}$  contains both 'good' and 'bad' outcomes, meaning that at these points the logarithm of the likelihood function diverges to negative infinity, violating conditions i and ii. If the information gained from measurements at lower Grover-depths is sufficient to restrict the likely estimates  $\tilde{a}$  to an interval around the amplitude a that does not include any  $a_k^{(d_j \to p_d(\theta_a)=0,1)}$  values, the validity of conditions i and ii is recovered in this interval, but if this is not the case than the Bernstein-von Mises theorem does not apply, explaining why using a number of shots  $N_{\text{shot}}$  according to Eq. (40) is not even approximately sufficient for the exceptional values. From this perspective, exceptional values should occur near the  $a_k^{(d_j \to p_d(\theta_a)=0,1)}$  at any Grover-depth  $d_j$  in the depth schedule, but in practice those at depths smaller than the maximum depth d have a much less significant impact due to the lower information content of measurements at those depths.

## 5. Depth jittering heuristic

As discussed in section 4, exactly which values of the amplitude value a are exceptional depends on exactly which Grover-depths appear in the Grover-depth schedule D, with a particular importance associated with the maximum depth d. With this in mind, we propose a heuristic method to produce a modified Grover-depth schedule D' in which the  $N_{\text{shot}}$  shots for some Grover-depths  $d_j \in D_{\text{EXP},\nu}$  are instead spread out over a small number of nearby depths, in order to mitigate the effects of exceptional values. We refer to this process as 'depth jittering', and the jittered Grover-depth schedule D' can be produced from the original Grover-depth schedule D via Algorithm 2.

Algorithm	<b>2</b>	Depth	jittering	algorithm
-----------	----------	-------	-----------	-----------

<b>Inputs</b> : <i>D</i> : Original Grover-depth schedule, with size at least 2, in ascending				
order.				
c: Logarithmic spread coefficient.				
<b>Outputs</b> : $D'$ : Jittered Grover-depth schedule, in ascending order.				
$F^{(\text{shot})}$ : Corresponding shot fractions.				
1: initialise $D' \coloneqq \{\}, F^{(\text{shot})} \coloneqq \{\}$				
2: for $d_j$ in reverse $(D)$ : #reverse order to prioritise jittering				
larger depths				
3: initialise jitter $\coloneqq$ False				
4: calculate depth_spread := Round $[\ln (cd_j)]$				
5: <b>if</b> $d_j = \max(D)$ :				
6: set lower $\coloneqq d_j - \texttt{depth\_spread}$				
7: set upper := $d_j$ #don't go beyond maximum depth				
8: <b>if</b> lower $> d_{j-1} + 1$ : set jitter $\coloneqq$ True				
elif $\min(D) < d_j < \max(D)$ :				
10: set lower := $d_j$ - depth_spread and upper := $d_j$ +				
depth_spread				
11: if lower $> d_{j-1} + 1$ and upper $< \min(D') - 1$ : set jitter :=				
True				
12: elif $d_j = \min(D)$ and not $d_j = 0$ : #jittering at depth 0 not				
necessary				
13: set lower := $\max(0, d_j - \text{depth}_s\text{pread}) \# \text{don't go below } 0$				
depth				
14: set upper $\coloneqq d_j + \texttt{depth\_spread}$				
15: <b>if</b> upper $< \min(D') - 1$ : set jitter := True				
16: if jitter:				
17: for $(d_j)'_k$ from upper to lower: #add in descending order to				
be reversed at end				
18: add $(d_j)'_k$ to $D'$ and $1/(upper - lower + 1)$ to $F^{(shot)}$				
19: else:				
20: add $d_j$ to $D'$ and 1 to $F^{(\text{shot})}$				
21: set $D' \coloneqq \texttt{reverse}(D')$ and $F^{(\text{shot})} \coloneqq \texttt{reverse}(F^{(\text{shot})})$				
22: return $D'$ and $F_{\text{shot}}$				

In short, Algorithm 2 takes each Grover-depth  $d_j$  in the original schedule D, as well as a 'logarithmic spread coefficient' c and adds all the depths between

~  $(d_j - \log(cd_j))$  and ~  $(d_j + \log(cd_j))$  to a new 'jittered' schedule D'. This is not done for zero-depth  $d_j = 0$ , as this depth has no exceptional values. This is also not done for depths  $d_j$  where this spread would overlap with the spreads from neighbouring depths  $d_{j-1}$  and  $d_{j+1}$ , where jittering the larger depths is prioritised in these cases. For the maximum Grover-depth d, only the part of the spread that is less than or equal to d is included, to avoid exceeding the maximum Grover-depth, while for the minimum Grover-depth the spread is truncated if necessary to avoid impossible depths below 0.

The use of depth-spreads of logarithmic size is a heuristic choice, motivated by the need to preserve the approximate scaling results presented in previous sections. Algorithm 2 also outputs a set of 'shot fractions'  $F^{(\text{shot})}$  that simply keeps track of the factor by which  $N_{\text{shot}}$  should be scaled down at each new of the new depths  $d'_j \in D'$ in order to effectively spread the  $N_{\text{shot}}$  shots that would have originally been run at depth  $d_j$  over the relevant jittered depths. Taking the examples of exponential schedules  $D_{\text{EXP},\nu}$  with maximum depths d = 16 and d = 50 (as in Fig. 4 and Fig. 5) and setting the spread coefficient c = 2 (this choice is arbitrary, but seems to work well in practice), the original schedules are

$$D_{\text{EXP},\nu} = \{0, 1, 2, 4, 8, 16\}$$
(50)

$$D_{\text{EXP},\nu} = \{0, 1, 2, 4, 7, 14, 26, 50\},\tag{51}$$

while the jittered schedules are

$$D'_{\text{EXP},\nu} = \{0; 1; 2; 4; 8; 13, 14, 15, 16\}$$

$$D'_{\text{EXP},\nu} = \{0; 1; 2; 4; 7; 11, 12, 13, 14, 15, 16, 17;$$
(52)

$$22, 23, 24, 25, 26, 27, 28, 29, 30; 45, 46, 47, 48, 49, 50\},$$
(53)

where semicolons have been used to separate the jittered depths into groups corresponding to the original depths.

It is also necessary to define jittered versions of the quantities  $S_D^{(1)}$  and  $S_D^{(2)}$  from Eq. (29) and Eq. (29). These are simply

$$S_{D'}^{(1)'} \equiv \sum_{j=0}^{q-1} F_j^{(\text{shot})}(2d'_j + 1)$$
(54)

and

$$S_{D'}^{(2)'} \equiv \sqrt{\sum_{j=0}^{q-1} F_j^{(\text{shot})} (2d'_j + 1)^2}.$$
(55)

With these definitions, the jittered number of shots  $N'_{\text{shot}}$  can be calculated according to Eq. (40), which is in general slightly different from the original  $N_{\text{shot}}$ . When the jittered MLQAE algorithm is run, the actual number of shots used at depth  $d'_j$  is  $\lceil F_j^{(\text{shot})} N'_{\text{shot}} \rceil$ .

To demonstrate that the extent of the speedup for MLQAE compared to classically sampling from the outputs of the algorithm  $\hat{A}$  is not significantly affected by the depth-jittering procedure, Fig. 8 shows the ratio of the number of calls  $N'_{\text{calls}}$  and  $N_{\text{calls}}$  to the algorithm  $\hat{A}$  in the case of a jittered and unjittered Grover-depth schedule



Figure 8. The ratio of the number of calls  $N'_{\text{calls}}$  and  $N_{\text{calls}}$  to the algorithm  $\hat{A}$  in the case of a jittered and unjittered Grover-depth schedule respectively as a function of maximum Grover-depth d, for three different choices of the target precision  $\epsilon$ . An acceptable failure probability of  $\delta = 0.01$  is used in all cases. The visible increase at the highest depths for the  $\epsilon = 10^{-4}$  case (solid green line) occurs because of the ceiling function applied in the expression  $[F_j^{(\text{shot})}N'_{\text{shot}}]$  for the actual number of shots used at depth  $d'_j$ , as  $N'_{\text{shot}}$  becomes small.

respectively as a function of maximum Grover-depth d, for three different choices of the target precision  $\epsilon$ . An acceptable failure probability of  $\delta = 0.01$  is used in all cases. The visible increase at the highest depths for the  $\epsilon = 10^{-4}$  case (solid green line) occurs because of the ceiling function applied in the expression  $\lceil F_j^{(\text{shot})} N'_{\text{shot}} \rceil$  for the actual number of shots used at depth  $d'_j$ , as  $N_{\text{shot}}$  becomes small. For the more strict precision requirements (the dashed red and dotted blue lines) where  $N_{\text{shot}}$  is not small at the larger Grover-depths, the ratio is convincingly tending toward a value of 1.

#### 5.1. Numerical validation

To show that the depth-jittering heuristic is effective, we have repeated the numerical experiments described in subsection 3.2 with exactly the same parameters, but using the depth-jittered schedule  $D'_{\text{EXP},\nu}$  instead of the original schedule  $D_{\text{EXP},\nu}$ . Fig. 9 should be compared with Fig. 4. It can be seen that the sharp features associated with exceptional values of the amplitude *a* have been largely removed by the depth-jittering procedure. Likewise, Fig. 10 should be compared with Fig. 5. It can be seen that the MLQAE algorithm with the depth-jittering heuristic now achieves the target precision  $\epsilon$  at an  $N'_{\text{shot}}$  value much closer to the value predicted by Eq. (40) for the exceptional values of the amplitude *a* (Fig. 10(c) and Fig. 10(d)), without significantly affecting the performance for the typical values of the amplitude *a* (Fig. 10(a) and Fig. 10(b)).

## 6. Summary and further work

In this paper, we have given new insights and made improvements to MLQAE, the maximum-likelihood quantum amplitude estimation algorithm first proposed by



**Figure 9.** Data produced by exactly the same procedure and with the same parameters as in Fig. 4, but with a jittered Grover-depth schedule  $D'_{\text{EXP},\nu}$ . The sharp features in Fig. 4 associated with exceptional values of the amplitude *a* have been largely removed by the depth-jittering procedure.



**Figure 10.** Data produced by exactly the same procedure and with the same parameters as in Fig. 5, but with a jittered Grover-depth schedule  $D'_{\text{EXP},\nu}$ . The MLQAE algorithm now achieves the target precision  $\epsilon$  at an  $N'_{\text{shot}}$  value much closer to the value predicted by Eq. (40) for the exceptional values of the amplitude a [(c) and (d)], without significantly affecting the performance for the typical values of the amplitude a [(a) and (b)].

Suzuki et al. [27]. Quantum amplitude estimation has the potential to be an important subroutine in quantum algorithms, offering quadratic speedups in many areas of science, engineering, and industry through its applications to quantum-enhanced Monte Carlo simulation [6], as well as to quantum machine learning [9, 13]. The maximum-likelihood-based quantum amplitude estimation algorithm is one of a number of recent proposals that use much simpler quantum circuits than earlier work [15], bringing quantum amplitude estimation closer to suitability for near- and mid-term quantum hardware.

The analysis in prior work focused on understanding how the average error scales with the number of calls to the quantum circuit whose output contains the amplitude we are trying to estimate, observing a quadratic speedup compared to simple classical sampling in the process. In section 3, we began our contribution by extending this analysis and showed how to calculate how many calls are required to target a particular precision with a particular choice of probability, rather than simply on average. As with the prior work, our analysis assumed that the conditions required for the Bernstein-von Mises theorem to hold are satisfied, and was therefore not completely rigorous. We also considered scenarios in which the circuit-depth cannot be made arbitrarily large, arguing that a trade-off between circuit-depth and quantum speedup developed in [25] can be achieved in a more a direct way.

While numerically validating our analysis in section 3 over the full range of target amplitudes, we observed that there are particular values of the amplitude, *exceptional values*, at which the algorithm fails to achieve the target precision. In section 4, we described at which values of the amplitude these occur, and explained in detail why they occur, both formally in terms of violations of the Bernstein-von Mises theorem and via a more intuitive picture in terms of the likelihood maximisation process.

Having explained in section 4 how the amplitudes which are the problematic exceptional values depended on the 'Grover-depth' (the number of applications of a 'Grover-like' operator present in the circuits that are run), we proposed in section 5 a heuristic method to mitigate the problem, which we refer to as 'depth-jittering'. This method involves taking the shots that would ordinarily all be performed on identical circuits at a single Grover-depth, and instead spreading them over a small number of nearby Grover-depths in order to bypass much of the sensitivity to Grover-depth. We show numerically that this process has only a negligible effect on the algorithm runtime, and that the algorithm now achieves approximately the expected precision across the full range of target amplitudes and the previously exceptional values are no longer problematic.

Our work leads naturally to a number of questions. Most importantly, the only way in which we consider the effects of the noise that would be unavoidably present in near- and mid-term implementations of this algorithm is by imposing limitations on the available circuit depths, but prior work [25, 27–29, 32] has shown how to modify the likelihood functions to account for the noise on the quantum device. It will be important to understand how these modifications would affect the exceptional values, as well as the effectiveness of our depth-jittering heuristic. Another possible extension to our work is exploring the possibility of performing 'adaptive' jittering, perhaps by detecting possible exceptional values through the low-depth measurements and only then jittering the high-depth runs if necessary. The runs at each Grover-depth could still be performed in parallel, and so this advantage of MLQAE is not lost. A further possible extension is to investigate whether there are numerical factors by which the number of measurement shots can be scaled up to reliably guard against the approximate nature of the theoretical analysis. Overall, the insights we offer in this work help to improve the understanding of the maximum-likelihood quantum amplitude estimation algorithm, and move it toward being a useful tool for practical algorithms in the near- and mid-term.

#### Acknowledgments

AC thanks Oscar Higgott, George Umbrarescu, and Tom Dodd for helpful and interesting discussions. The authors gratefully acknowledge support from the UKRI EPSRC Grant No. EP/T001062/1 via the UK Quantum Technology Hub in Computing and Simulation.

#### Appendix A. Fisher information content of measurement record

In this appendix, we show that that the Fisher information carried by the measurement record h regarding the amplitude a is

$$I(a) = \frac{N_{\text{shot}}}{a(1-a)} \sum_{j=0}^{q-1} (2d_j + 1)^2,$$

as in Eq. (23).

The probability of measuring a 'good' state after performing the MLQAE circuit of Grover-depth  $d_j$  as a function of the amplitude a is

$$p_{d_j}(a) = \sin^2 \left[ (2d_j + 1)\theta_a \right]$$
$$= \sin^2 \left[ (2d_j + 1) \arcsin \sqrt{a} \right]$$

Denoting the outcome of the single measurement as  $m_{d_j} = 1,0$  for a 'good' or 'bad' outcome respectively, the likelihood of the estimate  $\tilde{a}$  being the correct amplitude a can be written as

$$L_{d_j}^{\text{(single shot)}}\left(\tilde{a}=a;m_{d_j}\right) = \left[p_{d_j}(\tilde{a})\right]^{m_{d_j}} \left[1-p_{d_j}(\tilde{a})\right]^{1-m_{d_j}}$$

In the region around the true amplitude a where the logarithm of the likelihood function is differentiable, the Fisher information gained from the measurement can be written as

$$I^{(\text{single shot})}(a) = \mathbb{E}_{m_{d_j}} \left[ \left( \frac{\partial}{\partial \tilde{a}} \log L_{d_j}^{(\text{single shot})} \right)^2 \middle| a \right]$$

where the expectation value is being taken over the measurement outcome  $m_{d_j}$  and then evaluated at the true amplitude a.

The derivative can be developed as

$$\begin{split} \frac{\partial}{\partial \tilde{a}} \log L_{d_j}^{(\text{single shot})} &= \frac{\frac{\partial}{\partial \tilde{a}} L_{d_j}^{(\text{single shot})}}{L_{d_j}^{(\text{single shot})}} \\ &= \frac{\frac{\partial}{\partial \tilde{a}} \left( \left[ p_{d_j}(\tilde{a}) \right]^{m_{d_j}} \right) \left[ 1 - p_{d_j}(\tilde{a}) \right]^{1 - m_{d_j}} + \left[ p_{d_j}(\tilde{a}) \right]^{m_{d_j}} \frac{\partial}{\partial \tilde{a}} \left( \left[ 1 - p_{d_j}(\tilde{a}) \right]^{1 - m_{d_j}} \right)}{\left[ p_{d_j}(\tilde{a}) \right]^{m_{d_j}}} \\ &= \frac{\frac{\partial}{\partial \tilde{a}} \left( \left[ p_{d_j}(\tilde{a}) \right]^{m_{d_j}} \right)}{\left[ p_{d_j}(\tilde{a}) \right]^{m_{d_j}}} + \frac{\frac{\partial}{\partial \tilde{a}} \left( \left[ 1 - p_{d_j}(\tilde{a}) \right]^{1 - m_{d_j}} \right)}{\left[ 1 - p_{d_j}(\tilde{a}) \right]^{1 - m_{d_j}}} \\ &= m_{d_j} \frac{\frac{\partial}{\partial \tilde{a}} \left[ p_{d_j}(\tilde{a}) \right]}{p_{d_j}(\tilde{a})} + (1 - m_{d_j}) \frac{\frac{\partial}{\partial \tilde{a}} \left[ 1 - p_{d_j}(\tilde{a}) \right]}{1 - p_{d_j}(\tilde{a})} \end{split}$$

where the last line is valid because the  $m_{d_j}$  is a binary variable. The two remaining derivatives can be developed as

$$\begin{split} \frac{\partial}{\partial \tilde{a}} \left[ p_{d_j}(\tilde{a}) \right] &= \frac{\partial}{\partial \tilde{a}} \left( \sin^2 \left[ (2d_j + 1) \arcsin \sqrt{\tilde{a}} \right] \right) \\ &= 2(2d_j + 1) \sin \left[ (2d_j + 1) \arcsin \sqrt{\tilde{a}} \right] \cos \left[ (2d_j + 1) \arcsin \sqrt{\tilde{a}} \right] \frac{\partial}{\partial \tilde{a}} \left( \arcsin \sqrt{\tilde{a}} \right) \\ &= \frac{2d_j + 1}{\sqrt{\tilde{a}(1 - \tilde{a})}} \sin \left[ (2d_j + 1) \arcsin \sqrt{\tilde{a}} \right] \cos \left[ (2d_j + 1) \arcsin \sqrt{\tilde{a}} \right] \\ &= \frac{2d_j + 1}{\sqrt{\tilde{a}(1 - \tilde{a})}} \sqrt{p_{d_j}(\tilde{a}) \left[ 1 - p_{d_j}(\tilde{a}) \right]} \end{split}$$

and

$$\begin{split} \frac{\partial}{\partial \tilde{a}} \left[ 1 - p_{d_j}(\tilde{a}) \right] &= \frac{\partial}{\partial \tilde{a}} \left( \cos^2 \left[ (2d_j + 1) \arcsin \sqrt{\tilde{a}} \right] \right) \\ &= -2(2d_j + 1) \cos \left[ (2d_j + 1) \arcsin \sqrt{\tilde{a}} \right] \sin \left[ (2d_j + 1) \arcsin \sqrt{\tilde{a}} \right] \frac{\partial}{\partial \tilde{a}} \left( \arcsin \sqrt{\tilde{a}} \right) \\ &= -\frac{2d_j + 1}{\sqrt{\tilde{a}(1 - \tilde{a})}} \cos \left[ (2d_j + 1) \arcsin \sqrt{\tilde{a}} \right] \sin \left[ (2d_j + 1) \arcsin \sqrt{\tilde{a}} \right] \\ &= -\frac{2d_j + 1}{\sqrt{\tilde{a}(1 - \tilde{a})}} \sqrt{p_{d_j}(\tilde{a}) \left[ 1 - p_{d_j}(\tilde{a}) \right]}, \end{split}$$

giving

$$\begin{split} \frac{\partial}{\partial \tilde{a}} \log L_{d_j}^{(\text{single shot})} &= \frac{2d_j + 1}{\sqrt{\tilde{a}(1 - \tilde{a})}} \sqrt{p_{d_j}(\tilde{a}) \left[1 - p_{d_j}(\tilde{a})\right]} \left[\frac{m_{d_j}}{p_{d_j}(\tilde{a})} - \frac{1 - m_{d_j}}{1 - p_{d_j}(\tilde{a})}\right] \\ &= \frac{\left(2d_j + 1\right) \left(m_{d_j} \left[1 - p_{d_j}(\tilde{a})\right] - \left[1 - m_{d_j}\right] \left[p_{d_j}(\tilde{a})\right]\right)}{\sqrt{\tilde{a}(1 - \tilde{a}) \left[p_{d_j}(\tilde{a})\right] \left[1 - p_{d_j}(\tilde{a})\right]}}. \end{split}$$

As  $m_{d_j}$  is a binary value, squaring this gives

$$\left(\frac{\partial}{\partial \tilde{a}} \log L_{d_j}^{(\text{single shot})}\right)^2 = \frac{\left(2d_j+1\right)^2 \left(m_{d_j} \left[1-p_{d_j}(\tilde{a})\right]^2 + \left[1-m_{d_j}\right] \left[p_{d_j}(\tilde{a})\right]^2\right)}{\tilde{a}(1-\tilde{a}) \left[p_{d_j}(\tilde{a})\right] \left[1-p_{d_j}(\tilde{a})\right]},$$

and, finally, the expectation value over  $m_{d_j}$  can be taken and evaluated at the true amplitude a to find the Fisher information as

$$I^{(\text{single shot})}(a) = \frac{(2d_j+1)^2 \left( \left[ p_{d_j}(a) \right] \left[ 1 - p_{d_j}(a) \right]^2 + \left[ 1 - p_{d_j}(a) \right] \left[ p_{d_j}(a) \right]^2 \right)}{a(1-a) \left[ p_{d_j}(a) \right] \left[ 1 - p_{d_j}(a) \right]}$$
$$= \frac{(2d_j+1)^2}{a(1-a)} \left( \left[ 1 - p_{d_j}(a) \right] + \left[ p_{d_j}(a) \right] \right)$$
$$= \frac{(2d_j+1)^2}{a(1-a)}.$$

As Fisher information is additive, the single shot Fisher information  $I^{\text{single shot}}(a)$  can be summed over the  $N_{\text{shot}}$  shots at each Grover-depth  $d_j$  to give the total fisher information

$$I(a) = \frac{N_{\text{shot}}}{a(1-a)} \sum_{j=0}^{q-1} (2d_j + 1)^2,$$

as required.

#### Appendix B. Analysis of the base for the exponential schedule

In subsection 3.1, we defined a slight modification to the exponential Grover-depth schedule  $D_{\text{EXP}}$  (from Eq. (31)) to accommodate a maximum Grover-depth d that is not a power of 2. The modification, as stated in Eq. (44), is given by

$$D_{\text{EXP},\nu} = \{d_0 = 0\} \cup \{d_j = \text{Round}(\nu^{j-1})\}_{j=1}^{q-1}$$

where the base  $\nu$  and number of depths q are chosen together such that  $\nu$  is the closest value to 2 such that  $\nu^{q-2} = d$ . However, our scaling analyses consider only the case where the maximum Grover-depth d is a power of 2; that is, for  $\nu = 2$ . We assume that the scaling results still apply for other values of the base  $\nu$  and justify this with the claim that "the difference  $|\nu - 2|$  becomes smaller over larger scales of maximum Grover-depth d". In this appendix, we briefly show that this is true.

By definition,

$$\nu_{\text{lower}} < \nu < \nu_{\text{upper}},$$

where

$$\nu_{\text{lower}}^{q-2} = 2^{q-3}$$
 $\nu_{\text{upper}}^{q-2} = 2^{q-1}.$ 

These bounds can be evaluated as

$$\nu_{\text{lower}} = 2^{\frac{q-3}{q-2}}$$
 $\nu_{\text{upper}} = 2^{\frac{q-1}{q-2}},$ 

both of which tend to 2 with increasing q, and thus with increasing Grover-depth d.

## References

- Ashley Montanaro. Quantum algorithms: an overview. npj Quantum Information, 2(1):15023, 2016. doi: 10.1038/npjqi.2015.23. URL https://doi. org/10.1038/npjqi.2015.23.
- [2] Chien-Hung Cho, Chih-Yu Chen, Kuo-Chin Chen, Tsung-Wei Huang, Ming-Chien Hsu, Ning-Ping Cao, Bei Zeng, Seng-Ghee Tan, and Ching-Ray Chang. Quantum computation: Algorithms and Applications. *Chinese Journal of Physics*, 72:248–269, 2021. doi: https://doi.org/10.1016/j.cjph. 2021.05.001. URL https://www.sciencedirect.com/science/article/pii/S0577907321001039.
- Richard P. Feynman. Simulating physics with computers. International Journal of Theoretical Physics, 21(6):467-488, 1982. doi: 10.1007/BF02650179. URL https://doi.org/10.1007/BF02650179.
- [4] Paul Benioff. The computer as a physical system: A microscopic quantum mechanical Hamiltonian model of computers as represented by Turing machines. *Journal of statistical physics*, 22(5):563-591, 1980. URL https://doi.org/10. 1007/BF01011339.
- [5] David Deutsch and Roger Penrose. 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. doi: 10. 1098/rspa.1985.0070. URL https://royalsocietypublishing.org/doi/abs/ 10.1098/rspa.1985.0070.
- [6] Ashley Montanaro. Quantum speedup of Monte Carlo methods. Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences, 471(2181):20150301, 2015. doi: 10.1098/rspa.2015.0301. URL https:// royalsocietypublishing.org/doi/abs/10.1098/rspa.2015.0301.
- [7] Patrick Rebentrost, Brajesh Gupt, and Thomas R. Bromley. Quantum computational finance: Monte Carlo pricing of financial derivatives. *Phys. Rev. A*, 98:022321, 2018. doi: 10.1103/PhysRevA.98.022321. URL https://link.aps.org/doi/10.1103/PhysRevA.98.022321.
- [8] Koichi Miyamoto and Kenji Shiohara. Reduction of qubits in a quantum algorithm for Monte Carlo simulation by a pseudo-random-number generator. *Phys. Rev. A*, 102:022424, 2020. doi: 10.1103/PhysRevA.102.022424. URL https://link.aps.org/doi/10.1103/PhysRevA.102.022424.
- [9] Nathan Wiebe, Ashish Kapoor, and Krysta M. Svore. Quantum algorithms for nearest-neighbor methods for supervised and unsupervised learning. *Quantum Info. Comput.*, 15(3-4):316-356, 2015. URL https://doi.org/10.26421/ QIC15.3-4-7.
- [10] Nathan Wiebe, Ashish Kapoor, and Krysta M. Svore. Quantum deep learning. Quantum Info. Comput., 16(7-8):541-587, 2016. URL https://doi.org/10. 26421/QIC16.7-8-1.
- [11] Nathan Wiebe, Ashish Kapoor, and Krysta M Svore. Quantum Perceptron Models. In Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS'16, page 4006–4014, Red Hook, NY, USA, 2016. Curran Associates Inc. URL https://dl.acm.org/doi/10.5555/ 3157382.3157545.

- [12] Tongyang Li, Shouvanik Chakrabarti, and Xiaodi Wu. Sublinear quantum algorithms for training linear and kernel-based classifiers. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 3815–3824. PMLR, 2019. URL https://proceedings.mlr. press/v97/li19b.html.
- [13] Iordanis Kerenidis, Jonas Landman, Alessandro Luongo, and Anupam Prakash. q-means: A quantum algorithm for unsupervised machine learning. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, Advances in Neural Information Processing Systems, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper/2019/ hash/16026d60ff9b54410b3435b403afd226-Abstract.html.
- [14] Hideyuki Miyahara, Kazuyuki Aihara, and Wolfgang Lechner. Quantum expectation-maximization algorithm. *Phys. Rev. A*, 101:012326, 2020. doi: 10.1103/PhysRevA.101.012326. URL https://link.aps.org/doi/10.1103/ PhysRevA.101.012326.
- [15] Gilles Brassard, Peter Høyer, Michele Mosca, and Alain Tapp. Quantum Amplitude Amplification and Estimation. *Contemp. Math. Ser. Millenn.*, 305: 53-74, 2002. doi: 10.1090/conm/305. URL http://dx.doi.org/10.1090/conm/ 305.
- [16] Lov K Grover. A fast quantum mechanical algorithm for database search. In Proceedings of the twenty-eighth annual ACM symposium on Theory of computing, pages 212-219, 1996. URL https://dl.acm.org/doi/pdf/10.1145/237814.
   237866.
- [17] A Yu Kitaev. Quantum measurements and the Abelian stabilizer problem, 1995. URL https://arxiv.org/abs/quant-ph/9511026. arXiv preprint quantph/9511026.
- [18] Scott Aaronson and Patrick Rall. Quantum Approximate Counting, Simplified, pages 24-32. SIAM, 2020. doi: 10.1137/1.9781611976014.5. URL https: //doi.org/10.1137/1.9781611976014.5.
- [19] Ramgopal Venkateswaran and Ryan O'Donnell. Quantum approximate counting with nonadaptive Grover iterations. In 38th International Symposium on Theoretical Aspects of Computer Science, pages 59:1-59:12, 2021. URL https: //arxiv.org/pdf/2010.04370.pdf.
- [20] Kouhei Nakaji. Faster amplitude estimation. Quantum Info. Comput., 20(13–14): 1109–1123, 2020. URL https://doi.org/10.26421/QIC20.13-14-2.
- [21] Dmitry Grinko, Julien Gacon, Christa Zoufal, and Stefan Woerner. Iterative quantum amplitude estimation. npj Quantum Information, 7(1):52, 2021. doi: 10.1038/s41534-021-00379-1. URL https://doi.org/10.1038/ s41534-021-00379-1.
- [22] Patrick Rall. Faster Coherent Quantum Algorithms for Phase, Energy, and Amplitude Estimation. Quantum, 5:566, 2021. doi: 10.22331/q-2021-10-19-566. URL https://doi.org/10.22331/q-2021-10-19-566.
- [23] Patrick Rall and Bryce Fuller. Amplitude Estimation from Quantum Signal Processing, 2022. URL https://arxiv.org/abs/2207.08628. arXiv preprint arXiv:2207.08628.

- [24] Alberto Manzano, Daniele Musso, and Álvaro Leitao. Real quantum amplitude estimation, 2022. URL https://arxiv.org/abs/2204.13641. arXiv preprint arXiv:2204.13641.
- [25] Tudor Giurgica-Tiron, Iordanis Kerenidis, Farrokh Labib, Anupam Prakash, and William Zeng. Low depth algorithms for quantum amplitude estimation. *Quantum*, 6:745, 2022. doi: 10.22331/q-2022-06-27-745. URL https://doi. org/10.22331/q-2022-06-27-745.
- [26] Shion Fukuzawa, Christopher Ho, Sandy Irani, and Jasen Zion. Modified Iterative Quantum Amplitude Estimation is Asymptotically Optimal, 2022. URL https://arxiv.org/abs/2208.14612. arXiv preprint arXiv:2208.14612.
- [27] Yohichi Suzuki, Shumpei Uno, Rudy Raymond, Tomoki Tanaka, Tamiya Onodera, and Naoki Yamamoto. Amplitude estimation without phase estimation. *Quantum Information Processing*, 19(2):75, 2020. doi: 10.1007/ s11128-019-2565-2. URL https://doi.org/10.1007/s11128-019-2565-2.
- [28] Eric G Brown, Oktay Goktas, and WK Tham. Quantum amplitude estimation in the presence of noise, 2020. URL https://arxiv.org/abs/2006.14145. arXiv preprint arXiv:2006.14145.
- [29] Tomoki Tanaka, Yohichi Suzuki, Shumpei Uno, Rudy Raymond, Tamiya Onodera, and Naoki Yamamoto. Amplitude estimation via maximum likelihood on noisy quantum computer. *Quantum Information Processing*, 20(9):293, 2021. doi: 10.1007/s11128-021-03215-9. URL https://doi.org/10.1007/ s11128-021-03215-9.
- [30] Steven Herbert, Roland Guichard, and Darren Ng. Noise-Aware Quantum Amplitude Estimation, 2021. URL https://arxiv.org/abs/2109.04840. arXiv preprint arXiv:2109.04840.
- [31] Shumpei Uno, Yohichi Suzuki, Keigo Hisanaga, Rudy Raymond, Tomoki Tanaka, Tamiya Onodera, and Naoki Yamamoto. Modified Grover operator for quantum amplitude estimation. *New Journal of Physics*, 23(8):083031, 2021. doi: 10.1088/ 1367-2630/ac19da. URL https://doi.org/10.1088/1367-2630/ac19da.
- [32] Tomoki Tanaka, Shumpei Uno, Tamiya Onodera, Naoki Yamamoto, and Yohichi Suzuki. Noisy quantum amplitude estimation without noise estimation. *Phys. Rev. A*, 105:012411, 2022. doi: 10.1103/PhysRevA.105.012411. URL https: //link.aps.org/doi/10.1103/PhysRevA.105.012411.
- [33] Kirill Plekhanov, Matthias Rosenkranz, Mattia Fiorentini, and Michael Lubasch. Variational quantum amplitude estimation. *Quantum*, 6:670, 2022. doi: 10.22331/ q-2022-03-17-670. URL https://doi.org/10.22331/q-2022-03-17-670.
- [34] Adam Callison and Nicholas Chancellor. Hybrid quantum-classical algorithms in the noisy intermediate-scale quantum era and beyond. *Phys. Rev. A*, 106:010101, 2022. doi: 10.1103/PhysRevA.106.010101. URL https://link.aps.org/doi/ 10.1103/PhysRevA.106.010101.
- [35] A. W. van der Vaart. Bernstein-von Mises Theorem. In Asymptotic Statistics, chapter 10.2, pages 140–146. Cambride University Press, Cambridge, UK, 1998.
- [36] B L Higgins, D W Berry, S D Bartlett, M W Mitchell, H M Wiseman, and G J Pryde. Demonstrating heisenberg-limited unambiguous phase estimation without adaptive measurements. New Journal of Physics, 11(7):073023, 2009. doi: 10.1088/1367-2630/11/7/073023. URL https://dx.doi.org/10.1088/ 1367-2630/11/7/073023.

- [37] Shelby Kimmel, Guang Hao Low, and Theodore J. Yoder. Robust calibration of a universal single-qubit gate set via robust phase estimation. *Phys. Rev.* A, 92:062315, Dec 2015. doi: 10.1103/PhysRevA.92.062315. URL https://link.aps.org/doi/10.1103/PhysRevA.92.062315.
- [38] Nathan Wiebe and Chris Granade. Efficient Bayesian Phase Estimation. Phys. Rev. Lett., 117:010503, 2016. doi: 10.1103/PhysRevLett.117.010503. URL https://link.aps.org/doi/10.1103/PhysRevLett.117.010503.
- [39] Emanuel Knill, Gerardo Ortiz, and Rolando D. Somma. Optimal quantum measurements of expectation values of observables. *Phys. Rev. A*, 75:012328, 2007. doi: 10.1103/PhysRevA.75.012328. URL https://link.aps.org/doi/ 10.1103/PhysRevA.75.012328.
- [40] Alicja Dutkiewicz, Barbara M Terhal, and Thomas E O'Brien. Heisenberg-limited quantum phase estimation of multiple eigenvalues with few control qubits, 2021. URL https://arxiv.org/abs/2107.04605. arXiv preprint arXiv:2107.04605.
- [41] Ilia Zintchenko and Nathan Wiebe. Randomized gap and amplitude estimation. Phys. Rev. A, 93:062306, 2016. doi: 10.1103/PhysRevA.93.062306. URL https://link.aps.org/doi/10.1103/PhysRevA.93.062306.
- [42] Salvador Elías Venegas-Andraca. Quantum walks: a comprehensive review. Quantum Information Processing, 11(5):1015–1106, 2012. doi: 10.1007/ s11128-012-0432-5. URL https://doi.org/10.1007/s11128-012-0432-5.
- [43] Daniel J. Egger, Claudio Gambella, Jakub Marecek, Scott McFaddin, Martin Mevissen, Rudy Raymond, Andrea Simonetto, Stefan Woerner, and Elena Yndurain. Quantum Computing for Finance: State-of-the-Art and Future Prospects. *IEEE Transactions on Quantum Engineering*, 1:1–24, 2020. doi: 10. 1109/TQE.2020.3030314. URL https://doi.org/10.1109/TQE.2020.3030314.
- [44] Daniel J. Egger, Ricardo García Gutiérrez, Jordi Cahué Mestre, and Stefan Woerner. Credit risk analysis using quantum computers. *IEEE Transactions* on Computers, 70(12):2136–2145, 2021. doi: 10.1109/TC.2020.3038063. URL https://doi.org/10.1109/TQE.2020.3030314.
- [45] Tudor Giurgica-Tiron, Sonika Johri, Iordanis Kerenidis, Jason Nguyen, Neal Pisenti, Anupam Prakash, Ksenia Sosnova, Ken Wright, and William Zeng. Low-depth amplitude estimation on a trapped-ion quantum computer. *Phys. Rev. Research*, 4:033034, 2022. doi: 10.1103/PhysRevResearch.4.033034. URL https://link.aps.org/doi/10.1103/PhysRevResearch.4.033034.
- [46] Chu-Ryang Wie. Simpler quantum counting. Quantum Info. Comput., 19(11–12): 0967–0983, 2019. URL https://doi.org/10.26421/QIC19.11-12-5.
- [47] Guoming Wang, Dax Enshan Koh, Peter D. Johnson, and Yudong Cao. Minimizing estimation runtime on noisy quantum computers. *PRX Quantum*, 2:010346, 2021. doi: 10.1103/PRXQuantum.2.010346. URL https://link.aps. org/doi/10.1103/PRXQuantum.2.010346.
- [48] Pooja Rao, Kwangmin Yu, Hyunkyung Lim, Dasol Jin, and Deokkyu Choi. Quantum amplitude estimation algorithms on IBM quantum devices. In Keith S. Deacon, editor, *Quantum Communications and Quantum Imaging XVIII*, volume 11507, pages 49 – 60. International Society for Optics and Photonics, SPIE, 2020. doi: 10.1117/12.2568748. URL https://doi.org/10.1117/12.2568748.

- [49] JM Blair, CA Edwards, and J Howard Johnson. Rational Chebyshev approximations for the inverse of the error function. *Mathematics of Computation*, 30(136):827–830, 1976. doi: 10.1090/S0025-5718-1976-0421040-7. URL http: //dx.doi.org/10.1090/S0025-5718-1976-0421040-7.
- [50] Ryan Babbush, Jarrod R. McClean, Michael Newman, Craig Gidney, Sergio Boixo, and Hartmut Neven. Focus beyond Quadratic Speedups for Error-Corrected Quantum Advantage. *PRX Quantum*, 2:010103, 2021. doi: 10.1103/PRXQuantum.2.010103. URL https://link.aps.org/doi/10.1103/ PRXQuantum.2.010103.