

Alternating Layered Variational Quantum Circuits Can Be Classically Optimized Efficiently Using Classical Shadows

Afrad Basheer*, Yuan Feng, Christopher Ferrie, Sanjiang Li

Centre for Quantum Software and Information, University of Technology Sydney, NSW 2007, Australia

Afrad.M.Basheer@student.uts.edu.au, {Yuan.Feng, Christopher.Ferrie, Sanjiang.Li}@uts.edu.au

Abstract

Variational quantum algorithms (VQAs) are the quantum analog of classical neural networks (NNs). A VQA consists of a parameterized quantum circuit (PQC) which is composed of multiple layers of ansatzes (simpler PQCs, which are an analogy of NN layers) that differ only in selections of parameters. Previous work has identified the alternating layered ansatz as potentially a new standard ansatz in near-term quantum computing. Indeed, shallow alternating layered VQAs are easy to implement and have been shown to be both trainable and expressive. In this work, we introduce a training algorithm with an exponential reduction in training cost of such VQAs. Moreover, our algorithm uses classical shadows of quantum input data, and can hence be run on a classical computer with rigorous performance guarantees. We demonstrate 2–3 orders of magnitude improvement in the training cost using our algorithm for the example problems of finding state preparation circuits and the quantum autoencoder.

Introduction

The past several years have seen tremendous progress in quantum hardware and one of the state-of-the-art quantum processors—IBM Osprey—now has 433 qubits (Gambetta 2022). However, it is widely believed that general-purpose fault-tolerant quantum computers are unlikely to happen soon. With a limited number of (noisy) qubits and non-perfect gates, near-term quantum devices like IBM Osprey can only run quantum circuits with small depth. One may wonder if quantum advantage can still be demonstrated in practical applications using these near-term quantum devices. Variational Quantum Algorithms (VQAs) are the most promising proposal for this purpose. VQAs encode a task in a parametrized quantum circuit (PQC), evaluate it using a near-term quantum computer, and optimize the parameters with a classical optimizer. For almost all applications for which quantum computers are desired, VQAs have been proposed. We refer the readers to Cerezo et al. (2021a) for an excellent survey of VQAs.

VQAs can be regarded as the quantum analog of classical neural networks (NNs). A VQA is composed of multiple layers of ansatzes, which are smaller PQCs with a

fixed architecture. An ansatz is analogous to a classical NN layer, and a VQA uses ansatzes that have the same architecture and differ only in selections of parameters. Many different ansatzes have been proposed in the literature. Popular choices include the Hardware Efficient Ansatz (Kandala et al. 2017), Quantum Alternating Operator Ansatz (QAOA) (Hadfield et al. 2019), etc.

Similar to classical NNs, recent works show that VQAs also have the problem of trainability. Indeed, vanishing gradients, also called *barren plateaus*, were first theoretically demonstrated for deep PQCs by McClean et al. (2018) and then for shallow PQCs by Cerezo et al. (2021b). We may need to impose restrictions on the ansatz and/or the cost function to avoid this barren plateaus issue. One such solution is proposed by Cerezo et al. (2021b) using the alternating layered ansatz, which has a brick-like structure and its two-qubit gates act on alternating pairs of neighboring qubits (see Fig. 1 for a simple example). It is proved that barren plateaus can be avoided for alternating layered VQAs provided that the depth of the PQC is $O(\log n)$, where n is the number of qubits, and the cost function is defined with local observables. Surprisingly, Nakaji and Yamamoto (2021) recently proved that the shallow alternating layered ansatz is almost as expressive as the Hardware Efficient Ansatz. Thus the alternating layered ansatz is both expressive and trainable. In addition, this ansatz has been investigated or implemented in works such as Hinsche et al. (2021), Wu et al. (2021), Arrasmith et al. (2021), Slattery, Villalonga, and Clark (2022). In particular, Cerezo et al. (2021b) has shown that quantum autoencoders can be optimized without any barren plateaus for up to 100 qubits using the alternating layered ansatz.

This work introduces a training algorithm with an exponential improvement in the number of copies of input states consumed during training an alternating layered VQA with shallow depth and local observables, where, and in the remainder of this paper, the number of copies of the input state equals the number of executions of the quantum device. The reason is, during each execution, we have to measure the quantum systems which inevitably destroy the quantum states. Thus a fresh copy of the input state is needed for each iteration. Moreover, the training can be done entirely on a classical computer efficiently (with computational cost depending only polynomially on n) without the need to imple-

*Corresponding Author

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: An illustration of alternating layered ansatzes where the parameterized sub-circuit $S(\theta_{32})$ is applied on the first and the last qubits. Here, θ is an order 3 tensor with each θ_{ij} being vectors of real parameters.

ment the PQC on a quantum device. This result is achieved by using the recently proposed classical shadow technique (Aaronson 2018; Huang et al. 2021b) for quantum state tomography, and working in the Heisenberg picture rather than the Schrödinger model.

Specifically, for an alternating layered ansatz $U(\theta)$, an input state ρ and an observable O , the VQAs of our interest estimate each evaluation of functions of the form

$$f_{\rho,O}(\theta) = \text{tr}(OU(\theta)\rho U(\theta)^\dagger) \quad (1)$$

using quantum computers. In contrast, our method can efficiently compute this classically on classical shadows of ρ . But note that all VQAs that uses alternating layered ansatzes need not have this specific form.

Our method, called *Alternating Layered Shadow Optimization*, or simply ALSO, outperforms standard alternating layered VQA in two aspects:

1. *Exponential savings on input state copies.* Note that the number of copies of the input state needed in the standard VQA scales linearly in the total number of function evaluations required. In contrast, to achieve a similar precision, ALSO only uses logarithmically many copies. This allows for more iterations and better approximations in the classical optimization algorithm for a given PQC. In addition, it allows for more hyperparameter tuning with very few copies of the input state, and the same set of shadows can be used for multiple similar optimization problems and alternating layered ansatzes.
2. *Easy implementation on quantum hardware.* ALSO only requires the quantum device to be able to carry out single-qubit Pauli basis measurements on the input states. But standard VQA requires the ability to apply CNOT gates and rotation gates on them, and measurement also.

We demonstrate the practical efficacy of our result with two important examples: finding state preparation circuits and quantum data compression using a quantum autoencoder. In both cases, we demonstrate that ALSO can match the results of the impossible ideal VQA that uses infinite copies, using a comparatively small number of copies of the input quantum state. We also show that, with the same number of copies of the input state, ALSO outperforms the standard VQA significantly.

Related Works

The idea of using classical algorithms, specifically classical machine learning algorithms, on classical shadows of quantum states, has been considered in Huang et al. (2021a), Huang et al. (2021b), where popular classical machine learning algorithms such as Support Vector Machines and Convolutional Neural Networks are trained on classically loaded shadows to solve certain important problems in quantum many-body physics. In comparison, the post-processing of classical shadows in ALSO aims to implement a VQA in a more resource-efficient manner.

Studying alternating layered VQAs as optimizations of local parameterized observables is already considered in Okada et al. (2022). Here, the locality of W_O (cf. Figure 3) is leveraged to implement variational quantum eigensolvers. ALSO can be seen as a generalization of this method because the input state can be arbitrary in our setting (due to the use of classical shadows).

Fontana et al. (2022) have experimentally shown that in certain cases, the quantum alternating operator ansatz (Farhi, Goldstone, and Gutmann 2014) and the Hamiltonian variational ansatz (Wiersema et al. 2020) can be trained by estimations in the Fourier basis (Schuld, Sweke, and Meyer 2021), with $\mathcal{O}(\text{poly}(n))$ copies of the input state and $\mathcal{O}(\text{poly}(n))$ classical computational cost. Similarly, Yao, Bukov, and Lin (2020) showed that reinforcement learning can be used to optimize the learning process of variational parameters in the former ansatz. In our work, we focus on the alternating layered ansatz and prove theoretically the existence of sample efficient training methods for this ansatz.

In Stokes et al. (2020), Boyd and Koczor (2022) and Wierichs, Gogolin, and Kastoryano (2020), new classical optimization algorithms are introduced and analyzed (the latter also uses classical shadows) and are shown to converge using much fewer iterations compared to the standard gradient descent. But since ALSO is agnostic towards the choice of the classical optimizer, our method can be used to boost the performance of these methods by significantly reducing the number of state copies used. Variational shadow quantum circuits, developed by Li, Song, and Wang (2021), extract local classical features by focusing on a series of local subcircuits. Although inspired by the classical shadow work, the approach itself does not use classical shadows.

Background

This section recalls background in quantum computing, variational quantum algorithms and classical shadows.

Quantum Computing

We use ‘ket’ notation such as $|\psi\rangle$ to represent complex column vectors. For any $|\psi\rangle \in \mathbb{C}^d$, $\langle\psi|$, called ‘bra’, is its complex conjugated transpose. For any $i \in \{0, 1, \dots, d-1\}$, $|i\rangle \in \mathbb{C}^d$ is the i^{th} standard basis vector. Denote by $\mathcal{L}(\mathbb{C}^d)$ the set of all linear operators acting on \mathbb{C}^d .

A *qubit* is the fundamentally implementable entity in quantum computing. A *state* is defined as any positive semi-definite operator $\rho \in \mathcal{L}(\mathbb{C}^d)$ such that $\text{tr}(\rho) = 1$ (all the diagonal elements sum up to 1). A qubit can admit any quantum

state $\rho \in \mathcal{L}(\mathbb{C}^2)$ as its value, similar to how a *bit* in classical computing can admit any value in $\{0, 1\}$. A quantum *system* or *register* is an array of qubits. To describe the state of a system of two qubits, $[q_1, q_2]$, we use states that act on the *tensor product* of the two 2-dimensional vector spaces, denoted as $\mathbb{C}^2 \otimes \mathbb{C}^2 \cong \mathbb{C}^4$. So, a system of n -qubits can be in any state in $\mathcal{L}(\mathbb{C}^{2^n})$. A state ρ is called a *pure state* if its rank is 1. In this case, ρ , as well as its dynamics, can be fully characterized by any normalized eigenvector associated with the eigenvalue 1.

An n -qubit *quantum gate* is defined as a unitary operator $U \in \mathcal{L}(\mathbb{C}^{2^n})$. The application of such a quantum gate on a system in a state $\rho \in \mathcal{L}(\mathbb{C}^{2^n})$ transforms the state of the system from ρ to another state $U\rho U^\dagger$. Some important gates that feature extensively in this work are

$$H = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}, S = \begin{bmatrix} 1 & 0 \\ 0 & i \end{bmatrix}, \text{CNOT} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}.$$

Let $Z = |0\rangle\langle 0| - |1\rangle\langle 1|$, $X = HZH$, and $Y = iXZ$. These single-qubit gates are called *Pauli operators*. For $P \in \{X, Y, Z\}$ and $\theta \in \mathbb{R}$, the P -rotation gate $R_P(\theta)$ is the gate $\cos(\theta/2)\mathbb{1} + i\sin(\theta/2)P$, where $\mathbb{1}$ is the identity operator.

An n -qubit *observable* is defined as any Hermitian operator $O \in \mathcal{L}(\mathbb{C}^{2^n})$. To “observe” information from a system prepared in a state ρ , we *measure* ρ using O . The measurement result is probabilistic and its expected value is $\text{tr}(O\rho)$.

An n -qubit operator $V \in \mathcal{L}(\mathbb{C}^{2^n})$ on register $I = [q_1, \dots, q_n]$ is k -local if there is a k -qubit sub-register $A = [q_{i_1}, q_{i_2}, \dots, q_{i_k}]$ such that $V = \tilde{V} \otimes \mathbb{1}$, where $\tilde{V} \in \mathcal{L}(\mathbb{C}^{2^k})$ acts on qubits in A and $\mathbb{1}$ is the identity operator on $I \setminus A$. For simplicity, we often write V as $\tilde{V}[A]$ and say that V acts *non-trivially* on qubits only in A .

Classical Shadows Using Pauli Basis Measurements

Let ρ be an n -qubit quantum state and let O_1, O_2, \dots, O_M be arbitrary n -qubit observables the classical descriptions of which are given. Using conventional quantum tomography techniques, $O(2^n)$ copies of ρ are required to estimate $\text{tr}(O_i\rho)$ for each O_i .

The classical shadow technique (Huang, Kueng, and Preskill 2020), developed from shadow tomography (Aaronson 2018), provides succinct classical descriptions of quantum states. Using this technique, $\text{tr}(O_i\rho)$ can be collectively predicted by consuming only $\mathcal{O}(\log M)$ copies of ρ . Moreover, when these observables belong to certain classes, the dependency on n is polynomial or constant.

When the observables are all local (with locality $k \ll n$), the classical shadow method reduces to a very simple protocol. The first step is to measure the individual qubits of ρ on a random Pauli basis. To this end, for each qubit i , we apply a gate U_i uniformly randomly chosen from $\{\mathbb{1}, H, HS^\dagger\}$, and then measure it in the computational basis. Let the measurement outcome be $u_i \in \{0, 1\}$. Then a *single-qubit classical shadow* of ρ is calculated (classically) as

$$\hat{\rho} = F(U_1^\dagger |u_1\rangle\langle u_1| U_1) \otimes \dots \otimes F(U_n^\dagger |u_n\rangle\langle u_n| U_n), \quad (2)$$

where $F(V) = 3V - \mathbb{1}$. As a fully separable matrix, $\hat{\rho}$ can be stored efficiently as $n \times 2 \times 2$ matrices. Furthermore, $\hat{\rho}$ gives an unbiased estimation of the unknown state ρ and hence $\text{tr}(O_i\hat{\rho})$ is an unbiased estimator of $\text{tr}(O_i\rho)$ for all i .

Specifically, we have:

Theorem 1. (Sack et al. 2022) *Let $\rho \in \mathbb{C}^{2^n}$ be a quantum state. Suppose $O_1, O_2, \dots, O_M \in \mathbb{C}^{2^n}$ are M k -local observables. For any $\delta, \epsilon \in (0, 1)$, let T be any integer not smaller than $\frac{4^{k+1}}{\epsilon^2} \cdot \log(\frac{2M}{\delta}) \max_i \|O_i\|_\infty^2$ and define shadow state $\hat{\rho}_T$ as $\hat{\rho}_T = \frac{1}{T} \sum_{j=1}^T \hat{\rho}^{(j)}$, where $\hat{\rho}^{(j)}$ are single-qubit classical shadows as in Eq. (2). Then, with probability at least $1 - \delta$ and for all i , we have $|\text{tr}(O_i\hat{\rho}_T) - \text{tr}(O_i\rho)| \leq \epsilon$.*

Note that the original version of this theorem required $\|O_i\|_\infty \leq 1$ for all i . However, this can be relaxed by dividing every matrix by $\max_i \|O_i\|_\infty$ and then estimating with precision $\epsilon / \max_i \|O_i\|_\infty$.

Moreover, each estimation $\text{tr}(O_i\hat{\rho}_T)$ can be classically computed very efficiently. Let $A_i = [q_{i_1}, \dots, q_{i_k}]$ be the sub-register that O_i acts non-trivially on and $O_i = \tilde{O}_i \otimes \mathbb{1}$ with $\tilde{O}_i \in \mathcal{L}(\mathbb{C}^{2^k})$. For the shadow $\hat{\rho}$ in Eq. (2), we only need to use the $k \times 2 \times 2$ matrices corresponding to the sub-register A_i . Denote by $\hat{\rho}_T|_{A_i}$ the classical shadow obtained by taking average of T such reduced shadows. Then we have $\text{tr}(O_i\hat{\rho}_T) = \text{tr}(\tilde{O}_i\hat{\rho}_T|_{A_i})$ and hence it can be computed with cost exponential only in k and independent of n .

Variational Quantum Algorithms

A VQA encodes the task under consideration as a parameterized quantum circuit (PQC), which is typically composed of multiple layers of ansatzes, i.e., smaller PQCs with the same architecture. Write $U(\theta)$ for the PQC, where θ is a real-valued vector of parameters. The VQA uses $U(\theta)$ to estimate a target function’s value and gradient at some point in its domain, and then optimizes the parameters of the PQC by feeding the circuit’s output to a classical optimizer.

In this paper, we focus on the basic function as specified in Eq. (1) and aim to find the parameters which maximize it. Note that $f_{\rho, O}(\theta)$ can be estimated through multiple measurements. Using techniques such as the parameter shift rule (Mitarai et al. 2018), finite differences, etc, one can also estimate the gradient of $f_{\rho, O}$ using a quantum computer. With this, we can find the optimum values of θ by using any classical optimization method. Many tasks such as variational quantum eigensolver, finding state preparation circuits, quantum autoencoder, etc can be reduced to finding the best θ which maximizes some $f_{\rho, O}(\theta)$.

Alternating Layered Shadow Optimization

In this section, we explain the key idea and theoretical results behind ALSO. Following Cerezo et al. (2021b), we request the observables to be k -local or linear combinations of a small number (polynomially dependent on n) of k -local observables for some $k \ll n$.

Alternating Layered Ansatz

The Alternating Layered Ansatz is the brick-like circuit structure presented in Fig. 1, where each $S(\theta_{ij})$ is a param-




Figure 2: The structure of $S(\gamma)$ used in the simulation. The two-qubit gate used here is the CNOT gate.




Figure 3: The structure of $W_O(\theta) = U(\theta)^\dagger O U(\theta)$ where the blue box is a 1-local observable, and all other boxes are S sub-circuits. Except for the red ones, all other sub-circuits cancel each other out, resulting in $W_O(\theta)$ being $2d$ -local.

eterized circuit acting on a small number of qubits. A simple example of S is given in Fig. 2. This work assumes that each S acts on two qubits and has p real parameters, but our idea can be easily extended to the general case. In Fig. 1, the total number of vertical blocks of S gates, written d , is called the *depth* of the ansatz. The circuit depicted in the figure has $d = 3$. In a specific vertical block j , each circuit $S(\theta_{ij})$ acts on qubits $2(i-1) \oplus j$ and its neighbor $2(i-1) \oplus j \oplus 1$. So, the final form of the circuit is given as

$$U(\theta) = \prod_{j=1}^d \prod_{i=1}^{n/2} S(\theta_{ij})[2(i-1) \oplus j, 2(i-1) \oplus j \oplus 1], \quad (3)$$

where \oplus denotes addition modulo n , $\theta \in \mathbb{R}^{\frac{n}{2} \times d \times p}$ is a tensor of real parameters where θ_{ij} is a p -dimensional real vector of parameters and $S(\theta_{ij})[k, l]$ means $S(\theta_{ij})$ acting on qubits k and l .

Method

We first explain our approach in a simpler model, with 1-local observables and alternating layered ansatzes built with 2-local circuits, and then extend the results to circuits and observables with arbitrary localities. The detailed proof of our results can be found in the Technical Appendix.

We start with a lemma that forms the backbone of ALSO.

Lemma 1. *Let d, S and U be defined as in Eq. (3), and $\theta \in \mathbb{R}^{\frac{n}{2} \times d \times p}$. For any n -qubit 1-local observable O , define $W_O(\theta) = U(\theta)^\dagger O U(\theta)$. Then we have $\|W_O(\theta)\|_\infty = \|O\|_\infty$ and $W_O(\theta)$ is $2d$ -local, that is, $W_O(\theta) = \widetilde{W}_O(\theta)[A]$ for some sub-register A of $2d$ qubits.*

Moreover, from Figure 3, we can see that for any θ , $\widetilde{W}_O(\theta)$ can be computed with cost exponential only in d using tensor contractions of the observable with all the S gates marked in red.

Let $U(\theta)$ be the alternating layered ansatz defined as in Eq. (3), and $O = \sum_{i=1}^M O_i$ be the observable, where each O_i

is 1-local. Assume that we are using an iterative optimization algorithm, one that takes as input a target function and outputs its optimizer, to find the maximizer of Eq. (1), and the whole optimization procedure requires C function evaluations of the form $f_{\rho, O}(\theta^{(1)}), f_{\rho, O}(\theta^{(2)}), \dots, f_{\rho, O}(\theta^{(C)})$. Lemma 1 says that each function evaluation can be seen as estimating the expectation of $2d$ -local observables, because $f_{\rho, O}(\theta) = \sum_i \text{tr}(O_i U(\theta) \rho U(\theta)^\dagger) = \sum_i \text{tr}(W_{O_i}(\theta) \rho)$.

Now using Theorem 1, we can estimate all the C function evaluations and the whole ALSO algorithm goes as follows:

1. Load $T = \mathcal{O}(\log(C) \cdot \text{poly}(n))$ classical shadows of ρ . Let $\hat{\rho}_T$ be the shadow state (cf. Theorem 1).
2. For all i , compute $\hat{\rho}_T|_{A_i}$, where A_i is the sub-register that W_{O_i} acts non-trivially on.
3. Use the iterative optimization algorithm to optimize the target function $\hat{f}_{\rho, O}(\theta) = \sum_i \text{tr}(\widetilde{W}_{O_i}(\theta) \hat{\rho}_T|_{A_i})$.

Note that the cost of classical computation is dominated by the computation of $\sum_i \text{tr}(\widetilde{W}_{O_i}(\theta) \hat{\rho}_T|_{A_i})$ and so it scales exponentially only in d . Hence, when $d = \mathcal{O}(\log n)$, the classical computational cost scales polynomially on n .

Sample Complexity

In this section, we discuss the sampling complexity of the protocol, that is, the range of values of T that guarantee good estimations of all the function evaluations. We show that when $d = \mathcal{O}(\log n)$, the sample complexity is $\mathcal{O}(\log(C) \cdot \text{poly}(n))$.

Theorem 2. *Let d, S and U be defined as in Eq. (3). Suppose ρ is an arbitrary n -qubit state and $O = \sum_{i=1}^M O_i$, where each O_i is an n -qubit 1-local observable. Then, for any $\delta, \epsilon \in (0, 1)$ and any C parameter tensors $\theta^{(1)}, \theta^{(2)}, \dots, \theta^{(C)}$, all values $f_{\rho, O}(\theta^{(c)})$ can be estimated using $\hat{f}_{\rho, O}(\theta^{(c)}) := \text{tr}(W_O(\theta^{(c)}) \hat{\rho}_T)$ with the guarantee*

$$\text{Prob} \left(\bigcap_{c=1}^C \left[\left| f_{\rho, O}(\theta^{(c)}) - \hat{f}_{\rho, O}(\theta^{(c)}) \right| \leq \epsilon \right] \right) \geq 1 - \delta \quad (4)$$

where

$$T \geq M^2 \log \left(\frac{2MC}{\delta} \right) \cdot \frac{4^{2d+1}}{\epsilon^2} \max_i \|O_i\|_\infty^2.$$

This is remarkable as, without using classical shadows, we may need to estimate $f_{\rho, O}(\theta)$ for any parameter tensor θ through measurements. Suppose K copies of ρ are consumed to estimate each of these values. Then we end up consuming CK copies of ρ , which can be exponentially larger than the number consumed by ALSO.

In our method, the measurements that we have to make are solely for computing the classical shadows and hence are independent of all $\theta^{(c)}$. Moreover, each measurement outcome can be reused multiple times. In the standard method of training VQAs, we are not able to reuse the measurement outcomes that are made as part of the optimization, because each measurement outcome is dependent on the input parameter $\theta^{(c)}$. This is a crucial reason as to why ALSO is

a much more appealing option to optimize these functions, especially from a practical perspective where one has to do hyperparameter tuning, find the right classical optimizer, etc.

One important point to note is that even though the constants look large, in practice, we need not necessarily require this many copies (classical shadows) of ρ . This is illustrated in our experimental results, where we are able to match the results of ideal VQA simulations (simulations that use infinite copies of the input state ρ) by using a number of copies of ρ orders of magnitude fewer than the number suggested by Theorem 2.




Figure 4: Plot showing time (in seconds) taken for a single function evaluation using ALSO.

The space complexity of the protocol is dominated by the storage of the matrices $\hat{\rho}_T|_{A_i}$. Since the dimension of each of these matrices is 4^d , we need $16^d M$ complex numbers to store all of them. For example, let $n, M = 50, d = 5$. Then we see that if we are using 128 bits to store each complex number, then we only require 838MB to store all matrices $\hat{\rho}_T|_{A_i}$. Time taken for single function evaluations is plotted in Figure 4. On the x -axis, we have the number of qubits, and on the y -axis, we have the time (in seconds) taken to compute a single function evaluation, averaged over 5 cases. In each case, $d = \lfloor \log n \rfloor$ and the observable O is a 1-local observable, with S being the circuit in Figure 2. We plot the results showing a polynomial dependence of time on the number of qubits for both ‘complex128’ and ‘complex64’ being used as datatypes in Python. The simulation was carried out on a laptop with 16GB RAM and 2.6GHz Intel i7 processor.

One can easily generalize Theorem 2 for arbitrarily local observables and circuits. In a similar setting, if we use k_0 -local parameterized circuits and an observable that is a sum of k_1 -local observables, then we can carry out an iterative optimization algorithm with all function evaluations satisfying Eq. (4) using

$$T \geq \frac{M^2}{\epsilon^2} \cdot \log\left(\frac{2MC}{\delta}\right) \cdot 4^{k_1 + (2k_0 - 2)d - 1} \cdot \max_i \|O_i\|_\infty^2$$

copies of the input state. This is because for each increment in depth, the locality of $W_{O_i}(\theta)$ increases by at most $2k_0 - 2$, starting from k_1 .

Applications

Here we discuss two practical applications of ALSO in the field of quantum information.

The State Preparation Problem

Let $\rho = |\psi\rangle\langle\psi| \in \mathcal{L}(\mathbb{C}^{2^n})$ be an n -qubit pure quantum state. The state preparation problem of ρ intends to find the parameters of an alternating layered ansatz which best approximates (heuristically) the state $|\psi\rangle$. That is, we would like to find a parameter tensor such that $1 - |\langle\psi|U(\theta)^\dagger|\mathbf{0}\rangle|^2$ is minimized over all θ , where $|\mathbf{0}\rangle \in \mathbb{C}^{2^n}$. This quantity is called the *infidelity* between the states $U(\theta)^\dagger|\mathbf{0}\rangle$ and $|\psi\rangle$. Infidelity measures how “different” two states are and admits values in the range $[0, 1]$, with an infidelity of 0 implying that $U(\theta)^\dagger|\mathbf{0}\rangle = |\psi\rangle$.

Hence, the problem can be framed as an objective function as $\min_{\theta} (1 - f_{|\psi\rangle\langle\psi|, |\mathbf{0}\rangle\langle\mathbf{0}|}(\theta))$ (or equivalently, $\max_{\theta} f_{|\psi\rangle\langle\psi|, |\mathbf{0}\rangle\langle\mathbf{0}|}(\theta)$), with $|\mathbf{0}\rangle\langle\mathbf{0}|$ being the observable and $\theta \in \mathbb{R}^{\frac{n}{2} \times d \times p}$. However, since $|\mathbf{0}\rangle\langle\mathbf{0}|$ acts on all n -qubits, it does not fit into our framework. Motivated by Cerezo et al. (2021b), we take $\min_{\theta} (1 - f_{|\psi\rangle\langle\psi|, J}(\theta))$ as the objective function, where $J = \frac{1}{n} \sum_{i=1}^n |0\rangle_i\langle 0|$ and $|0\rangle_i\langle 0|$ is a 1-local operator that applies $|0\rangle\langle 0| \in \mathcal{L}(\mathbb{C}^2)$ on the i^{th} qubit only and $\mathbb{1}$ on all other qubits. In this case, the observable J is a sum of 1-local observables $O_i = \frac{1}{n} |0\rangle_i\langle 0|$. Hence, this problem fits in our framework and we can use ALSO to optimize it.

Quantum Autoencoder

Autoencoder is a popular dimensionality reduction technique in classical machine learning (Hinton and Zemel 1993). Using deep neural networks, autoencoders learn low dimensional representations of high dimensional input data, which should ideally keep hold of the original characteristics of the data. This can also be seen as a form of data compression. Recently, there have been numerous works on extending this concept to quantum data (Romero, Olson, and Aspuru-Guzik 2017; Wan et al. 2017; Verdon, Pye, and Broughton 2018; Pepper, Tischler, and Pryde 2019; Lamata et al. 2018). We focus on the version presented in Romero, Olson, and Aspuru-Guzik (2017), more specifically, its implementation using alternating layered ansatzes described in Cerezo et al. (2021b).

The idea behind this version of quantum autoencoder is to compress n -qubit quantum states into $n_B < n$ qubit states. Consider an ensemble of n -qubit states $\mathcal{Z} = \{(p_i, |\psi_i\rangle)\}$ with each state being prepared in registers A and B having n_A and n_B qubits respectively, where $n = n_A + n_B$. Let $\rho_{\mathcal{Z}} = \sum_i p_i |\psi_i\rangle\langle\psi_i|$. As a measure of the compression effect, we consider $1 - f_{\rho_{\mathcal{Z}}, \mathbb{1}_{[A]} \otimes J_{[B]}}(\theta)$, where $J = \frac{1}{n_B} \sum_{i=1}^{n_B} |0\rangle_i\langle 0|$ is defined, as in the state preparation problem, to be a sensible cost function that not only forces the population of all the states to be in the qubits in the register B , but also involves the same 1-local observables that we have used for state preparation. Again, this cost function has been used in Cerezo et al. (2021b).

Simulation Results

In this section, we discuss the experimental results comparing the performance of ALSO and the standard VQA in the two use cases discussed above.




Figure 5: Simulation results for state preparation.




Figure 6: Simulation results for quantum autoencoder.

Experiments Set-Up

For all experiments, each brick-like sub-circuit $S(\theta_{ij})$ (cf. Fig. 1) has the form given in Fig. 2. The simulation results presented in this section (except for Table 1) have used Simultaneous Perturbation Stochastic Approximation (Spall 1992) (SPSA), where the converging sequences used for state preparation and quantum autoencoder are, respectively, $c_r = a_r = r^{-0.5}$ and $c_r = a_r = r^{-0.3}$.

In the following, we denote by ALSO- T the ALSO algorithm that uses T shadows and by VQA- K the VQA algorithm that consumes per function evaluation K state copies (for state preparation) or K samples from $\mathcal{Z} = \{(p_i, |\psi_i\rangle)\}$ (for quantum autoencoder). In addition, we write VQA-infinite for the VQA algorithm which has access to an infinite number of state copies.

Let R be the total number of iterations of SPSA. Since SPSA requires 2 function evaluations per iteration, for state preparation and quantum autoencoder, VQA- K will consume $2KRn$ and $2KRn_B$ state copies respectively.

State Preparation Experiments

For the state preparation problem, we first consider the case when $n = 8$, i.e., the target state is an 8-qubit state. In each experiment, the target state is compatible with an alternating layered ansatz. We repeat the experiments for five different target states and our results are shown in Fig. 5(a,b), where each plot corresponds to five different instances of a problem. At any value on the x -axis, we plot the mean of infidelity/cost values across the five different experiments that

were carried out. The coloured area of the plot is marked on top and bottom by the mean plus and minus the standard deviation of the 5 values at each point respectively.

In Fig. 5(a), VQA-10 consumes $2KRn = 2 \times 10 \times 3000 \times 8 = 4.8 \times 10^5$ state copies and in a similar manner, the other VQA algorithms consume 2.4×10^6 and 4.8×10^6 state copies, which are 4.8x, 24x, and 48x of that ALSO consumes. Furthermore, from Fig. 5(b), we can see that ALSO closely matches the outcome of VQA-infinite with only 5×10^5 state copies.

Moving on from the 8-qubit scenario, we then carry out similar experiments for 30-qubit systems. Fig. 5(c) shows the results, where, as in Cerezo et al. (2021b), all states involved are computational basis states.

We note in this case, VQA-10 consumes $2 \times 10 \times 9000 \times 30 = 5.4 \times 10^6$ state copies while ALSO- T remains unchanged with the change of n from 8 to 30. From the figure, it is clear that the similar conclusion for 8-qubit state preparation also applies to 30-qubit state preparation. In particular, ALSO (with 5×10^5 samples) significantly outperforms VQA with 100x more samples.

Quantum Autoencoder Experiments

For quantum autoencoder, similar experiments are carried out for both 8- and 30-qubit systems. Ensembles containing two pure states $|\psi_1\rangle$ and $|\psi_2\rangle$ are chosen with $p_1 = 1/3$ and $p_2 = 2/3$, and n_B is set as 4 and 10, respectively, for 8- and 30-qubit systems. We repeat the experiments for five different ensembles. The results are summarised in Fig. 6,

which have the same explanation as those in Fig. 5. Note that the cost values plotted here are the actual cost $f_{\rho, O}(\theta)$ and not their estimations. From the figure, we can see similar conclusion we have obtained for state preparation also holds for quantum autoencoder. It seems that in this case ALSO with 10^5 samples significantly outperforms VQA with 48×10^5 samples and ALSO with 5×10^5 samples can often beat VQA-infinite!

Resource Consumption for the Same Objective

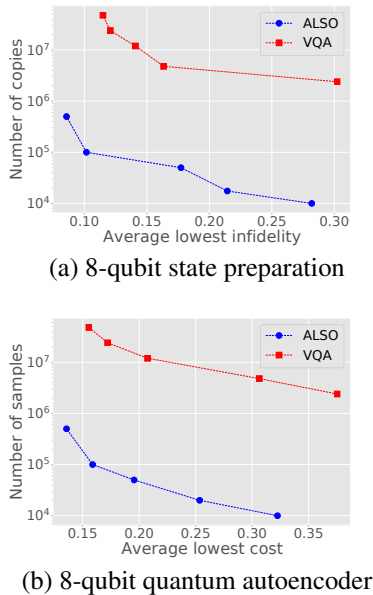


Figure 7: Resource requirement for the same objectives.

In above, we compared the performance of ALSO and VQA algorithms with predetermined resources. The efficiency of ALSO over VQA can also be illustrated by comparing the resource consumption required for the same objective. Given an objective, which can be either the average lowest infidelity or the average lowest cost, we carry out experiments to check how many state copies or samples are required for ALSO or VQA to achieve the objective. The results are presented in Fig. 7, where each point represents the average of five instances. It is clear that ALSO achieves a huge advantage in the number of state copies or samples that were required to achieve the specific levels of quality.

More Iterations by Using Powell’s Method

All the simulation results discussed above have used SPSA to find the optimal parameters. In each case, we set the maximum iterations to be the same for ALSO and VQA. From a practical point of view, this is unfair as ALSO can carry out more iterations with the given number of state copies or samples.

To further demonstrate this advantage of ALSO, we turn to Powell’s method (Powell 1964) to optimize the parameters. We carry out 8-qubit state preparation as well as quantum autoencoder optimizations and the results are presented

	state preparation		quantum autoencoder	
	#copies	infidelity	#samples	cost
VQA- 10^2	5×10^5	0.921	5.6×10^5	0.494
VQA- 10^3	1.3×10^7	0.348	4.6×10^6	0.408
VQA- 10^4	3.3×10^8	0.094	10^8	0.250
VQA- 10^5	4×10^9	0.069	2.1×10^9	0.188
ALSO	5×10^5	0.004	5×10^5	0.117

Table 1: Results under a different classical optimizer (Powell’s method).

in Table 1. In the infidelity (cost) columns, each entry is an average optimal infidelity (cost) of 5 instances of the problem, and we give an approximation of the average number of copies consumed (except for ALSO where exactly 5×10^5 copies are consumed) to achieve these values in the #copies (#samples) columns.

We set 5×10^4 as an upper limit on the total number of function evaluations for VQA. But, since ALSO does not consume any copies for more iterations, we don’t set any limit in the case of ALSO. As we can see, our approach greatly outperforms VQA in this case. Interestingly, in the case of VQA, the optimizers terminated in $5 \times 10^3 - 3 \times 10^4$ function evaluations in most cases. Only for the state preparation problem and with 10^5 copies consumed per function evaluation, we saw the optimizer exceeding the 5×10^4 limit. We also observe that VQA with Powell’s method performs very poorly compared to VQA with SPSA when K is small, which is possibly due to the inherent ability of SPSA to deal with noisy functions.

Conclusion and Future Direction

In this work, we proposed ALSO — an efficient method to train alternating layered VQAs that is exponentially better than the standard way of training VQAs in terms of the number of copies of input state consumed (or, in some other applications, number of executions of the quantum computer). The saving of state copies is especially useful when multiple rounds of the same optimization algorithm are required for various choices of hyperparameters, or when one has to experiment with different algorithms altogether. Moreover, ALSO is implementable using fewer and simpler quantum operations; in fact, only single-qubit measurements according to Pauli bases (in the classical shadow preparation stage) are required in ALSO. Another interesting benefit of the classical shadow technique, and so ALSO in particular, is that the produced classical shadows can be reused in different (independent) tasks. For example, the same set of classical shadows can be used in both finding the state preparation circuits and building quantum autoencoders.

In terms of future directions, we are trying to design similar resource efficient protocols for other trainable ansatzes such as the Quantum CNN ansatz (Cong, Choi, and Lukin 2019; Pesah et al. 2021). Another topic of interest would be to extend this method to techniques in other areas of quantum information that use large amounts of copies or executions such as device calibration and error correction.

Acknowledgements

We thank Afham and Richard Kueng for valuable discussions. This work was partially supported by the National Key R&D Program of China (Grant No.: 2018YFA0306704) and the Australian Research Council (Grant No.s: DP180100691, DP220102059). AB was partially funded by the Sydney Quantum Academy.

References

- Aaronson, S. 2018. Shadow Tomography of Quantum States. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018*, 325–338. New York, NY, USA: Association for Computing Machinery. ISBN 9781450355599.
- Arrasmith, A.; Cerezo, M.; Czarnik, P.; Cincio, L.; and Coles, P. J. 2021. Effect of barren plateaus on gradient-free optimization. *Quantum*, 5: 558.
- Boyd, G.; and Koczor, B. 2022. Training variational quantum circuits with CoVaR: covariance root finding with classical shadows. *arXiv:2204.08494*.
- Cerezo, M.; Arrasmith, A.; Babbush, R.; Benjamin, S. C.; Endo, S.; Fujii, K.; McClean, J. R.; Mitarai, K.; Yuan, X.; Cincio, L.; and Coles, P. J. 2021a. Variational quantum algorithms. *Nature Reviews Physics*, 3(9): 625–644.
- Cerezo, M.; Sone, A.; Volkoff, T.; Cincio, L.; and Coles, P. J. 2021b. Cost function dependent barren plateaus in shallow parametrized quantum circuits. *Nature Communications*, 12(1).
- Cong, I.; Choi, S.; and Lukin, M. D. 2019. Quantum convolutional neural networks. *Nature Physics*, 15(12): 1273–1278.
- Farhi, E.; Goldstone, J.; and Gutmann, S. 2014. A Quantum Approximate Optimization Algorithm. *arXiv:1411.4028*.
- Fontana, E.; Rungger, I.; Duncan, R.; and Cirstoiu, C. 2022. Efficient recovery of variational quantum algorithms landscapes using classical signal processing. *arXiv:2208.05958*.
- Gambetta, J. 2022. Quantum-centric supercomputing: The next wave of computing. *IBM Research Blog*.
- Hadfield, S.; Wang, Z.; O’Gorman, B.; Rieffel, E.; Venturelli, D.; and Biswas, R. 2019. From the Quantum Approximate Optimization Algorithm to a Quantum Alternating Operator Ansatz. *Algorithms*, 12(2): 34.
- Hinsche, M.; Ioannou, M.; Nietner, A.; Haferkamp, J.; Quek, Y.; Hangleiter, D.; Seifert, J.-P.; Eisert, J.; and Sweke, R. 2021. Learnability of the output distributions of local quantum circuits. *arXiv:2110.05517*.
- Hinton, G. E.; and Zemel, R. S. 1993. Autoencoders, Minimum Description Length and Helmholtz Free Energy. In *Proceedings of the 6th International Conference on Neural Information Processing Systems, NIPS’93*, 3–10. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.
- Huang, H.-Y.; Broughton, M.; Mohseni, M.; Babbush, R.; Boixo, S.; Neven, H.; and McClean, J. R. 2021a. Power of data in quantum machine learning. *Nature Communications*, 12(1).
- Huang, H.-Y.; Kueng, R.; and Preskill, J. 2020. Predicting many properties of a quantum system from very few measurements. *Nature Physics*, 16(10): 1050–1057.
- Huang, H.-Y.; Kueng, R.; Torlai, G.; Albert, V. V.; and Preskill, J. 2021b. Provably efficient machine learning for quantum many-body problems. *arXiv:2106.12627*.
- Kandala, A.; Mezzacapo, A.; Temme, K.; Takita, M.; Brink, M.; Chow, J. M.; and Gambetta, J. M. 2017. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature*, 549: 242–246.
- Lamata, L.; Alvarez-Rodriguez, U.; Martín-Guerrero, J. D.; Sanz, M.; and Solano, E. 2018. Quantum autoencoders via quantum adders with genetic algorithms. *Quantum Science and Technology*, 4(1): 014007.
- Li, G.; Song, Z.; and Wang, X. 2021. VSQ: Variational Shadow Quantum Learning for Classification. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(9): 8357–8365.
- McClean, J. R.; Boixo, S.; Smelyanskiy, V. N.; Babbush, R.; and Neven, H. 2018. Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9(1).
- Mitarai, K.; Negoro, M.; Kitagawa, M.; and Fujii, K. 2018. Quantum circuit learning. *Phys. Rev. A*, 98: 032309.
- Nakaji, K.; and Yamamoto, N. 2021. Expressibility of the alternating layered ansatz for quantum computation. *Quantum*, 5: 434.
- Okada, K. N.; Osaki, K.; Mitarai, K.; and Fujii, K. 2022. Identification of topological phases using classically-optimized variational quantum eigensolver. *arXiv:2202.02909*.
- Pepper, A.; Tischler, N.; and Pryde, G. J. 2019. Experimental Realization of a Quantum Autoencoder: The Compression of Qutrits via Machine Learning. *Phys. Rev. Lett.*, 122: 060501.
- Pesah, A.; Cerezo, M.; Wang, S.; Volkoff, T.; Sornborger, A. T.; and Coles, P. J. 2021. Absence of Barren Plateaus in Quantum Convolutional Neural Networks. *Physical Review X*, 11(4).
- Powell, M. J. D. 1964. An efficient method for finding the minimum of a function of several variables without calculating derivatives. *The Computer Journal*, 7(2): 155–162.
- Romero, J.; Olson, J. P.; and Aspuru-Guzik, A. 2017. Quantum autoencoders for efficient compression of quantum data. *Quantum Science and Technology*, 2(4): 045001.
- Sack, S. H.; Medina, R. A.; Michailidis, A. A.; Kueng, R.; and Serbyn, M. 2022. Avoiding Barren Plateaus Using Classical Shadows. *PRX Quantum*, 3(2).
- Schuld, M.; Sweke, R.; and Meyer, J. J. 2021. Effect of data encoding on the expressive power of variational quantum-machine-learning models. *Physical Review A*, 103(3).
- Slattery, L.; Villalonga, B.; and Clark, B. K. 2022. Unitary block optimization for variational quantum algorithms. *Phys. Rev. Research*, 4: 023072.

- Spall, J. 1992. Multivariate stochastic approximation using a simultaneous perturbation gradient approximation. *IEEE Transactions on Automatic Control*, 37(3): 332–341.
- Stokes, J.; Izaac, J.; Killoran, N.; and Carleo, G. 2020. Quantum Natural Gradient. *Quantum*, 4: 269.
- Verdon, G.; Pye, J.; and Broughton, M. 2018. A Universal Training Algorithm for Quantum Deep Learning. *arXiv:1806.09729*.
- Wan, K. H.; Dahlsten, O.; Kristjánsson, H.; Gardner, R.; and Kim, M. S. 2017. Quantum generalisation of feedforward neural networks. *npj Quantum Information*, 3(1).
- Wierichs, D.; Gogolin, C.; and Kastoryano, M. 2020. Avoiding local minima in variational quantum eigensolvers with the natural gradient optimizer. *Phys. Rev. Research*, 2: 043246.
- Wiersema, R.; Zhou, C.; de Sereville, Y.; Carrasquilla, J. F.; Kim, Y. B.; and Yuen, H. 2020. Exploring Entanglement and Optimization within the Hamiltonian Variational Ansatz. *PRX Quantum*, 1(2).
- Wu, A.; Li, G.; Wang, Y.; Feng, B.; Ding, Y.; and Xie, Y. 2021. Towards Efficient Ansatz Architecture for Variational Quantum Algorithms. *arXiv:2111.13730*.
- Yao, J.; Bukov, M.; and Lin, L. 2020. Policy Gradient based Quantum Approximate Optimization Algorithm. In Lu, J.; and Ward, R., eds., *Proceedings of The First Mathematical and Scientific Machine Learning Conference*, volume 107 of *Proceedings of Machine Learning Research*, 605–634. PMLR.