

Variational quantum amplitude estimation

Kirill Plekhanov, Matthias Rosenkranz, Mattia Fiorentini, and Michael Lubasch

Cambridge Quantum Computing Limited, SW1P 1BX London, United Kingdom

We propose to perform amplitude estimation with the help of constant-depth quantum circuits that variationally approximate states during amplitude amplification. In the context of Monte Carlo (MC) integration, we numerically show that shallow circuits can accurately approximate many amplitude amplification steps. We combine the variational approach with maximum likelihood amplitude estimation [Y. Suzuki *et al.*, *Quantum Inf. Process.* 19, 75 (2020)] in variational quantum amplitude estimation (VQAE). VQAE typically has larger computational requirements than classical MC sampling. To reduce the variational cost, we propose adaptive VQAE and numerically show in 6 to 12 qubit simulations that it can outperform classical MC sampling.

1 Introduction

Amplitude estimation [1] is a powerful algorithm that can achieve a quadratic quantum speedup over classical Monte Carlo (MC) methods [2]. It has a wide range of applications, e.g. in quantum chemistry [3, 4], machine learning [5–7], and finance [8–10] where it can help with tasks such as risk analysis [11, 12] and the pricing of financial derivatives [13, 14].

The original amplitude estimation procedure [1] has hardware requirements that are challenging for current quantum devices and, therefore, reducing these requirements is currently an active area of research. Crucial breakthroughs were obtained in recent proposals which succeeded in replacing the hardware-intensive components of traditional amplitude estimation – controlled multi-qubit gates and quantum Fourier transform – by classical post-processing [15–18]. Alternatively, one can systematically reduce the circuit depth by interpolating between classical MC methods and amplitude estimation [19]. Additionally, classical pre-processing can replace costly quantum arithmetic [20] and Bayesian inference can be used to boost the algorithmic efficiency in the presence of device errors [21, 22].

In this article, we address the question whether the quantum computational requirements for amplitude estimation can be further decreased by making use of

Kirill Plekhanov: kirill.plekhanov@cambridgequantum.com

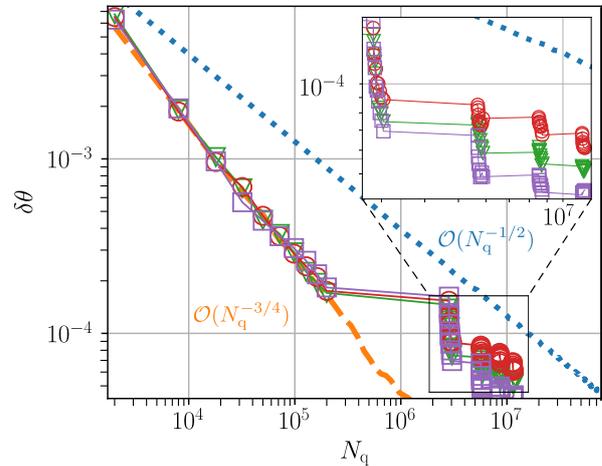


Figure 1: Amplitude estimation error $\delta\theta$ as a function of the computational cost, i.e. the number of queries N_q . We compare adaptive VQAE (symbols with lines guide to the eye) to MLAE (dashed orange) and classical MC sampling (dotted blue). We calculate the rescaled mean value of a shifted Cauchy-Lorentz (red circles), Gaussian (green triangles), and log-normal (purple squares) probability distribution. In VQAE, the first 10 amplitude estimates are computed via MLAE, then one step of the variational optimization is performed, which is followed by the next iteration of 10 MLAE steps. This procedure of variational approximation followed by MLAE is repeated three more times, resulting in a final error $\delta\theta \approx 6 \cdot 10^{-5}$. We see that, for this error, N_q is up to an order of magnitude smaller in VQAE than in classical MC sampling. Throughout this calculation, VQAE’s circuit depth is the depth of the initial state plus at most 10 times the depth of the query operator, whereas MLAE’s circuits have the depth of the initial state plus, at the end, 50 times the depth of the query operator.

variational quantum algorithms [23–25]. We present variational quantum amplitude estimation (VQAE) in which the depth of the entire quantum circuit is always kept below a desired maximum value by means of variational optimization. VQAE is based on maximum likelihood amplitude estimation (MLAE) [15]. We present a naïve and an adaptive VQAE algorithm. Adaptive VQAE rescales the amplitude to reduce the cost of the variational optimization. The advantage of VQAE over MLAE is that the maximum circuit depth of VQAE is independent of the total number of MLAE steps, whereas in MLAE this depth grows linearly with the number of MLAE steps. The advantage of VQAE over classical MC sampling is that VQAE can have a lower computational cost. Figure 1

shows that, for the problems considered here, VQAE outperforms classical MC sampling and additionally keeps the overall circuit depth below a fixed value.

This article is organized as follows. Firstly, in Section 2, we define the problems considered here and explain the original quantum algorithm for amplitude estimation as well as the classical MC approach. Then, in Section 3 we present our variational methods, study variational errors of constant-depth quantum circuits, and develop naïve and adaptive VQAE. We conclude this article and discuss potential next steps in Section 4.

2 Background

In this section, we first define the problem that we are interested in. Next, we explain quantum amplitude estimation and classical MC sampling.

2.1 Problem definition

Throughout this article, we focus on the calculation of expectation values

$$\mathbb{E}_p[f] = \sum_x p(x)f(x) \quad (1)$$

where the sum runs over 2^n equidistant values of $x \in [0, 1)$, $p(x)$ represents a probability distribution and $f(x)$ a real-valued function. Here n is the qubit count of the wave function that encodes $p(x)$ and $f(x)$ in its amplitudes. We consider three probability distributions: a Gaussian

$$p_G(x) = \frac{1}{\mathcal{N}_G} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right), \quad (2)$$

Cauchy-Lorentz

$$p_{C-L}(x) = \frac{1}{\mathcal{N}_{C-L}} \frac{\sigma}{(x-\mu)^2 + \sigma^2}, \quad (3)$$

and log-normal distribution

$$p_{l-n}(x) = \frac{1}{\mathcal{N}_{l-n}(c_0 + c_1x)} \exp\left(-\frac{(\ln(c_0 + c_1x) - \mu)^2}{2\sigma^2}\right). \quad (4)$$

The normalization constants \mathcal{N}_G , \mathcal{N}_{C-L} and \mathcal{N}_{l-n} are chosen so that $\sum_x p(x) = 1$.

We choose the following parameters for our analysis. We fix the total number of qubits encoding $f(x)$ and $p(x)$ to $n = 5$. In our calculations with the Gaussian and Cauchy-Lorentz distribution, we use $\mu = 0.5$ and $\sigma = 0.1$. In our calculations with the log-normal distribution, we use $c_0 = 0$, $c_1 = 10$, $\mu = 1.5$, and $\sigma = 0.2$. For the function $f(x)$, we use

$$f(x) = Cx \quad (5)$$

with some $C > 0$. For this choice of parameters, the expectation value (1) is approximately $\mathbb{E}_p[f] \approx 0.5 C$ for all distributions.

2.2 Quantum amplitude estimation

Let us present a way to encode the solution to (1) on a quantum computer. We assume $f(x)$ and $p(x)$ are functions that map $[0, 1)$ to $[0, 1]$. We consider real numbers $x \in [0, 1)$ that satisfy

$$x = \sum_i \frac{x_i}{2^{n-i}}, \quad x_i \in \{0, 1\} \quad (6)$$

and that we identify with n -bit strings $\{x_i, i = 1, 2, \dots, n\}$. Each bit string shall correspond to a quantum state $|x\rangle = |x_1, x_2, \dots, x_n\rangle$ in the computational basis of a n -qubit register. Additionally we have a quantum circuit \mathcal{A} that acts on a register of $n+1$ qubits and produces a state $|\chi_0\rangle_{n+1} = \mathcal{A}|0\rangle_{n+1}$ such that

$$|\chi_0\rangle_{n+1} = \sqrt{1-a} |\psi_{\text{bad}}\rangle_n |0\rangle + \sqrt{a} |\psi_{\text{good}}\rangle_n |1\rangle. \quad (7)$$

Here $|\psi_{\text{bad}}\rangle_n$ and $|\psi_{\text{good}}\rangle_n$ are two normalized quantum states of a n -qubit register which is connected to one additional ancilla qubit. We define the good state

$$|\psi_{\text{good}}\rangle_n = \frac{1}{\sqrt{a}} \sum_x \sqrt{p(x)f(x)} |x\rangle_n \quad (8)$$

so that $a = \mathbb{E}_p[f]$ of Eq. (1) coincides with the probability of measuring the ancilla qubit in the state $|1\rangle$.

To determine a , the amplitude estimation algorithm uses the Grover operator $\mathcal{Q} = -\mathcal{R}_\chi \mathcal{R}_{\text{good}}$ [1, 26, 27] where

$$\begin{aligned} \mathcal{R}_\chi &= \mathbb{I} - 2|\chi_0\rangle_{n+1} \langle \chi_0|_{n+1} = \mathbb{I} - 2\mathcal{A}|0\rangle_{n+1} \langle 0|_{n+1} \mathcal{A}^\dagger \\ \mathcal{R}_{\text{good}} &= \mathbb{I} - 2|\psi_{\text{good}}\rangle_n |1\rangle \langle \psi_{\text{good}}|_n \langle 1| \end{aligned} \quad (9)$$

are reflections in a two-dimensional subspace \mathcal{H}_χ spanned by states $|\psi_{\text{bad}}\rangle_n |0\rangle$ and $|\psi_{\text{good}}\rangle_n |1\rangle$. We define $a = \sin^2(\theta)$ and explicitly write out the action of \mathcal{Q} :

$$\begin{aligned} |\chi_m\rangle_{n+1} &= Q^m |\chi_0\rangle_{n+1} = \cos[(2m+1)\theta] |\psi_{\text{bad}}\rangle_n |0\rangle \\ &\quad + \sin[(2m+1)\theta] |\psi_{\text{good}}\rangle_n |1\rangle. \end{aligned} \quad (10)$$

Therefore, the subspace \mathcal{H}_χ is stable under the action of \mathcal{Q} and the only effect of \mathcal{Q} is to rotate by an angle of 2θ . The original amplitude estimation algorithm then uses quantum phase estimation to find the eigenvalues of \mathcal{Q} equal to $\exp(\pm 2i\theta)$ and provides an estimate of a with an error

$$\epsilon \leq 2\pi \frac{\sqrt{a(1-a)}}{2N_q} + \frac{\pi^2}{4N_q^2}, \quad (11)$$

with a probability of at least $8/\pi^2$ [1] where N_q ($2N_q$) is the total number of times the operator \mathcal{A} (\mathcal{Q}) has to be applied.

Traditional amplitude estimation has high requirements on quantum hardware because it uses quantum

phase estimation. This algorithm needs the quantum Fourier transform and multiple controlled Q^m operations where $\{m = 1, 2, 4, \dots, 2^M\}$. The depth of the corresponding quantum circuit is mostly determined by the depth of the last controlled Q^m operator for which $m = 2^M$. In general, the total circuit depth scales like the total number of queries $\mathcal{O}(N_q) \sim \mathcal{O}(1/\epsilon)$ inversely proportional to the desired error ϵ .

To avoid these deep quantum circuits, several recent articles propose new ways to carry out amplitude estimation circumventing quantum phase estimation [15, 17, 18] and circuits of depth $\mathcal{O}(1/\epsilon)$ [19]. One proposal is MLAE [15] in which one combines measurements of the states $|\chi_m\rangle_{n+1}$ with a maximum likelihood estimation of a . For an exponential schedule $\{m = 1, 2, 4, \dots, 2^M\}$, this algorithm has the query cost $N_q \sim \mathcal{O}(1/\epsilon)$. A linear schedule $\{m = 1, 2, 3, \dots, M\}$ increases the query cost to $N_q \sim \mathcal{O}(\epsilon^{-4/3})$. Note that in this case N_q scales quadratically with the maximum circuit depth M . Following the same idea of reducing the hardware requirements, the authors of Ref. [19] present two algorithms with computational cost $N_q \sim \mathcal{O}(1/\epsilon^{\beta+1})$ for quantum circuits of reduced depth $\mathcal{O}(1/\epsilon^{1-\beta})$. These algorithms are controlled by an external parameter β which allows one to interpolate between the quantum regime at $\beta = 0$ and the classical MC regime at $\beta = 1$.

2.3 Classical MC sampling

We perform classical MC sampling in the following way. We sample from the state $|\chi_0\rangle_{n+1}$ of Eq. (7) and measure the ancilla qubit. We compute a as the relative frequency of measuring the ancilla qubit in the state $|1\rangle$. This calculation of a has the error $\epsilon = \sqrt{a(1-a)/N_q}$ so that the total number of queries required for a certain error ϵ is $N_q \sim \mathcal{O}(1/\epsilon^2)$ [28].

Comparing this query cost with the previous ones, we find that traditional amplitude estimation as well as MLAE with exponential schedule achieve a quadratic quantum speedup over classical MC sampling. Both MLAE with linear schedule and the algorithms in [19] obtain a reduced quantum speedup.

Note that, throughout this article, the query complexity is defined in terms of \mathcal{A} operators, with two applications of \mathcal{A} required per application of Q , see Eq. (9). Also, the depth of quantum circuits is measured in units of \mathcal{A} , so that the depth of $|\chi_0\rangle_{n+1}$ is equal to one and the depth of Q is equal to two. Additionally, in the following we assume that \mathcal{A} and Q are given, i.e. we do not address questions e.g. relating to their efficient quantum circuit representation.

3 Variational algorithms

Here we present our variational algorithms. We first explain the general VQAE formalism, then our naïve

implementation, and finally the adaptive VQAE approach.

3.1 General formalism

The VQAE algorithm is based on the maximum likelihood framework of Ref. [15] with linearly incremental sequence $\{m = 1, 2, 3, \dots, M\}$. In this framework, the depth of the quantum circuit implementing the state $|\chi_m\rangle_{n+1} = Q^m |\chi_0\rangle_{n+1}$ scales with m as $2m+1$. To prevent the circuit depth from increasing indefinitely, we add to this framework a variational step during which states $|\chi_m\rangle_{n+1}$ are periodically approximated by a variational quantum state of depth one. We note that this strategy will not always work and the corresponding approximation can have a large error. The variational approach, however, allows us to compute the approximation error so that we can identify when the strategy works. We perform the variational approximation every k -th power of Q , with $0 < k < M$. For all the other iterations, we simply apply the corresponding power of Q to the variational state. This results in Algorithm 1.

Algorithm 1 Variational quantum amplitude estimation

Require: functions f and p , integer k

- Use f and p to encode $|\phi_{i=0}\rangle_{n+1} = |\chi_0\rangle_{n+1}$ and $Q = -\mathcal{R}_\chi \mathcal{R}_{\text{good}}$ according to Eqs. (7) and (9)
- for** $0 \leq m \leq M$ **do**
 - Set $i = \lfloor m/k \rfloor$, $j = m \% k$
 - ▷ sampling
 - Sample the circuit $Q^j |\phi_i\rangle_{n+1}$ and collect h samples.
 - Save the number of times the ancilla qubit is $|1\rangle$ in a variable h_m
 - ▷ end sampling
 - ▷ variational approximation
 - if** $j = k - 1$ **then**
 - Perform the variational approximation
 - $|\phi_{i+1}\rangle_{n+1} \approx Q^k |\phi_i\rangle_{n+1}$
 - end if**
 - ▷ end variational approximation
- end for**
 - Use $\{h_m\}$ to carry out the maximum likelihood estimation

Here $\lfloor \cdot \rfloor$ denotes the floor function and $\%$ the modulo operation. The resulting approximation of $|\chi_m\rangle_{n+1}$ corresponds to the state $Q^j |\phi_i\rangle_{n+1}$ with $i = \lfloor m/k \rfloor$ and $j = m \% k$. The depth of this approximation reaches the minimum of one when $j = 0$ and the maximum of $2k - 1$ when $j = k - 1$.

The maximum likelihood post-processing [15] consists in maximizing the likelihood function

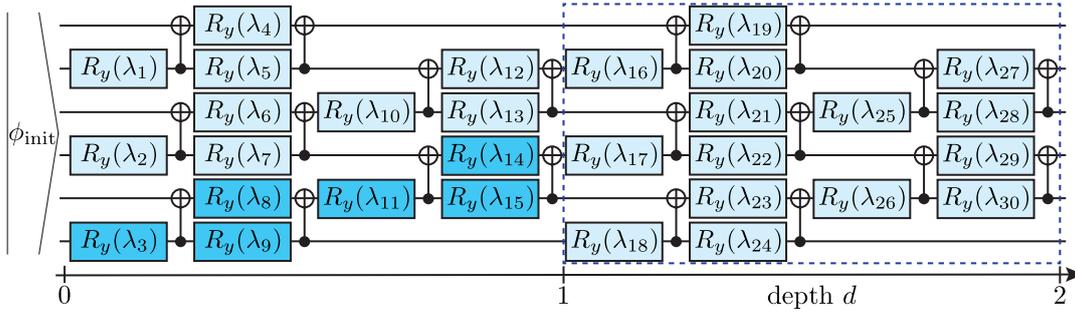


Figure 2: Variational PQC ansatz of depth d for 6 qubits. The variational parameters λ_j are in single-qubit rotation gates $R_y(\lambda_j) = \exp(-i\lambda_j\sigma_y/2)$ where σ_y denotes the y Pauli matrix. In naïve VQAE, $|\phi_{\text{init}}\rangle_{n+1} = |0\rangle_{n+1}$ and the circuit structure inside the dashed box is repeated d times. In adaptive VQAE, $|\phi_{\text{init}}\rangle_{n+1} = |\chi_0\rangle_{n+1}$ and only the dark blue gates with adjacent CNOTs form our variational ansatz.

$L(\{h_m\}, x) = \prod_m L_m(h_m, x)$ with

$$L_m(h_m, x) = [\sin^2((2m+1)x)]^{h_m} [\cos^2((2m+1)x)]^{h-h_m}, \quad (12)$$

so that the estimate of the phase θ becomes

$$\hat{\theta} = \arg \max_x \left(\ln L(\{h_m\}, x) \right). \quad (13)$$

Our implementations of the maximum likelihood estimation use $h = 2 \times 10^3$ samples. The minimization of $L(\{h_m\}, x)$ is accomplished by means of a brute-force search algorithm that uses 5×10^3 grid points.

We variationally approximate states $\mathcal{Q}^k |\phi_i\rangle_{n+1}$ by minimizing $\| |\phi_{\text{var}}(\boldsymbol{\lambda})\rangle_{n+1} - \mathcal{Q}^k |\phi_i\rangle_{n+1} \|^2$ which is equivalent to maximizing the objective function

$$\mathcal{F}(\boldsymbol{\lambda}) = \Re \left\langle \phi_{\text{var}}(\boldsymbol{\lambda}) |_{n+1} \mathcal{Q}^k |\phi_i\rangle_{n+1} \right\rangle \quad (14)$$

with respect to the variational parameters $\boldsymbol{\lambda}$. The optimal solution can be formally written as

$$|\phi_{i+1}\rangle_{n+1} = |\phi_{\text{var}}(\tilde{\boldsymbol{\lambda}})\rangle, \quad \tilde{\boldsymbol{\lambda}} = \arg \max_{\boldsymbol{\lambda}} \mathcal{F}(\boldsymbol{\lambda}). \quad (15)$$

We notice that the depth of the quantum circuit required to compute $\mathcal{F}(\boldsymbol{\lambda})$ is the largest circuit depth used by the algorithm. This quantum circuit is composed of the parts encoding $|\phi_{\text{var}}(\boldsymbol{\lambda})\rangle_{n+1}$ and $|\phi_i\rangle_{n+1}$, each having depth one, and an operator \mathcal{Q}^k of depth $2k$, resulting in a total depth of $2k+2$. In general, the variational quantum state $|\phi_{\text{var}}(\boldsymbol{\lambda})\rangle_{n+1}$ is a parameterized quantum circuit (PQC)

$$|\phi_{\text{var}}(\boldsymbol{\lambda})\rangle_{n+1} = \mathcal{U}_{\text{var}}(\boldsymbol{\lambda}) |\phi_{\text{init}}\rangle_{n+1} = \prod_j e^{-i\lambda_j G_j} |\phi_{\text{init}}\rangle_{n+1}, \quad (16)$$

where G_j are Hermitian operators acting on the $(n+1)$ -qubit register and $|\phi_{\text{init}}\rangle_{n+1}$ is some initial state. For our purposes, we are interested in hardware-efficient quantum circuits that produce real-valued quantum states. We use the PQC shown in Fig. 2 that is composed of d layers with 15 parameterized

single-qubit rotation gates and 10 CNOT gates per layer.

One single variational update of a PQC consists of n_s sweeps over all circuit parameters, during which all parameters are updated simultaneously. To perform the optimization, it is convenient to introduce a coordinate-wise version of Eq. (14) for the j -th parameter

$$f_j(x) = \mathcal{F}(\lambda_1, \lambda_2, \dots, \lambda_{j-1}, x, \lambda_{j+1}, \dots). \quad (17)$$

The optimization of the parameterized state in Eq. (16) can then be performed via a particle swarm approach [29, 30], the coordinate-wise update [31–35], or gradient based methods with the parameter-shift rule [36–41]

$$\frac{df_j(\lambda_j)}{d\lambda_j} = f_j(\lambda_j + \pi/4) - f_j(\lambda_j - \pi/4). \quad (18)$$

We obtained the best results using the gradient based approach with the Adam optimizer [42]. Therefore this technique is being used throughout this article for the computation of all results. Each gradient calculation requires two evaluations of the coordinate-wise objective function $f_j(\lambda_j \pm \pi/4)$. On a quantum computer, f_j can be determined via the Hadamard test [43]. In our numerical simulations, we emulate the measurement of the Hadamard circuit by first evaluating the exact value of f_j and then sampling it using a binomial distribution with the probability $(1 + f_j)/2$ and n_f independent Bernoulli trials [28].

The variational approximation step significantly affects the total number of queries N_q used by VQAE. In MLAE with a linearly incremental sequence, the total number of queries is equal to

$$N_q = \sum_{m=1}^M h(2m+1) = hM(M+2), \quad (19)$$

where $2m+1$ is the depth of the quantum circuit encoding $|\chi_m\rangle_{n+1} = \mathcal{Q}^m |\chi_0\rangle_{n+1}$. In VQAE, the total number of queries is composed of two separate contributions. The first one accounts for the sampling of

the quantum circuits $\mathcal{Q}^j |\phi_i\rangle_{n+1}$ and we denote it by N_{samp} . The corresponding section in Algorithm 1 is labelled by “sampling”. The second contribution corresponds to the variational approximation cost, which we denote by N_{var} . It is associated with the section in Algorithm 1 labelled by “variational approximation”. We assume that the number of queries required per variational approximation is independent of the iteration number m and changes only as a function of the desired variational error as well as the depth of the circuit for the objective function. We denote the cost of a single variational update as $N_{\text{var}/1}(2k+2)$, where $(2k+2)$ is the depth of the objective function $\mathcal{F}(\boldsymbol{\lambda})$ and $N_{\text{var}/1}$ is the number of quantum circuits per variational update that need to be run by the algorithm. As a result, the total number of variational queries becomes

$$N_{\text{var}} = N_{\text{var}/1}(2k+2)\lfloor M/k \rfloor \sim \mathcal{O}(k\lfloor M/k \rfloor), \quad (20)$$

where $\lfloor M/k \rfloor$ is the total number of variational updates required to approximate \mathcal{Q}^M . The number of sampling queries is equal to

$$\begin{aligned} N_{\text{samp}} &= \lfloor M/k \rfloor \sum_{j=0}^{k-1} h(2j+1) + \sum_{j=0}^{M\%k} h(2j+1) \\ &= hk(k+2)\lfloor M/k \rfloor + h(M\%k)(M\%k+2), \end{aligned} \quad (21)$$

where the last term accounts for the situation when k is not a divisor of M . In the limit $M \gg k$, $N_{\text{var}} \sim \mathcal{O}(M)$ and $N_{\text{samp}} \sim \mathcal{O}(kM)$. Note that both contributions scale like $\mathcal{O}(M)$ which is quadratically better than the scaling $\mathcal{O}(M^2)$ of MLAE in Eq. (19).

3.2 Naïve VQAE

In our naïve implementation of the VQAE algorithm, the initial state of the PQC in Eq. (16) and Fig. 2 is $|\phi_{\text{init}}\rangle_{n+1} = |0\rangle_{n+1}$.

Let us first explore the expressive power of the corresponding variational state $|\phi_{\text{var}}(\boldsymbol{\lambda})\rangle$. To this end, we perform amplitude amplifications followed by variational approximations of the resulting state with $k=1$ and $M=50$. To evaluate the quality of the variational approximation, we calculate the infidelity

$$\mathcal{I}_m = 1 - \langle \chi_m |_{n+1} \mathcal{Q}^j |\phi_i\rangle_{n+1}, \quad m = i \cdot k + j, \quad (22)$$

where for $k=1$ we have $j=0$ and $i=m$. Figure 3(a) shows the results of such calculations performed for different depths d of the PQC for $m=10$ and Fig. 3(b) shows the infidelity as a function of m for $d=4$. We observe that the accuracy of the variational ansatz increases with the depth and saturates at $d \approx 4$. The infidelity increases linearly with m . This behaviour is seen for all probability distributions considered.

Next, we present the amplitude estimation results of naïve VQAE. Figure 4 shows the convergence of

$\delta\theta$ as a function of N_{q} , under the assumption that $N_{\text{var}/1} = 0$ and $k=1$. The resulting error is compared with the one of classical MC sampling which scales like $\delta\theta \sim \mathcal{O}(N_{\text{q}}^{-1/2})$ and the one of MLAE which scales like $\delta\theta \sim \mathcal{O}(N_{\text{q}}^{-3/4})$. Interestingly, we find that the convergence of $\delta\theta$ changes as a function of M . For small values of M , it follows the ideal VQAE scaling $\delta\theta \sim \mathcal{O}(N_{\text{q}}^{-3/2})$ as if the variational approximation is performed without error. We emphasize that this scaling is cubically better than the one of classical MC sampling. The second convergence regime is observed for larger values of M . In this regime, the error follows the MC scaling with $\delta\theta \sim \mathcal{O}(N_{\text{q}}^{-1/2})$. To understand this behaviour, we first notice that the MLAE error decreases with M , while the variational error increases instead. In the regime when the MLAE error is larger than the variational error, the scaling of $\delta\theta$ is the best achievable MLAE scaling $\delta\theta \sim \mathcal{O}(N_{\text{q}}^{-3/2})$. When the MLAE error is smaller than the variational error, the convergence of $\delta\theta$ is dominated by the latter. The accumulation of the variational error can be modelled via a random process, in which each variational approximation results in a random error of zero mean and some variance σ^2 . After M steps of the algorithm, $\lfloor M/k \rfloor = M$ (as $k=1$ here) variational approximations were performed resulting in a final error of variance $M\sigma^2$. Hence, an ideal MLAE estimation of the angle θ will produce a relative error $\sqrt{M}\sigma/[(2M+1)\theta]$ scaling as $\mathcal{O}(M^{-1/2})$. In our simulations, we find that the transition from the first regime – where $\delta\theta \sim \mathcal{O}(N_{\text{q}}^{-3/2})$ – to the second regime – where $\delta\theta \sim \mathcal{O}(M^{-1/2})$ – occurs at $M \sim 20$.

Finally, we take into account the cost of the variational approximation, to obtain a more complete assessment of the algorithmic performance of naïve VQAE. To estimate the cost of a single variational update, we write down the number of circuits needed to be run for each variational update as $N_{\text{var}/1} = 2n_f n_s n_p$ where n_p is the number of parameters of a PQC, n_s is the total number of sweeps through all the parameters of the PQC, and n_f is the number of Bernoulli trials per evaluation of the objective function. The factor 2 comes from the fact that two evaluations of the objective function are required for each evaluation of the gradient in Eq. (18). For the PQC in Fig. 2 with $d=4$, the number of parameters is $n_p = 60$. Additionally, we choose $n_f \sim n_s \sim 100$ so that $N_{\text{var}/1} \sim 1.2 \times 10^6$ and $N_{\text{var}} \sim 4.8 \times 10^6 M$. This large variational cost is the dominant part in the calculation of the total number of queries N_{q} . Ultimately, it leads to a performance of naïve VQAE that is worse than the one of classical MC sampling. Reducing any of n_f , n_s , or n_p decreases the variational cost but also increases the variational error which then leads to a worse final amplitude estimation error.

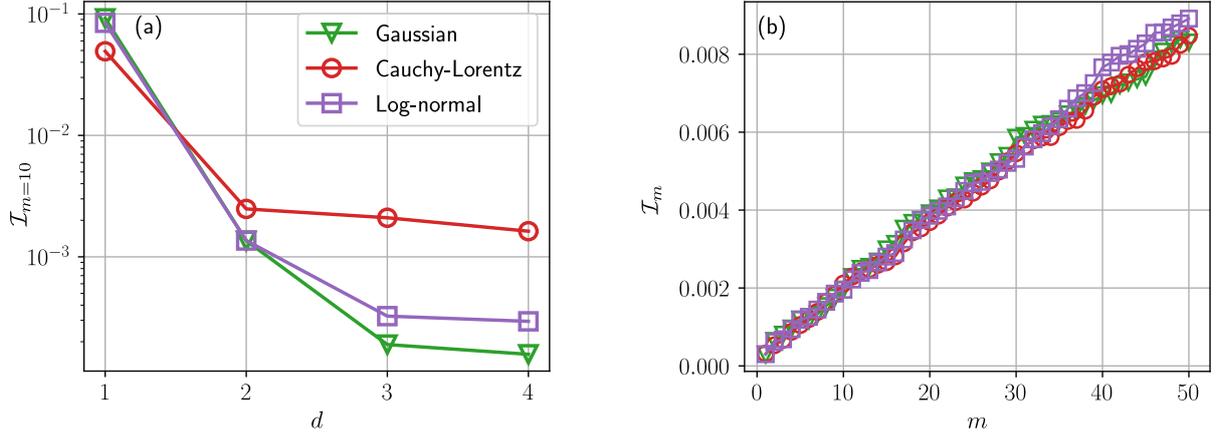


Figure 3: Infidelity \mathcal{I}_m in Eq. (22) for the probability distributions of Eqs. (2)-(4) (see legend) shown (a) as a function of d for $m = 10$ and (b) as a function of m for $d = 4$. We see that the infidelity decreases with increasing d , due to the corresponding increase of the expressive power of the PQC. The infidelity increases linearly as a function of m slowly with a slope of ≈ 0.00017 that is approximately the same for the three distributions. These results are obtained via naïve VQAE with $k = 1$, $M = 50$, and $n_s = 1000$, using the numerically exact gradient without sampling and Adam with the initial learning rate $\beta = 0.1$. We consider 100 randomly initialized PQC and (a) shows one example calculation and (b) the average over all 100 calculations.

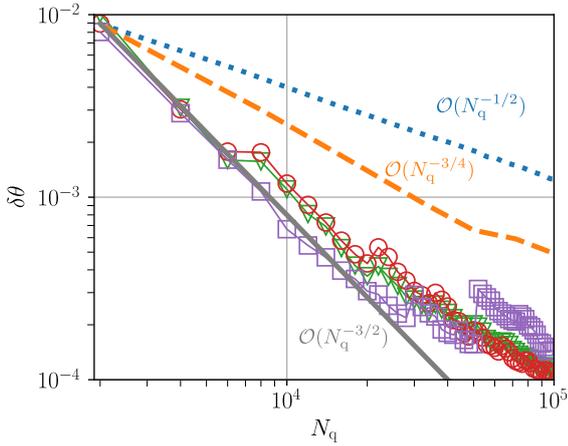


Figure 4: Amplitude estimation error $\delta\theta$ as a function of the number of queries N_q obtained using naïve VQAE with $k = 1$ under the assumption of zero variational cost $N_{\text{var}/1} = 0$. We observe that for small M , the error follows the ideal VQAE scaling $\delta\theta \sim \mathcal{O}(N_q^{-3/2})$ (solid gray line). For larger values of M , the scaling changes to the MC scaling $\delta\theta \sim \mathcal{O}(N_q^{-1/2})$ (dotted blue line). The result is also compared to the MLAE scaling $\delta\theta \sim \mathcal{O}(N_q^{-3/4})$ (dashed orange line). The legend as well as the simulation parameters are the same as in Fig. 3.

3.3 Adaptive VQAE

To reduce the variational cost of VQAE, in the following we present the adaptive VQAE algorithm. In this algorithm, the function f is rescaled such that the k -th power of the Grover operator is close to the identity and only then the variational optimization is carried out. Then the PQC ansatz needs to be just slightly different from the initial state $|\phi_{\text{init}}\rangle$ and the optimization needs fewer samples than naïve VQAE.

To introduce the adaptive VQAE algorithm, we first note that the amplitude $a = \mathbb{E}_p[f]$ – see Eq. (1) – is linear in f , meaning that rescaling the function f with a proportionality constant r also rescales the amplitude a :

$$a' = \mathbb{E}_p[f'] = \mathbb{E}_p[rf] = ra. \quad (23)$$

The rescaled function f' can then be used to encode a new quantum state $|\chi'_0\rangle_{n+1}$ and a new Grover operator \mathcal{Q}' , provided that $0 \leq f'(x) \leq 1$ for all x , which is required for the successful state preparation via Eq. (7). Restricted by this constraint, the rescaling factor has to satisfy $0 \leq rf(x) \leq 1$ for all x . To proceed further, we make the observation that the new Grover operator \mathcal{Q}' implements a rotation by an angle $2\theta'$ in the subspace $\mathcal{H}_{\chi'}$ spanned by good and bad renormalized states, as shown in Eq. (9). Under the commensurability condition

$$a' = \sin^2(\theta'), \quad \theta' = \pi l/k, \quad l \in \mathbb{Z}, \quad (24)$$

applying the renormalized Grover operator \mathcal{Q}'^k results in performing l full rotations in $\mathcal{H}_{\chi'}$. Such a commensurability condition can be achieved by fixing the renormalization factor as

$$r = a'/a = \sin^2(\theta')/a \quad (25)$$

where θ' is uniquely determined by the choice of the desired power k and some integer number l . Hence, we conclude that for a proper choice of r satisfying $0 \leq r \leq 1/\max_x f(x)$, it is possible to rescale the function f so that the k -th power of the corresponding Grover operator acts as identity in the subspace of good and bad states, *i.e.* $\mathcal{Q}'^k = \mathbb{I}_{\mathcal{X}'}$ in theory.

In practice, however, looking at Eq. (25) we see that finding the exact renormalization factor r requires exact knowledge of the initial amplitude a which, of course, we do not have. However, as we show in the following, a loose estimate a_1 , obtained from a moderate number of MC samples of the initial state $|\chi_0\rangle_{n+1}$, is sufficient to get $\mathcal{Q}'^k \approx \mathbb{I}_{\mathcal{X}'}$ and use it in adaptive VQAE. Assuming that such a loose amplitude estimate is provided, a loose renormalization factor can then be expressed as $r_1 = a'/a_1$, with a' defined as in Eq. (24). Because of this imprecise estimation, the actual value of the amplitude after rescaling becomes $a'_1 = \sin^2(\theta'_1) = r_1 a$ and the Grover operator performs l full rotations only approximately, *i.e.* the previous exact identity transforms into $\mathcal{Q}'^k \approx \mathbb{I}_{\mathcal{X}'}$ with a typical phase error per Grover rotation of

$$\delta\theta' = \theta' - \theta'_1 = \theta' - \arcsin \sqrt{r_1 a}. \quad (26)$$

For an unbiased loose estimate with zero average, $\delta\theta'$ can be interpreted as a random error of zero mean. After k Grover rotations, this error becomes k times as large.

Next, we use the VQAE algorithm to estimate the amplitude a'_1 by means of the Grover operator \mathcal{Q}' and the initial state $|\chi'_0\rangle_{n+1}$. The variational approximation is performed at every k -th step, when the overlap

$$\langle \chi'_0 |_{n+1} | \chi'_k \rangle_{n+1} = \cos(2k\theta'_1) = \cos(2k\delta\theta) \quad (27)$$

is expected to be the largest. Here we use that $\theta'_1 + \delta\theta = \theta' = \pi l/k$. Additionally, we assume that the PQC has the initial state $|\phi_{\text{init}}\rangle_{n+1} = |\chi'_0\rangle_{n+1}$ so that the variational quantum state of Eq. (16) reads

$$|\phi_{\text{var}}(\boldsymbol{\lambda})\rangle_{n+1} = \mathcal{U}_{\text{var}}(\boldsymbol{\lambda}) |\chi'_0\rangle_{n+1}. \quad (28)$$

Having the PQC initialized to the identity at the beginning of each optimization step, the only role of the variational quantum circuit is to correct the deviation of Eq. (26) originating from an imprecise value of the renormalization constant r_1 and to bring the overlap of Eq. (27) as close to one as possible. As a consequence, the variational optimization always starts from a good solution and therefore, in general, converges quicker to a better solution than naïve VQAE. This leads to a significant reduction in variational cost of adaptive VQAE compared to the naïve version of the algorithm.

Finally, at the end of the calculation, a maximum likelihood estimation of $a'_1 = \sin^2(\theta'_1) = r_1 a$ is obtained. To go back to the original formulation of the

problem and compare the results, we use the inverse transformation

$$\theta = \arcsin \sqrt{a}, \quad a = a'_1/r_1, \quad (29)$$

where the renormalization constant r_1 has to be exactly the same as the one used for the function rescaling in order for the prior and posterior rescaling errors to cancel out. This last step concludes the adaptive VQAE algorithm which is summarized in terms of pseudocode as Algorithm 2.

Algorithm 2 Adaptive VQAE

Require: functions f and p , integer k

- Use f to get a loose estimate a_1
- Calculate $r_1 = a'/a_1$ and $f' = r_1 f$
- Use f' and p to encode $|\chi'_0\rangle_{n+1}$ and $\mathcal{Q}' = -\mathcal{R}_{\mathcal{X}'} \mathcal{R}_{\text{good}}$
- Use Algorithm 1 to estimate the amplitude $a'_1 = r_1 a$ associated with the state $|\chi'_0\rangle_{n+1}$
- Get the estimate for the original problem as $a = a'_1/r_1$

We analyze the performance of the adaptive VQAE algorithm with a simplified variational ansatz consisting of only six single-qubit rotation gates and four CNOT gates, as shown in Fig. 2 in dark blue color. This simplified ansatz has only six parameters in total, which significantly reduces the number of variational queries as well as the effects of the noise due to finite sampling. We determine the loose estimate of the amplitude a_1 via 5×10^5 MC samples. As a result, much smaller values of infidelity are achieved for n_f being an order of magnitude smaller than in our naïve VQAE computations. We also note that for smaller values of a , the initial MC estimation of a' gets worse and, as a consequence, more sweeps are required to ensure the convergence of the variational ansatz.

Our results for adaptive VQAE are presented in Fig. 1, where we show the convergence of $\delta\theta$ as a function of N_q for $k = 10$. The simulations use Adam with the initial learning rate $\beta = 10^{-3}$, $n_f = 100$, $n_s = 100$, and $n_p = 6$, resulting in $N_{\text{var}/1} = 2n_f n_s n_p = 1.2 \times 10^5$. As in Fig. 4, we compare with the classical MC scaling $\delta\theta \sim \mathcal{O}(N_q^{-1/2})$ and the MLAE scaling $\delta\theta \sim \mathcal{O}(N_q^{-3/4})$. The major difference of the adaptive VQAE as compared to all previously studied methods is a large starting cost which corresponds to the amount of MC samples required for the evaluation of a_1 . This starting cost, however, represents only an additive contribution to N_q and, hence, is insignificant in the regime of our interest when the number of queries gets large. Additionally, we find that, thanks to a significant improvement of the number of query calls and the overall precision of the variational state, the resulting final error $\delta\theta$ of the adaptive VQAE algorithm surpasses the classical MC error.

Interestingly, we observe that in the regime of small k , the performance of the adaptive VQAE algorithm

decreases. This has several reasons. Firstly, the precision of the maximum likelihood estimation decreases when the angle $\theta' = \pi l/k$ (where $l = 1$ in our case) becomes larger than $\pi/4$, *i.e.* for $k \leq 4$. Hence, to perform an estimation with such small values of k , a different statistical inference technique has to be considered. Secondly, for small values of k , the rescaling factor can become much larger than one and then leads to more efficient classical MC sampling. Therefore, in calculations with $C = 0.1$ classical MC sampling performs better than adaptive VQAE for $k \leq 5$, corresponding to $r \gtrsim 7.508$.

To understand how adaptive VQAE performs for increasing qubit counts n , we have run the same simulations as in Fig. 1 for $n = 8, 10$, and 12 . The results for the Gaussian probability distribution are shown in Fig. 5. For the shifted Cauchy-Lorentz and log-normal probability distributions we obtained results (not shown) lying on top of the ones in Fig. 5. We find no significant dependence on n in any of our results. This is surprising as the cost function Eq. (14) is global and therefore the vanishing gradient problem [44] should lead to worse results for larger values of n . We conjecture that the equally good performance of adaptive VQAE for all considered values of n is due to the simple variational ansatz as well as the specific problems studied here. Adaptive VQAE uses the simple ansatz shown in Fig. 2 that consists of only 6 variational angles independent of n . Additionally, the ansatz is composed of nearest-neighbour CNOTs and single-qubit R_y rotation gates, *i.e.* not exact local 2-designs as in [44]. The problems studied here are expectation value calculations where an increased qubit count n leads to an increased number of grid points 2^n for the discretized approximation, see Eq. (1). For the smooth functions considered here, we anticipate that the expectation value converges rapidly with increasing number of grid points.

4 Discussion

In this article, we provide numerical evidence that variational quantum algorithms based on constant-depth quantum circuits can be more efficient than classical MC sampling in the context of amplitude estimation. The quantum circuits used for our numerical demonstrations, however, are still challenging for this generation of gate-based quantum computers. Therefore, an exciting next step is to find other problems and applications for which VQAE has low quantum hardware requirements and can be realized on actual quantum devices.

We can imagine future applications for VQAE in several areas, including combinatorial optimization, quantum machine learning, and quantum chemistry. In the context of combinatorial optimization, VQAE enables us to use constant-depth quantum circuits to carry out Grover search, which can find the opti-

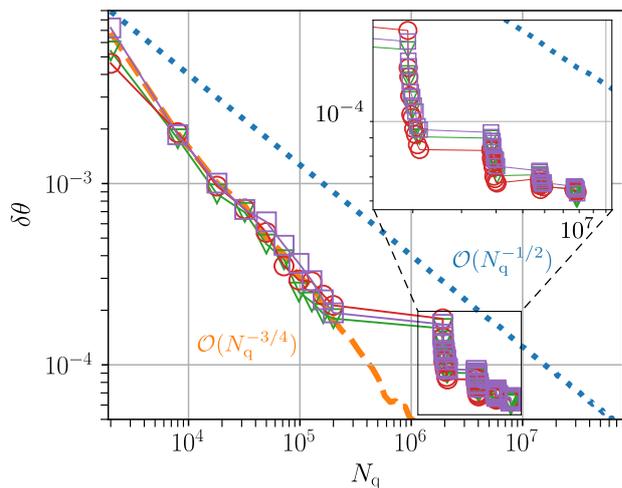


Figure 5: Amplitude estimation error $\delta\theta$ as a function of the number of queries N_q for adaptive VQAE with $n = 8$ (green triangles), 10 (red circles), and 12 (purple squares) qubits for the Gaussian probability distribution in Eq. (2). The remainder of the legend and the optimization procedure are the same as in Fig. 1.

mal solution with a quadratic quantum speedup over brute-force search. Here it is also enticing to study whether such a variational Grover search algorithm can benefit from filtering operators [45]. In relation to quantum machine learning, VQAE has the potential to make it possible for current quantum devices to accelerate inference in Bayesian networks [46], which can then be compared with state-of-the-art variational quantum algorithms for inference [47]. With regards to quantum chemistry, the concept of VQAE can be combined with variational quantum phase estimation (VQPE) [48–50] to realize VQPE with shallow circuits on actual quantum hardware. In this context, one interesting application is to use accurate quantum chemistry results obtained with a quantum computer to train an ansatz for the exchange-correlation energy in density functional theory by means of machine learning [51–53].

We anticipate that the efficiency of the VQAE algorithm can be further increased. Firstly, it would be interesting to analyse whether a local cost function exists – which can help mitigate the negative effect of barren plateaus [44] – that improves the variational optimization and reduces the required number of variational queries. Secondly, the performance of our variational algorithm crucially depends on the maximum likelihood estimation procedure. It would be interesting to investigate whether alternative approaches perform better, *e.g.* iterative QAE [18] or QoPrime AE [19].

Acknowledgements

KP and ML are grateful to David Amaro and Marcello Benedetti for helpful discussions.

References

- [1] Gilles Brassard, Peter Hoyer, Michele Mosca, and Alain Tapp. Quantum amplitude amplification and estimation. *Contemp. Math.*, 305:53–74, 2002. DOI: [10.1090/conm/305](https://doi.org/10.1090/conm/305).
- [2] Ashley Montanaro. Quantum speedup of Monte Carlo methods. *Proc. Math. Phys. Eng. Sci.*, 471(2181), 2015. DOI: [10.1098/rspa.2015.0301](https://doi.org/10.1098/rspa.2015.0301).
- [3] Emanuel Knill, Gerardo Ortiz, and Rolando D. Somma. Optimal quantum measurements of expectation values of observables. *Phys. Rev. A*, 75:012328, Jan 2007. DOI: [10.1103/PhysRevA.75.012328](https://doi.org/10.1103/PhysRevA.75.012328).
- [4] Ivan Kassal, Stephen P Jordan, Peter J Love, Masoud Mohseni, and Alán Aspuru-Guzik. Polynomial-time quantum algorithm for the simulation of chemical dynamics. *Proc. Natl. Acad. Sci. USA*, 105(48):18681–18686, 2008. DOI: [10.1073/pnas.0808245105](https://doi.org/10.1073/pnas.0808245105).
- [5] 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, mar 2015. ISSN 1533-7146. URL <https://dl.acm.org/doi/10.5555/2871393.2871400>.
- [6] Nathan Wiebe, Ashish Kapoor, and Krysta M. Svore. Quantum Deep Learning. *Quantum Info. Comput.*, 16(7-8):541–587, may 2016. ISSN 1533-7146. URL <https://dl.acm.org/doi/10.5555/3179466.3179467>.
- [7] Iordanis Kerenidis, Jonas Landman, Alessandro Luongo, and Anupam Prakash. q-means: A quantum algorithm for unsupervised machine learning. *Adv. Neural Inf. Process. Syst.*, 32, 2019. URL <https://proceedings.neurips.cc/paper/2019/file/16026d60ff9b54410b3435b403afd226-Paper.pdf>.
- [8] Román Orús, Samuel Mugel, and Enrique Lizaso. Quantum computing for finance: Overview and prospects. *Rev. Phys.*, 4:100028, 2019. ISSN 2405-4283. DOI: [10.1016/j.revip.2019.100028](https://doi.org/10.1016/j.revip.2019.100028).
- [9] 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 Trans. Quantum Eng.*, 1:1–24, 2020. DOI: [10.1109/TQE.2020.3030314](https://doi.org/10.1109/TQE.2020.3030314).
- [10] Adam Bouland, Wim van Dam, Hamed Joorati, Iordanis Kerenidis, and Anupam Prakash. Prospects and challenges of quantum finance. *arXiv:2011.06492*, 2020. DOI: [10.48550/arXiv.2011.06492](https://doi.org/10.48550/arXiv.2011.06492).
- [11] Stefan Woerner and Daniel J. Egger. Quantum risk analysis. *npj Quantum Inf.*, 5:15, February 2019. DOI: [10.1038/s41534-019-0130-6](https://doi.org/10.1038/s41534-019-0130-6).
- [12] M. C. Braun, T. Decker, N. Hegemann, S. F. Kerstan, and C. Schäfer. A Quantum Algorithm for the Sensitivity Analysis of Business Risks. *arXiv:2103.05475*, March 2021. DOI: [10.48550/arXiv.2103.05475](https://doi.org/10.48550/arXiv.2103.05475).
- [13] Patrick Reberntrost, Brajesh Gupta, and Thomas R. Bromley. Quantum computational finance: Monte Carlo pricing of financial derivatives. *Phys. Rev. A*, 98:022321, Aug 2018. DOI: [10.1103/PhysRevA.98.022321](https://doi.org/10.1103/PhysRevA.98.022321).
- [14] Nikitas Stamatopoulos, Daniel J. Egger, Yue Sun, Christa Zoufal, Raban Iten, Ning Shen, and Stefan Woerner. Option Pricing using Quantum Computers. *Quantum*, 4:291, July 2020. ISSN 2521-327X. DOI: [10.22331/q-2020-07-06-291](https://doi.org/10.22331/q-2020-07-06-291).
- [15] Yohichi Suzuki, Shumpei Uno, Rudy Raymond, Tomoki Tanaka, Tamiya Onodera, and Naoki Yamamoto. Amplitude estimation without phase estimation. *Quantum Inf. Process.*, 19(2):75, January 2020. DOI: [10.1007/s11128-019-2565-2](https://doi.org/10.1007/s11128-019-2565-2).
- [16] Tomoki Tanaka, Yohichi Suzuki, Shumpei Uno, Rudy Raymond, Tamiya Onodera, and Naoki Yamamoto. Amplitude estimation via maximum likelihood on noisy quantum computer. *Quantum Inf. Process.*, 20(9):293, Sep 2021. ISSN 1573-1332. DOI: [10.1007/s11128-021-03215-9](https://doi.org/10.1007/s11128-021-03215-9).
- [17] Scott Aaronson and Patrick Rall. Quantum approximate counting, simplified. In *Symposium on Simplicity in Algorithms*, pages 24–32. SIAM, 2020. DOI: [10.1137/1.9781611976014.5](https://doi.org/10.1137/1.9781611976014.5).
- [18] Dmitry Grinko, Julien Gacon, Christa Zoufal, and Stefan Woerner. Iterative quantum amplitude estimation. *npj Quantum Inf.*, 7(1):1–6, 2021. DOI: [10.1038/s41534-021-00379-1](https://doi.org/10.1038/s41534-021-00379-1).
- [19] Tudor Giurgica-Tiron, Iordanis Kerenidis, Farrokh Labib, Anupam Prakash, and William Zeng. Low depth algorithms for quantum amplitude estimation. *arXiv:2012.03348*, December 2020. URL <https://arxiv.org/abs/2012.03348>.
- [20] Steven Herbert. Quantum Monte-Carlo Integration: The Full Advantage in Minimal Circuit Depth. *arXiv:2105.09100*, May 2021. DOI: [10.48550/arXiv.2105.09100](https://doi.org/10.48550/arXiv.2105.09100).
- [21] Guoming Wang, Dax Enshan Koh, Peter D. Johnson, and Yudong Cao. Minimizing Estimation Runtime on Noisy Quantum Computers. *PRX Quantum*, 2:010346, Mar 2021. DOI: [10.1103/PRXQuantum.2.010346](https://doi.org/10.1103/PRXQuantum.2.010346).
- [22] Amara Katabarwa, Alex Kunitsa, Borja Peropadre, and Peter Johnson. Reducing runtime and error in VQE using deeper and noisier quantum circuits. *arXiv:2110.10664*, 2021. DOI: [10.48550/arXiv.2110.10664](https://doi.org/10.48550/arXiv.2110.10664).
- [23] Marcello Benedetti, Erika Lloyd, Stefan Sack, and Mattia Fiorentini. Parameterized quantum circuits as machine learning models. *Quan-*

- tum Sci. Technol.*, 4(4):043001, nov 2019. DOI: [10.1088/2058-9565/ab4eb5](https://doi.org/10.1088/2058-9565/ab4eb5).
- [24] M. Cerezo, Andrew Arrasmith, Ryan Babbush, Simon C. Benjamin, Suguru Endo, Keisuke Fujii, Jarrod R. McClean, Kosuke Mitarai, Xiao Yuan, Lukasz Cincio, and Patrick J. Coles. Variational quantum algorithms. *Nat. Rev. Phys.*, 3(9):625–644, Sep 2021. ISSN 2522-5820. DOI: [10.1038/s42254-021-00348-9](https://doi.org/10.1038/s42254-021-00348-9).
- [25] Kishor Bharti, Alba Cervera-Lierta, Thi Ha Kyaw, Tobias Haug, Sumner Alperin-Lea, Abhinav Anand, Matthias Degroote, Hermann Heimonen, Jakob S. Kottmann, Tim Menke, Wai-Keong Mok, Sukin Sim, Leong-Chuan Kwek, and Alán Aspuru-Guzik. Noisy intermediate-scale quantum algorithms. *Rev. Mod. Phys.*, 94(1), Feb 2022. ISSN 1539-0756. DOI: [10.1103/RevModPhys.94.015004](https://doi.org/10.1103/RevModPhys.94.015004).
- [26] Gilles Brassard and Peter Hoyer. An exact quantum polynomial-time algorithm for Simon’s problem. In *Proceedings of the Fifth Israeli Symposium on Theory of Computing and Systems*, pages 12–23. IEEE, 1997. DOI: [10.1109/ISTCS.1997.595153](https://doi.org/10.1109/ISTCS.1997.595153).
- [27] Lov K Grover. Quantum computers can search rapidly by using almost any transformation. *Phys. Rev. Lett.*, 80(19):4329, 1998. DOI: [10.1103/PhysRevLett.80.4329](https://doi.org/10.1103/PhysRevLett.80.4329).
- [28] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical Recipes in C*. Cambridge University Press, Cambridge, USA, second edition, 1992. URL <https://dl.acm.org/doi/10.5555/148286>.
- [29] J. Kennedy and R. Eberhart. Particle swarm optimization. In *Proceedings of ICNN’95 - International Conference on Neural Networks*, volume 4, pages 1942–1948, 1995. DOI: [10.1109/ICNN.1995.488968](https://doi.org/10.1109/ICNN.1995.488968).
- [30] Mohammad Reza Bonyadi and Zbigniew Michalewicz. Particle Swarm Optimization for Single Objective Continuous Space Problems: A Review. *Evol. Comput.*, 25(1):1–54, 03 2017. ISSN 1063-6560. DOI: [10.1162/EVCO_r_00180](https://doi.org/10.1162/EVCO_r_00180).
- [31] Javier Gil Vidal and Dirk Oliver Theis. Calculus on parameterized quantum circuits. *arXiv:1812.06323*, dec 2018. URL <https://arxiv.org/abs/1812.06323>.
- [32] Robert M. Parrish, Joseph T. Iosue, Asier Ozaeta, and Peter L. McMahon. A Jacobi Diagonalization and Anderson Acceleration Algorithm For Variational Quantum Algorithm Parameter Optimization. *arXiv:1904.03206*, 2019. DOI: [10.48550/arXiv.1904.03206](https://doi.org/10.48550/arXiv.1904.03206).
- [33] Ken M. Nakanishi, Keisuke Fujii, and Synge Todo. Sequential minimal optimization for quantum-classical hybrid algorithms. *Phys. Rev. Research*, 2:043158, Oct 2020. DOI: [10.1103/PhysRevResearch.2.043158](https://doi.org/10.1103/PhysRevResearch.2.043158).
- [34] Mateusz Ostaszewski, Edward Grant, and Marcello Benedetti. Structure optimization for parameterized quantum circuits. *Quantum*, 5:391, 2021. DOI: [10.22331/q-2021-01-28-391](https://doi.org/10.22331/q-2021-01-28-391).
- [35] Marcello Benedetti, Mattia Fiorentini, and Michael Lubasch. Hardware-efficient variational quantum algorithms for time evolution. *Phys. Rev. Research*, 3(3):033083, 2021. DOI: [10.1103/PhysRevResearch.3.033083](https://doi.org/10.1103/PhysRevResearch.3.033083).
- [36] Elad Hazan, Kfir Y Levy, and Shai Shalev-Shwartz. Beyond convexity: Stochastic quasi-convex optimization. *arXiv:1507.02030*, 2015. URL <https://arxiv.org/abs/1507.02030>.
- [37] Yudai Suzuki, Hiroshi Yano, Rudy Raymond, and Naoki Yamamoto. Normalized Gradient Descent for Variational Quantum Algorithms. *arXiv:2106.10981*, 2021. DOI: [10.48550/arXiv.2106.10981](https://doi.org/10.48550/arXiv.2106.10981).
- [38] Jun Li, Xiaodong Yang, Xinhua Peng, and Chang-Pu Sun. Hybrid Quantum-Classical Approach to Quantum Optimal Control. *Phys. Rev. Lett.*, 118(15):150503, April 2017. DOI: [10.1103/PhysRevLett.118.150503](https://doi.org/10.1103/PhysRevLett.118.150503).
- [39] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii. Quantum circuit learning. *Phys. Rev. A*, 98(3):032309, September 2018. DOI: [10.1103/PhysRevA.98.032309](https://doi.org/10.1103/PhysRevA.98.032309).
- [40] Maria Schuld, Ville Bergholm, Christian Gogolin, Josh Izaac, and Nathan Killoran. Evaluating analytic gradients on quantum hardware. *Phys. Rev. A*, 99(3):032331, mar 2019. DOI: [10.1103/PhysRevA.99.032331](https://doi.org/10.1103/PhysRevA.99.032331).
- [41] Leonardo Banchi and Gavin E Crooks. Measuring analytic gradients of general quantum evolution with the stochastic parameter shift rule. *Quantum*, 5:386, 2021. DOI: [10.22331/q-2021-01-25-386](https://doi.org/10.22331/q-2021-01-25-386).
- [42] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. *arXiv:1412.6980*, December 2014. URL <https://arxiv.org/abs/1412.6980>.
- [43] Dorit Aharonov, Vaughan Jones, and Zeph Landau. A polynomial quantum algorithm for approximating the Jones polynomial. *Algorithmica*, 55(3):395–421, 2009. DOI: [10.1007/s00453-008-9168-0](https://doi.org/10.1007/s00453-008-9168-0).
- [44] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles. Cost function dependent barren plateaus in shallow parameterized quantum circuits. *Nat. Commun.*, 12:1791, 2021. DOI: [10.1038/s41467-021-21728-w](https://doi.org/10.1038/s41467-021-21728-w).
- [45] David Amaro, Carlo Modica, Matthias Rosenkranz, Mattia Fiorentini, Marcello Benedetti, and Michael Lubasch. Filtering variational quantum algorithms for combinatorial optimization. *Quantum Sci. Technol.*, 7(1):015021, Jan 2022. ISSN 2058-9565. DOI: [10.1088/2058-9565/ac3e54](https://doi.org/10.1088/2058-9565/ac3e54).

- [46] Guang Hao Low, Theodore J. Yoder, and Isaac L. Chuang. Quantum inference on Bayesian networks. *Phys. Rev. A*, 89:062315, Jun 2014. DOI: [10.1103/PhysRevA.89.062315](https://doi.org/10.1103/PhysRevA.89.062315).
- [47] Marcello Benedetti, Brian Coyle, Mattia Fiorentini, Michael Lubasch, and Matthias Rosenkranz. Variational inference with a quantum computer. *Phys. Rev. Applied*, 16:044057, Oct 2021. DOI: [10.1103/PhysRevApplied.16.044057](https://doi.org/10.1103/PhysRevApplied.16.044057).
- [48] Robert M. Parrish and Peter L. McMahon. Quantum Filter Diagonalization: Quantum Eigendecomposition without Full Quantum Phase Estimation. *arXiv:1909.08925*, 2019. DOI: [10.48550/arXiv.1909.08925](https://doi.org/10.48550/arXiv.1909.08925).
- [49] Nicholas H. Stair, Renke Huang, and Francesco A. Evangelista. A Multireference Quantum Krylov Algorithm for Strongly Correlated Electrons. *J. Chem. Theory Comput.*, 16: 2236, 2020. DOI: [10.1021/acs.jctc.9b01125](https://doi.org/10.1021/acs.jctc.9b01125).
- [50] Katherine Klymko, Carlos Mejuto-Zaera, Stephen J. Cotton, Filip Wudarski, Miroslav Urbanek, Diptarka Hait, Martin Head-Gordon, K. Birgitta Whaley, Jonathan Moussa, Nathan Wiebe, Wibe A. de Jong, and Norm M. Tubman. Real time evolution for ultracompact Hamiltonian eigenstates on quantum hardware. *arXiv:2103.08563*, 2021. DOI: [10.48550/arXiv.2103.08563](https://doi.org/10.48550/arXiv.2103.08563).
- [51] John C. Snyder, Matthias Rupp, Katja Hansen, Klaus-Robert Müller, and Kieron Burke. Finding Density Functionals with Machine Learning. *Phys. Rev. Lett.*, 108:253002, Jun 2012. DOI: [10.1103/PhysRevLett.108.253002](https://doi.org/10.1103/PhysRevLett.108.253002).
- [52] Michael Lubasch, Johanna I. Fuks, Heiko Appel, Angel Rubio, J. Ignacio Cirac, and Mari-Carmen Bañuls. Systematic construction of density functionals based on matrix product state computations. *New J. Phys.*, 18(8):083039, aug 2016. DOI: [10.1088/1367-2630/18/8/083039](https://doi.org/10.1088/1367-2630/18/8/083039).
- [53] Sebastian Dick and Marivi Fernandez-Serra. Machine learning accurate exchange and correlation functionals of the electronic density. *Nat. Commun.*, 11:3509, 2020. DOI: [10.1038/s41467-020-17265-7](https://doi.org/10.1038/s41467-020-17265-7).