

PAPER • OPEN ACCESS

Natural evolutionary strategies for variational quantum computation

To cite this article: Abhinav Anand *et al* 2021 *Mach. Learn.: Sci. Technol.* **2** 045012

View the [article online](#) for updates and enhancements.

You may also like

- [A method of rotating accelerometer gravity gradiometer for centrifugal gradient detection](#)
Mingbiao Yu and Tijing Cai
- [A cradle-shaped gradient coil to expand the clear-bore width of an animal MRI scanner](#)
K M Gilbert, J S Gati, L M Klassen et al.
- [A Comparison of Multiphase Magnetic Field Tracers in a High Galactic Latitude Region of the Filamentary Interstellar Medium](#)
J. L. Campbell, S. E. Clark, B. M. Gaensler et al.



PAPER

OPEN ACCESS

Natural evolutionary strategies for variational quantum computation

RECEIVED
17 December 2020REVISED
24 February 2021ACCEPTED FOR PUBLICATION
30 March 2021PUBLISHED
19 July 2021Abhinav Anand^{1,*} , Matthias Degroote^{1,2}  and Alán Aspuru-Guzik^{1,2,3,4,*} ¹ Chemical Physics Theory Group, Department of Chemistry, University of Toronto, 80 St, George Street, Toronto, ON M5S 3H6 Canada² Department of Computer Science, University of Toronto, 40 St. George Street, Toronto, ON M5S 2E4, Canada³ CIFAR AI chair, Vector Institute, 661 University Ave. Suite 710, Toronto, ON M5G 1M1, Canada⁴ Canadian Institute for Advanced Research (CIFAR) Lebovic Fellow, 661 University Ave, Toronto, ON M5G 1M1, Canada

* Author to whom any correspondence should be addressed.

E-mail: abhinav.anand@mail.utoronto.ca and aspuru@utoronto.ca**Keywords:** natural evolutionary strategies, optimization, variational quantum algorithmsOriginal Content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

**Abstract**

Natural evolutionary strategies (NES) are a family of gradient-free black-box optimization algorithms. This study illustrates their use for the optimization of randomly initialized parameterized quantum circuits (PQCs) in the region of vanishing gradients. We show that using the NES gradient estimator the exponential decrease in variance can be alleviated. We implement two specific approaches, the exponential and separable NES, for parameter optimization of PQCs and compare them against standard gradient descent. We apply them to two different problems of ground state energy estimation using variational quantum eigensolver and state preparation with circuits of varying depth and length. We also introduce batch optimization for circuits with larger depth to extend the use of ES to a larger number of parameters. We achieve accuracy comparable to state-of-the-art optimization techniques in all the above cases with a lower number of circuit evaluations. Our empirical results indicate that one can use NES as a hybrid tool in tandem with other gradient-based methods for optimization of deep quantum circuits in regions with vanishing gradients.

1. Introduction

Rapid developments in quantum hardware [1–3] have led to the development of a wide variety of algorithms to perform useful tasks on the available devices. The limited coherence times and number of qubits of the devices places limitations on the circuits that can be reliably executed. A class of algorithms which are well-suited for these noisy intermediate scale quantum [4] devices are the variational quantum algorithms (VQAs) [5]. These algorithms use classical and quantum computers in tandem, using the quantum computers for state preparation with parametrized unitaries and measurement, and then optimizing the parameters of the unitaries on a classical computer in an iterative manner.

VQAs have been used for a variety of tasks such as simulating molecules and many body systems [6–8], compression of quantum states [9–11], classification tasks [12–14], generative modelling [15–20], and solving combinatorial optimization problems [21] among others [22–24]. However, these algorithms suffer in cases when there is little or no knowledge about a good initial point. For such problems, the optimization of the unitaries using gradient-based methods become impossible for deep quantum circuits, as the average and variance of the gradient of the objective function decay exponentially as a function of the number of qubits and layer in the circuits [25]. More recently it was reported [26–28] that the barren plateaus are not only present in deep circuits, but are dependent on the cost function used for optimization and that optimization of problems with global cost function is not possible for randomly parameterized shallow circuits either. Furthermore, another study [29] showed that the variance of the gradient is bounded by the width of the causal cone of the circuit corresponding to the various Pauli terms in the decomposition of the cost function. In another recent study [30], the authors investigated the use of gradient-free optimizers in a

barren plateau region of the landscape, and show that it becomes increasingly difficult to optimize as the number of shots required grows exponentially.

Recently, there have been different studies that have tried to solve the barren plateau problem. One approach circumvents the problem by formulating better initial guesses for the values of the parameters [31]. A different study [32] investigated ways to boost the gradients for certain problems by using correlated parameters. In this study we use an approach based on natural evolution strategies (NES) and batch optimization for avoiding barren plateaus in the optimization landscapes.

Evolutionary strategies (ES) [33, 34] have been applied extensively in classical machine learning tasks as a viable black-box optimization tool for high-dimensional problems. These strategies are promising for optimization of quantum circuits with respect to both the number of function evaluations and the scaling with the size of the circuit. The gradient is estimated by a number of function evaluations that does not increase with the number of parameters. Moreover, these function evaluations are fully independent and could be executed in parallel. This means that these methods are suitable for high-dimensional problems. The optimization of parametrized quantum circuits in quantum computing is an example of such high-dimensional problem. The application of these optimization strategies to VQAs and the adaptation to the specific challenges arise in this context pose an interesting research question. They can be further improved by using estimates of natural search gradients, fitness shaping, and adaptive sampling [35, 36] for both highly correlated (considering the full co-variance matrix) and separable (only considering a diagonal co-variance matrix) problems. In a recent work [37], the authors illustrated the use of NES for approximate combinatorial optimization problems, and show that they can achieve state-of-the-art performance. More recently in a work [38] NES was used to investigate the problem of manipulating Majorana bound-states in a topological superconductor, where the authors show that these strategies can be used efficiently to study quantum mechanical many-body systems. In this study we use two modified strategies, separable NES (sNES) and exponential NES (xNES) [36, 39], for optimizing randomly initialized parameterized quantum circuits (PQCs) by considering circuit parameters as multivariate normal distributions. We show that optimization of randomly initialized parameterised quantum circuits can be performed with significantly less circuit evaluations using NES, while achieving comparable accuracy to gradient-based methods. We also show that they can be used to amplify gradients and improve parameter initialization for gradient-based approaches.

The rest of the paper is organised as follows, we provide the theory behind NES in section 2, the details of the numerical experiments in section 3, the result from optimization of randomly parameterized quantum circuits (RPQCs) in section 4, and finally present some concluding remarks in section 5.

2. Natural evolution strategies

Evolution strategies were first introduced by Rechenberg [33] and Schwefel [34], to cope with high-dimensional continuous-valued domains. It has remained an active field of research, and has been developed extensively over the years to include the full covariance matrix [40] for efficient optimization. In what follows, we describe a specific subclass of evolution strategies called the NES for iteratively updating a search distribution by ascending the surface along estimated search gradients. The reader is redirected to [35, 39] for a detailed description of NES.

2.1. Canonical natural evolution strategies

The NES uses search gradients to iteratively update the distribution of a stochastic variable \mathbf{z} . The search gradient can be defined for any distribution for which the derivatives of the log-density can be calculated. If we consider a distribution with density $\pi(\mathbf{z}|\theta)$ with θ as the parameter and $f(\mathbf{z})$ as the fitness for \mathbf{z} sampled from the distribution, then we can write the expected fitness as

$$J(\theta) = E_{\theta}[f(\mathbf{z})] = \int f(\mathbf{z})\pi(\mathbf{z}|\theta)d\mathbf{z}. \quad (1)$$

Using the ‘log-likelihood trick’ we can then approximate the search gradient from samples $\mathbf{z}_1 \dots \mathbf{z}_k$ as

$$\nabla_{\theta} J(\theta) \approx \frac{1}{k} \sum_{n=1}^k f(\mathbf{z}_n) \nabla_{\theta} \log \pi(\mathbf{z}_n|\theta) \quad (2)$$

where k is the population size. This equation has the benefit that it does not require derivatives of the fitness function. One can estimate the gradient for different commonly used search distributions including

Gaussian and Cauchy distributions, and use standard gradient descent to iteratively update the parameters of the distribution:

$$\theta = \theta + \eta \nabla_{\theta} J(\theta), \quad (3)$$

where η is the learning rate.

For a Gaussian prior distribution on the perturbation [41], the procedure is shown in algorithm 1:

Algorithm 1. Canonical ES.

```

input:  $f, \mu_{init}, \sigma_{init}$ 
while stopping criterion is not met do
  for  $n = 1..k$  do
    draw sample  $s_n \sim N(0, I)$ 
     $z_n = \mu + \sigma_{init} s_n$ 
    evaluate the fitness  $f(z_n)$ 
  end
  compute gradients:
   $\nabla J = \frac{1}{k} \sum_{n=1}^k f(z_n) \cdot s_n$ 
  update parameters:
   $\mu = \mu + \eta \cdot \frac{1}{\sigma_{init}} \nabla J$ 
end

```

where f is the fitness function, η is the learning rate, σ_{init} is the fixed variance, s_n is a Gaussian perturbation sampled from the normal distribution $N(0, I)$. This procedure replaces the normal gradient with the evaluation of the objective function blurred with Gaussian noise and samples the stochastic gradient with the perturbed parameters. It reduces to the simultaneous perturbation stochastic approximation (SPSA) [42] for the case of a symmetric perturbation forward and backward for a single Gaussian perturbation. As this method does not follow the exact gradient of the cost function but a stochastic estimator, it does not necessarily get stuck on the same barren plateaus as the analytical gradient. We will show in the numerical section that this is the case in all the simulations we performed.

Instead of using the standard gradients for updates, modern ES implementations use the natural gradients because of their numerous advantages [43, 44]. The natural gradient can be calculated using the Fisher matrix F which can be estimated from the search gradient as

$$F \approx \frac{1}{k} \sum_{n=1}^k \nabla_{\theta} \log \pi(z_n | \theta) \nabla_{\theta} \log \pi(z_n | \theta)^T, \quad (4)$$

and the distribution is now updated as:

$$\theta = \theta + \eta \cdot F^{-1} \nabla_{\theta} J(\theta). \quad (5)$$

The canonical NES can be further improved by fitness shaping and adaptive sampling, which are used in xNES and sNES [35, 36, 39].

2.2. Exponential natural evolution strategies

The multivariate Gaussian distribution is one of the most prominent search distribution used for ES. Historically this distribution has received the most attention because of its analytical properties and the resulting elegance of the equations. Based on these techniques a new algorithm known as xNES can be formulated, for which we provide the outline in algorithm 2, as presented in [36, 39]. It goes beyond the standard NES by introducing an exponential parametrization by which it makes the natural gradient easier to compute:

Algorithm 2. xNES.

```

input:  $f, \boldsymbol{\mu}_{init}, \boldsymbol{\Sigma}_{init} = \mathbf{A}\mathbf{A}^T$ 
initialize:
 $\sigma = \sqrt[d]{|\det(\mathbf{A})|}$ 
 $\mathbf{B} = \mathbf{A}/\sigma$ 
while stopping criterion is not met do
  for  $n = 1..k$  do
    draw sample  $\mathbf{s}_n \sim N(0, \mathbf{I})$ 
     $\mathbf{z}_n = \boldsymbol{\mu} + \sigma\mathbf{B}^T\mathbf{s}_n$ 
    evaluate the fitness  $f(\mathbf{z}_n)$ 
  end
  sort  $\{(\mathbf{s}_n, \mathbf{z}_n)\}$  with respect to  $f(\mathbf{z}_n)$  and compute utilities  $u_n$ 
  compute gradients:
   $\nabla_{\boldsymbol{\mu}}J = \sum_{n=1}^k u_n \cdot \mathbf{s}_n$ 
   $\nabla_{\mathbf{M}}J = \sum_{n=1}^k u_n \cdot (\mathbf{s}_n\mathbf{s}_n^T - \mathbf{I})$ 
   $\nabla_{\sigma}J = \text{tr}(\nabla_{\mathbf{M}}J)/d$ 
   $\nabla_{\mathbf{B}}J = \nabla_{\mathbf{M}}J - \nabla_{\sigma}J \cdot \mathbf{I}$ 
  update parameters:
   $\boldsymbol{\mu} = \boldsymbol{\mu} + \eta_{\boldsymbol{\mu}} \cdot \sigma\mathbf{B} \cdot \nabla_{\boldsymbol{\mu}}J$ 
   $\sigma = \sigma \cdot \exp(\eta_{\sigma}/2 \cdot \nabla_{\sigma}J)$ 
   $\mathbf{B} = \mathbf{B} \cdot \exp(\eta_{\mathbf{B}}/2 \cdot \nabla_{\mathbf{B}}J)$ 
end

```

where f is the fitness function, d is the dimension of the problem, \mathbf{A} is the covariance matrix, $\boldsymbol{\mu}$ is the mean, k is the number of walkers, \mathbf{s}_n is the sample in local coordinates, \mathbf{z}_n is the sample in the task coordinates, $N(0, \mathbf{I})$ is a normal distribution, $\eta_{\mathbf{B}}$ is the covariance learning rate, $\eta_{\boldsymbol{\mu}}$ is the mean learning rate, η_{σ} is the scale learning rate, and u_n is the utility function used for fitness shaping which is defined as

$$u_n = \frac{\max(0, \log(\frac{k}{2} + 1) - \log(n))}{\sum_{j=1}^k \max(0, \log(\frac{k}{2} + 1) - \log(j))} - \frac{1}{k}. \quad (6)$$

The stopping criterion for the xNES algorithm is met when the maximum value of the covariance matrix reaches below a threshold of 10^{-8} . The xNES strategy works well across a wide variety of optimization problems [39], however the computational cost increases significantly with increasing variable dimension.

2.3. Separable natural evolution strategies

For high-dimensional problems, the more suitable variant is the sNES. Here, the correlation between the distribution variables is not considered (the covariance matrix \mathbf{A} is forced to be diagonal), implying that the parameter distributions of the PQCs are considered independent of each other. The sNES algorithm is expected to perform at least as well as xNES on separable problems, to perform worse than xNES for highly correlated problems and to outperform xNES for very high-dimensional problems [39]. The sNES algorithm is outlined in algorithm 3, as in [39]:

Algorithm 3. sNES.

```

input:  $f, \boldsymbol{\mu}_{init}, \boldsymbol{\sigma}_{init}$ 
while stopping criterion is not met do
  for  $n = 1..k$  do
    draw sample  $\mathbf{s}_n \sim N(0, \mathbf{I})$ 
     $\mathbf{z}_n = \boldsymbol{\mu} + \boldsymbol{\sigma}\mathbf{s}_n$ 
    evaluate the fitness  $f(\mathbf{z}_n)$ 
  end
  sort  $\{(\mathbf{s}_n, \mathbf{z}_n)\}$  with respect to  $f(\mathbf{z}_n)$  and compute utilities  $u_n$ 
  compute gradients:
   $\nabla_{\boldsymbol{\mu}}J = \sum_{n=1}^k u_n \cdot \mathbf{s}_n$ 
   $\nabla_{\boldsymbol{\sigma}}J = \sum_{n=1}^k u_n \cdot (\mathbf{s}_n^2 - 1)$ 
  update parameters:
   $\boldsymbol{\mu} = \boldsymbol{\mu} + \eta_{\boldsymbol{\mu}} \cdot \boldsymbol{\sigma} \cdot \nabla_{\boldsymbol{\mu}}J$ 
   $\boldsymbol{\sigma} = \boldsymbol{\sigma} \cdot \exp(\eta_{\boldsymbol{\sigma}}/2 \cdot \nabla_{\boldsymbol{\sigma}}J)$ 
end

```

where f is the fitness function, μ is the mean, σ is the standard deviation, k is the number of walkers, s_n is the sample in local coordinates, z_n is the sample in the task coordinates, $N(0, I)$ is a normal distribution, η_μ is the mean learning rate, η_σ is the deviation learning rate, and u_n is the utility function used for fitness shaping as defined in equation (6). The stopping criterion for the sNES algorithm is met when the maximum value of the standard deviation vector reaches below a threshold of 10^{-8} .

2.4. NES for quantum algorithms

A VQA generally consists of finding the optimal parameters θ^* of the PQC $U(\theta)$ to minimize the associated loss function. The traditional way of finding the optimal parameters include gradient-based (free) descent towards the surface minima. We consider an approach where the parameters of the circuit are sampled from a distribution, which is updated over time to achieve the global minima. This fits well within the NES framework, where one uses a search distribution to sample the parameters and update it iteratively to reach the optimum solution. The NES framework has also been shown to perform efficiently for very high-dimensional classical problems, thus can be an important tool for optimization of quantum algorithms as they deal with a high number of parameters and challenging optimization surfaces.

We carry the optimization of PQCs using NES, by considering the parameters θ of the PQCs as centres of a search distribution $\pi(z|\theta)$. This search distribution is then used in NES to sample different points z and estimate the gradient on the distribution parameters. The parameters of the search distribution θ are then optimized over time by using the estimated gradient to find the optimal parameters, θ^* for the PQC. We use two specific flavours, xNES and sNES, as illustrated in algorithms 2 and 3 for the optimization as they most naturally fit the requirements for the optimization of PQCs.

3. Application to quantum circuits

In this section, we illustrate the use of the NES for optimization of VQAs. We consider two different ansatz for the problem of state preparation. The first ansatz is a RPQC with repeating parameterized and entangling layers, as

$$U(\theta) = U(\theta_1 \dots \theta_L) = \prod_{l=1}^L V(\theta_l) W_l, \quad (7)$$

where L is the number of layers in the circuit, $V(\theta_l) = \bigotimes_{j=1}^Q R_{ij}(\theta_{ij})$ is a tensor product of randomly chosen pauli rotation gates R_{ij} , Q is the number of qubits in the circuit and W_l is a layer of entangling CZ gates between adjacent qubits. The second ansatz is an alternate layered parameterized quantum circuit (ALPQC) which consists of two layers of CZ gates between alternate nearest neighbour qubits with layers of R_Y rotation gates between the entangling layers:

$$U(\theta) = U(\theta_1 \dots \theta_L) = \prod_{l=1}^L V(\theta_l) W_l V'(\theta_l) W_l', \quad (8)$$

where L is the number of layers in the circuit, $V(\theta_n) = \bigotimes_{j=1}^{Q-1} R_Y(\theta)$ and $V'(\theta_l) = \bigotimes_{j=2}^Q R_Y(\theta)$ represents the layer of the rotation gates, and W_l and W_l' represent layers of entangling CZ gates between adjacent qubits.

The circuits are rotated out of the computational basis by adding a layer of $R_Y(\pi/4)$ gates, and are shown in figure 1. The objective Hamiltonian operator H is chosen to be a projector $|00\dots 00\rangle\langle 00\dots 00|$ on the vacuum state, and the fitness function is chosen as $f(\theta) = (1 - \langle \psi | H | \psi \rangle)^2$, where $|\psi\rangle = U(\theta) |00\dots 00\rangle$ is defined as the state after the application of the parameterized unitary on the vacuum state.

The parameters of the quantum circuits are used to define a search distribution—a multi-normal Gaussian with centres as the circuit parameters, which is then optimized using NES as discussed in section 2. For the xNES parameter distributions, the parameter distributions are correlated through the covariance matrix, while for the sNES the distributions are independent. The estimates of the fitness function are carried out by calculating the average outcome of the circuit at k different task coordinate values. The number k of these evaluations, which we will reference as walkers from here on, is independent of the number of parameters and qubits in the circuit and can be fixed to a constant for a given problem. We perform different numerical simulations to show the application of NES for estimating the ground state energy of the water molecule in the STO-3G basis using a unitary coupled cluster (UCC) ansatz with randomly initialized parameters. We also consider a more abstract example of state preparation. In both cases we focus on the initial part of the optimization landscape to demonstrate the ability to optimize in cases with randomly initialized parameters.

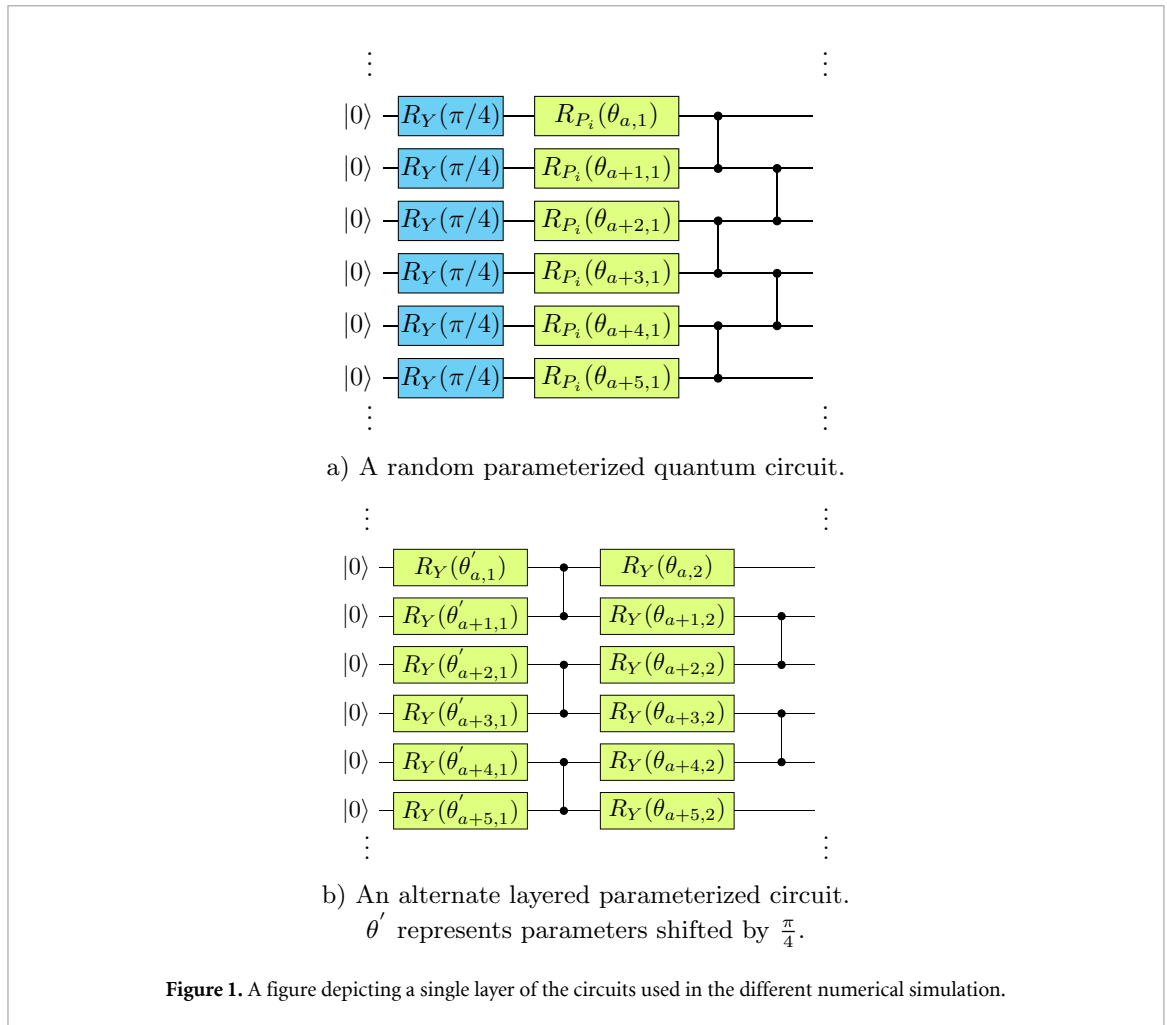


Figure 1. A figure depicting a single layer of the circuits used in the different numerical simulation.

Table 1. A table containing the values of all the parameters employed for NES strategies used for optimization. d represents the dimension of the batch.

Parameters	Values
η_μ	1
$\eta_s = \eta_B$	$\frac{(9+3\log(d))}{5d\sqrt{d}}$
η_σ	$\frac{(3+\log(d))}{5d\sqrt{d}}$

4. Numerical simulations

The circuit parameters are initialized as a multinormal distribution with randomly chosen centres and standard deviations of 0.1. The parameters used for the optimization are given in table 1 are the same as considered in [39]. These values for the different learning rates have been chosen based on the different investigations carried out in the classical machine learning literature to select the optimal value of these hyper-parameters for problems of varying dimensions [39]. The value for the number of walkers k , has been set to 16 in most simulation after analysing the effect of choosing a particular number is investigated in section 4.4. However, the value that is used widely in the machine learning community [39] based on the problem dimension is $(4 + 3\log d)$. The whole framework is implemented using TEQUILA [45] an open source package, which uses QULACS [46] as the backend for the execution of all the numerical simulations.

4.1. Search direction for NES

The problem for randomly initialized unstructured quantum circuits is that both the mean and variance tend to zero for random initializations [25]. This implies that all directions to search for the optimal parameters are equally likely, thus the difficulty in optimization. Adding random noise can not change the average value of the gradient but we want to demonstrate that we have a valid search direction in the optimization landscape. The search direction is defined by an estimate of the gradient of the search distribution on its

parameters, and we show that this estimate does not have a vanishing variance. For this we have repeated the calculations in [25] for the circuit in figure 1(a) for 18 qubits with a local cost function. We average each point over 10 000 random initializations. For each random initialization, we have calculated the estimate of the search gradient with respect to the first parameter with a varying number of walkers k from 1 to 8. Figure 2 shows results of the calculations for the mean of the surrogate search gradients obtained with σ_{init} taking values $\pi/8, \pi/16, \pi/32$.

The plots show that variance can be amplified arbitrarily as the width of the Gaussian noise is reduced. This should not be surprising as this parameter appears in the denominator in algorithm 1. Decreasing the width of the noise also means that the NES samples less far around the current parameters and less information of the optimization surface is included. As the number of walkers k is increased, the stochastic approximation of the gradient approaches the true value of the gradient better and the variance is reduced. Both the width of the noise and the number of walkers are important and need to be chosen appropriately. In what follows we will let the width of the noise be a variable to optimize and the number of walkers a hyperparameter of our algorithm.

This section demonstrates that we can get an arbitrary amplification of the gradient for randomly initialized quantum circuits by replacing the exact gradient with a stochastic surrogate. That does not necessarily mean that optimization would be possible following this gradient. It is possible that our estimator leads to finding local minima or poor convergence behaviour. In the next sections we show that optimization following this gradient is possible for a wide array of problems.

4.2. State preparation

We carry out numerical simulations for optimization of different quantum circuits with varying number of qubits and layers to prepare the corresponding $|00\dots00\rangle$ state. The number of walkers for all the simulations is chosen to be 16 unless stated otherwise. We plot the results from the simulation using NES and the two different ansätze RPQCs and ALPQCs with different layers in figure 3.

It can be seen from the plots in figure 3 that the optimization is possible in all cases when using sNES, but xNES fails to find the global optimum in cases with large number of parameters. This finding agrees with the expected behaviour of sNES for these high-dimensional and multi-modal problems with highly redundant global optima, as pointed out for various cases earlier [39]. The failure of xNES for high-dimensional problems suggests that it is not advantageous to retain the full covariance matrix for optimization and that including the correlations between these parameters leads to bad performance. Thus, instead of considering the full covariance matrix for optimization of deep quantum circuits, a viable strategy could be using a block matrix correlating only some parameters, which we employ in section 4.4. It is evident, that the rate of convergence of the optimization slows down with increasing number of parameters for both techniques which is expected as the search space becomes very large.

In figure 4, we compare the performance of NES with that of gradient descent for the optimization of RPQCs. The plots show that optimization with sNES performs equally well for all the cases in terms of finding the global minimum, but that xNES performance gets worse for circuits with more layers. An important point to make here is that, an update of the parameters using NES only requires k circuit evaluations, while a standard gradient descent calculation with the parameter shift procedure typically needs twice the number of parameters ($2QL$) function evaluations. Even though gradient descent reaches the global minimum in less iterations, it does so at a much higher cost in function evaluations. An interesting observation from the plots is that the initial rate of convergence with NES for circuits with 10 qubits or more is higher than that using gradient descent in many of the simulations, but similar on average. This is important in the context of the barren plateaus, where typically no suitable search direction can be found.

It can be also seen from the plots that the NES work very similar for the different numerical simulations. However the variation in the loss values in simulations using gradient descent is very large indicating the need for a better strategies in the initial part of the optimization. It opens the possibility for hybrid strategies, where the NES can be used to provide the initial direction for gradient descent for optimizing large RPQCs.

4.3. VQE for the H₂O molecule

In the context of chemistry, PQCs are commonly used to represent the wave function of molecules. Although the circuits here typically use structured physically motivated ansätze, the number of parameters becomes large as the system size grows. The motivation behind choosing this optimization problem is to investigate the applicability of NES to problems with physically inspired ansätze, and compare their performance with gradient-based methods for such problems. In this simulation we estimate the ground state energy of a water molecule at equilibrium geometry using the STO-3G basis (14-qubit) with the UCCSD ansatz (34 parameters). The parameters of the circuit are initialized randomly, which in certain cases lead to convergence issues due to presence of barren plateaus or getting stuck in a local minima as observed in [47].

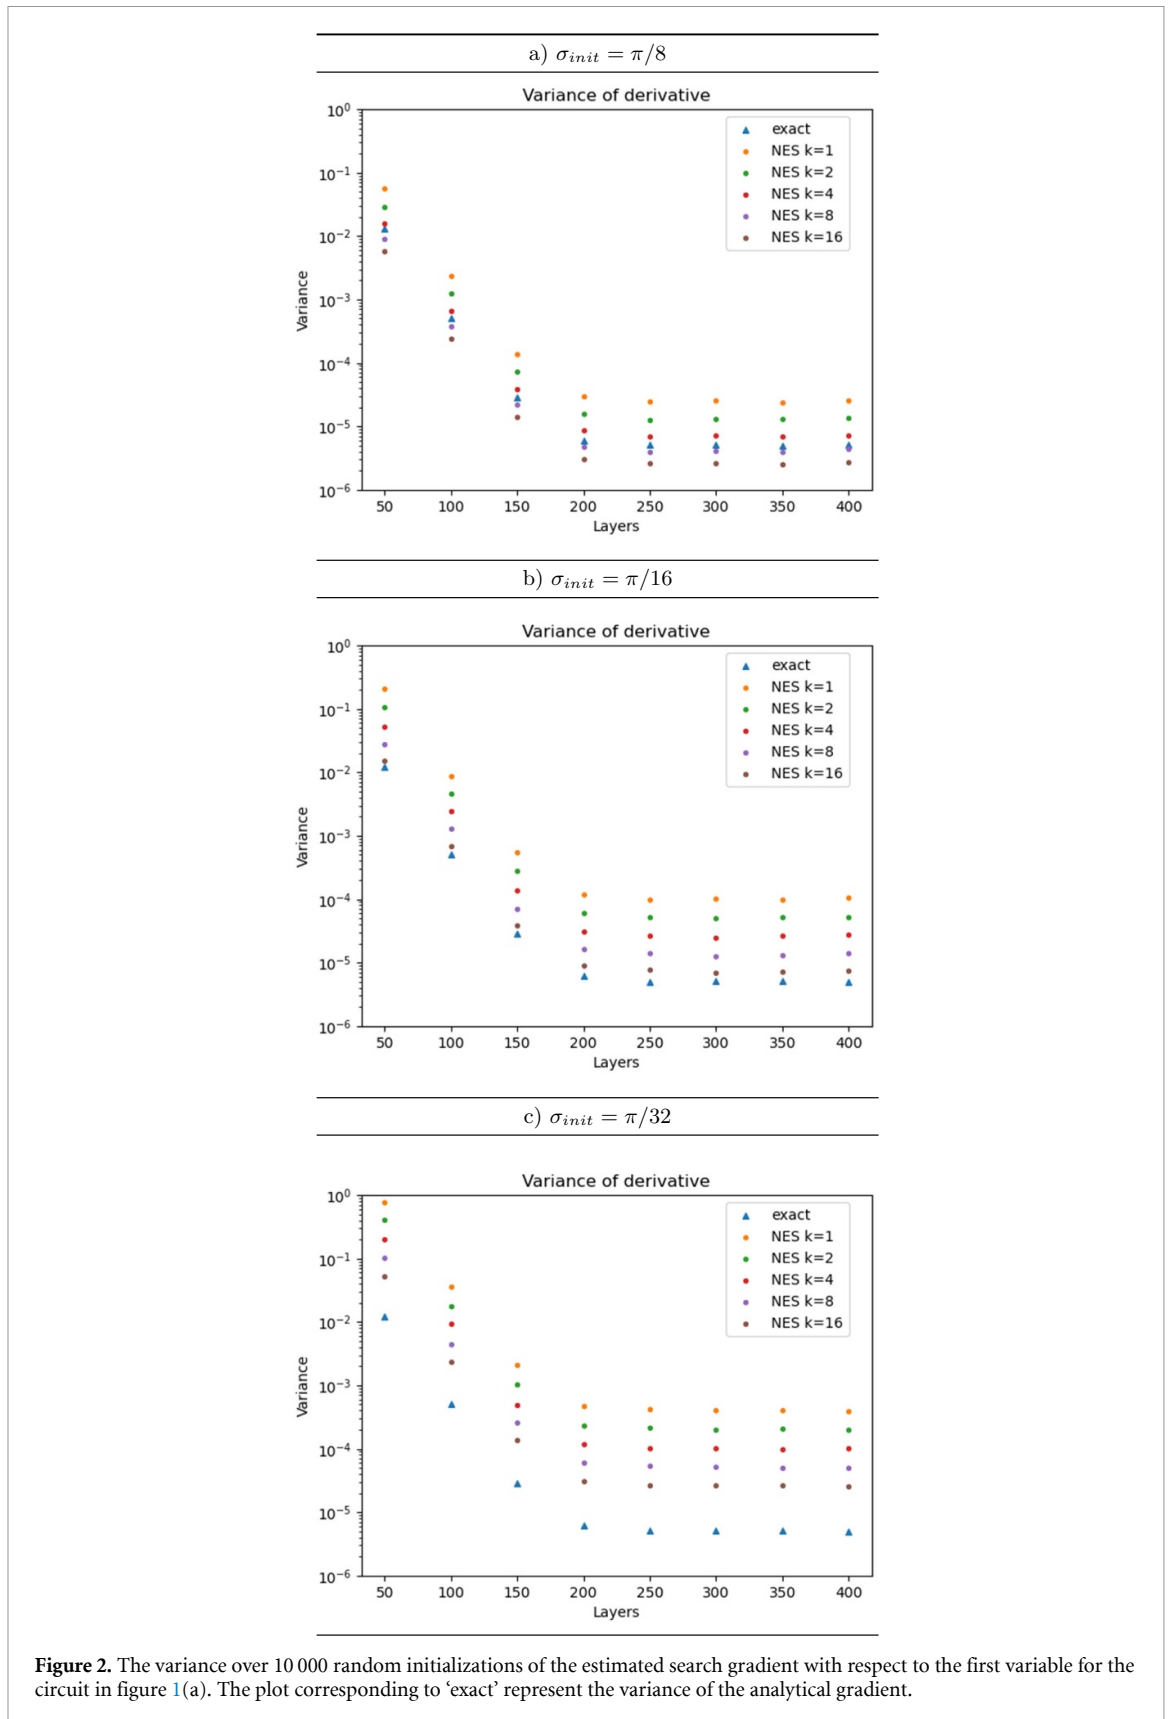


Figure 2. The variance over 10 000 random initializations of the estimated search gradient with respect to the first variable for the circuit in figure 1(a). The plot corresponding to 'exact' represent the variance of the analytical gradient.

We plot the results from the numerical experiments in figure 5, and it can be seen from the plots that we can use both sNES and xNES to optimize the parameters for estimating the ground state energy of water molecule using variational quantum eigensolver (VQE). We also carry out the optimisation with a gradient-based optimizer, and plot the results in figure 5. It can be seen from the plots that the performances of sNES, xNES and gradient descent are pretty similar. However, the xNES performs equally well for this

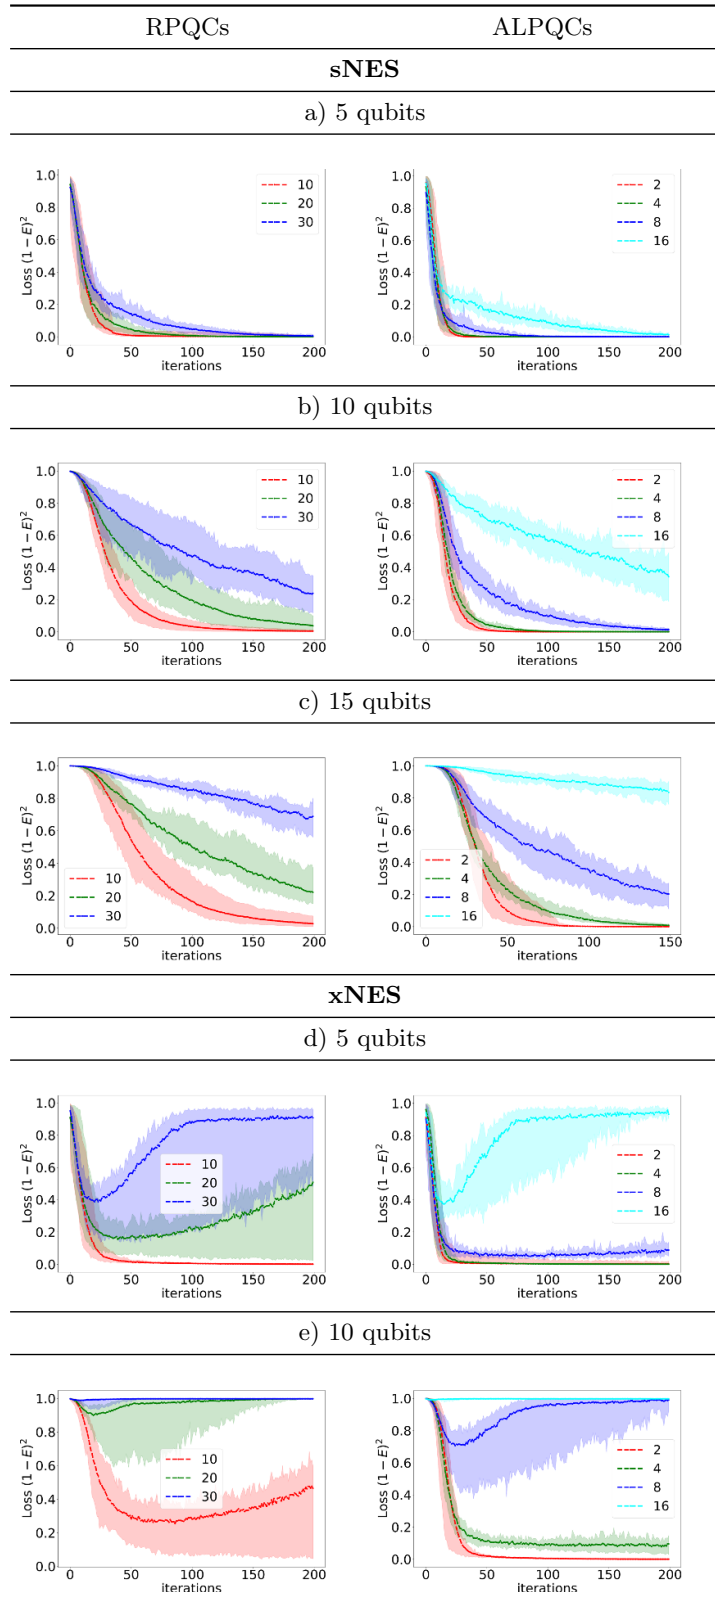
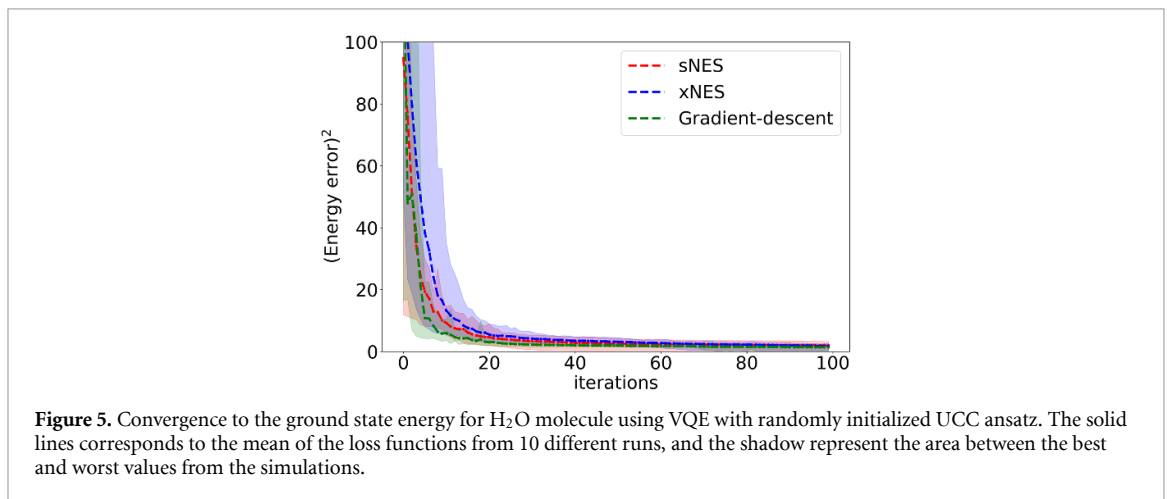
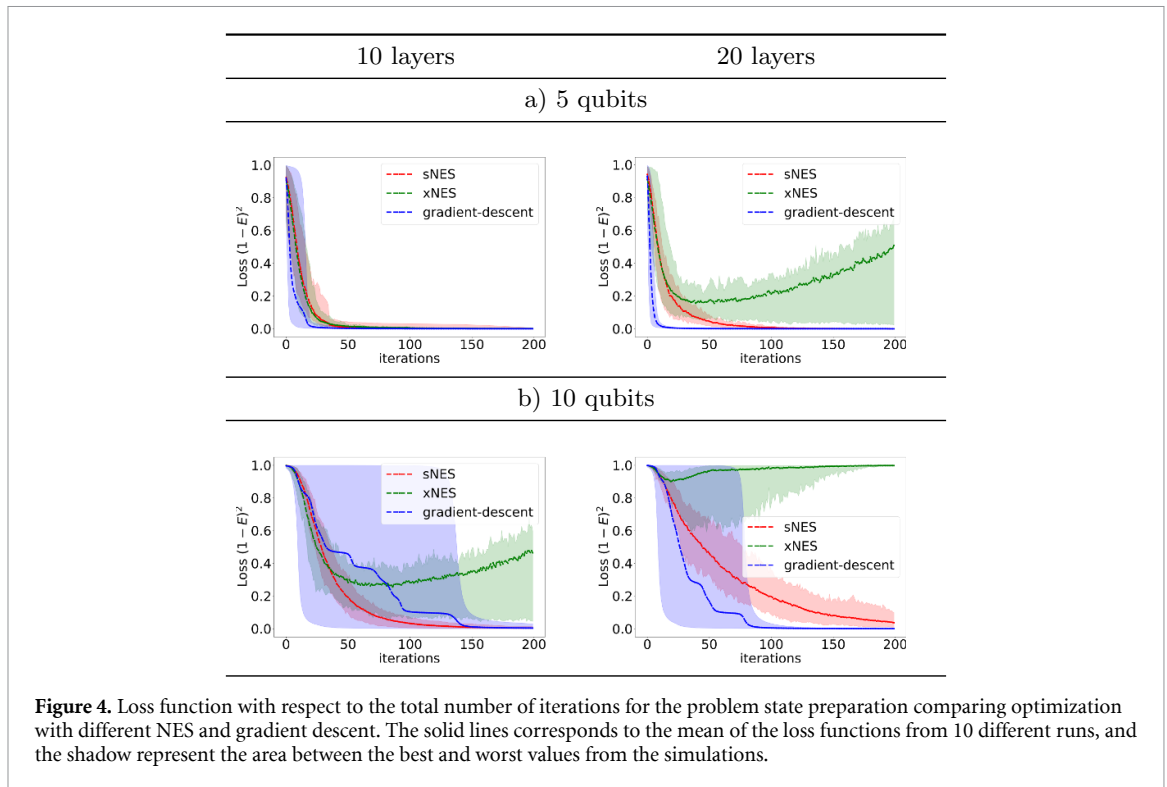


Figure 3. Loss functions with respect to the total number of iterations for different state preparation problems of varying circuit size. The solid lines corresponds to the mean of the loss functions from 10 different runs, and the shadow represent the area between the best and worst values from the simulations. The colours in the plots correspond to different number of layers used in the parameterized quantum circuits for the problem.



example as sNES, even though the size of the water molecule is beyond the limit for which we previously found that xNES loses performance compared to sNES. This suggests that in this case keeping the covariance matrix gives an advantage and that the parameters for this problem are actually correlated. The degree of randomness in the circuit topology can be a factor in selecting the optimal NES flavour.

4.4. Batch optimization

As seen in figure 3 the rate of convergence decreases with increasing number of parameters. This leads to long optimization runs for growing circuits. At the same time the overhead of the full NES strategy grows. To alleviate these effects, we investigate the use of batch optimization strategy for optimizing RPQCs of larger depth, as investigated in another recent study [48] where the authors investigate the use of layerwise learning techniques for optimisation of quantum neural networks. In batch optimization, only a limited number (the batch size) of parameters are adjusted every iteration to lower the cost. We carry out various simulations to analyse the effect of the size of the batches, number of walkers and different partitioning strategies on the convergence of the optimization.

We first investigate the effect of the size of batches and the number of walkers on the optimization by carrying out different numerical simulations with varying batch size. We use RPQCs with 10 qubits and 50 layers for the simulations and the batches of parameters are selected randomly but fixed over the entire

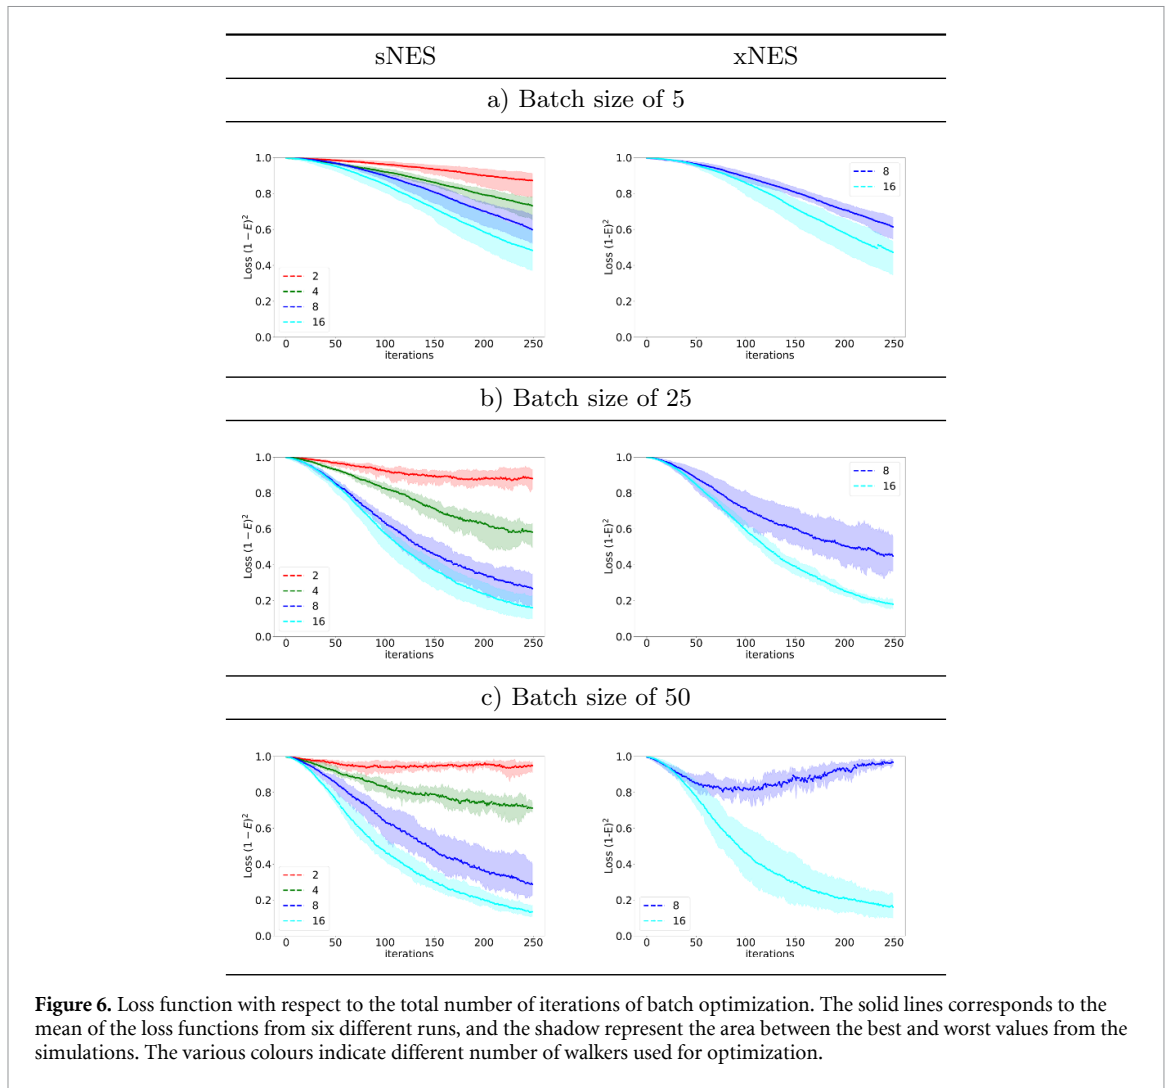


Figure 6. Loss function with respect to the total number of iterations of batch optimization. The solid lines corresponds to the mean of the loss functions from six different runs, and the shadow represent the area between the best and worst values from the simulations. The various colours indicate different number of walkers used for optimization.

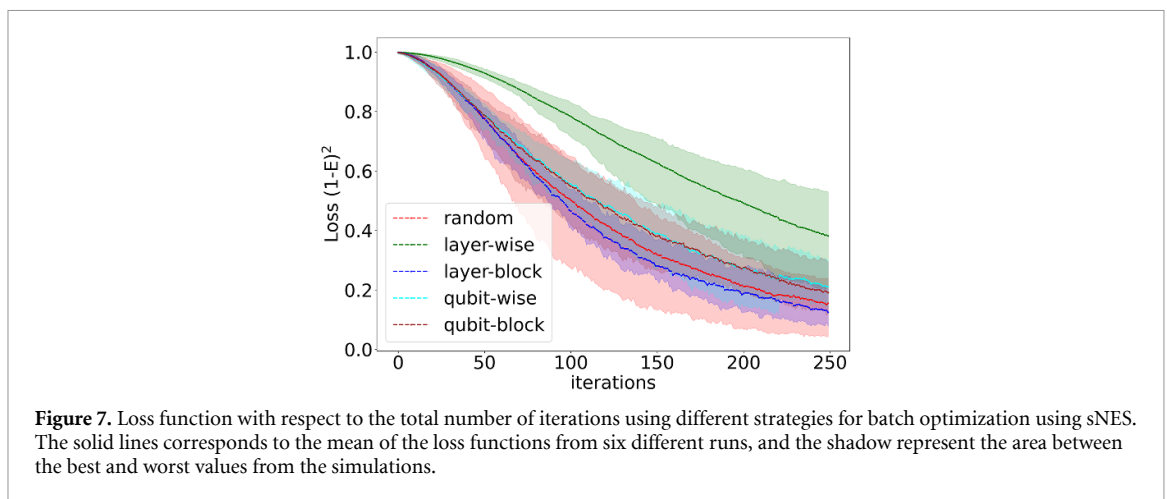


Figure 7. Loss function with respect to the total number of iterations using different strategies for batch optimization using sNES. The solid lines corresponds to the mean of the loss functions from six different runs, and the shadow represent the area between the best and worst values from the simulations.

calculation. The results in figure 6 show clearly that a minimum batch size is required. The xNES needs more walkers to give meaningful results than the sNES. The curves show both a convergence in the batch size and the number of walkers at around 16 walkers and a batch size of 50. We will use these values in the following calculations, it is expected that these values have some problem dependence and need to be tuned accordingly.

We next carry out simulations to investigate the effect of different strategies for generating the batches. The different strategies we investigate are dividing the circuit parameters randomly, layer-wise, qubit-wise, layer-block-wise, and qubit-block-wise. We plot the results from simulations of 10 qubit 50 layered RPQCs in

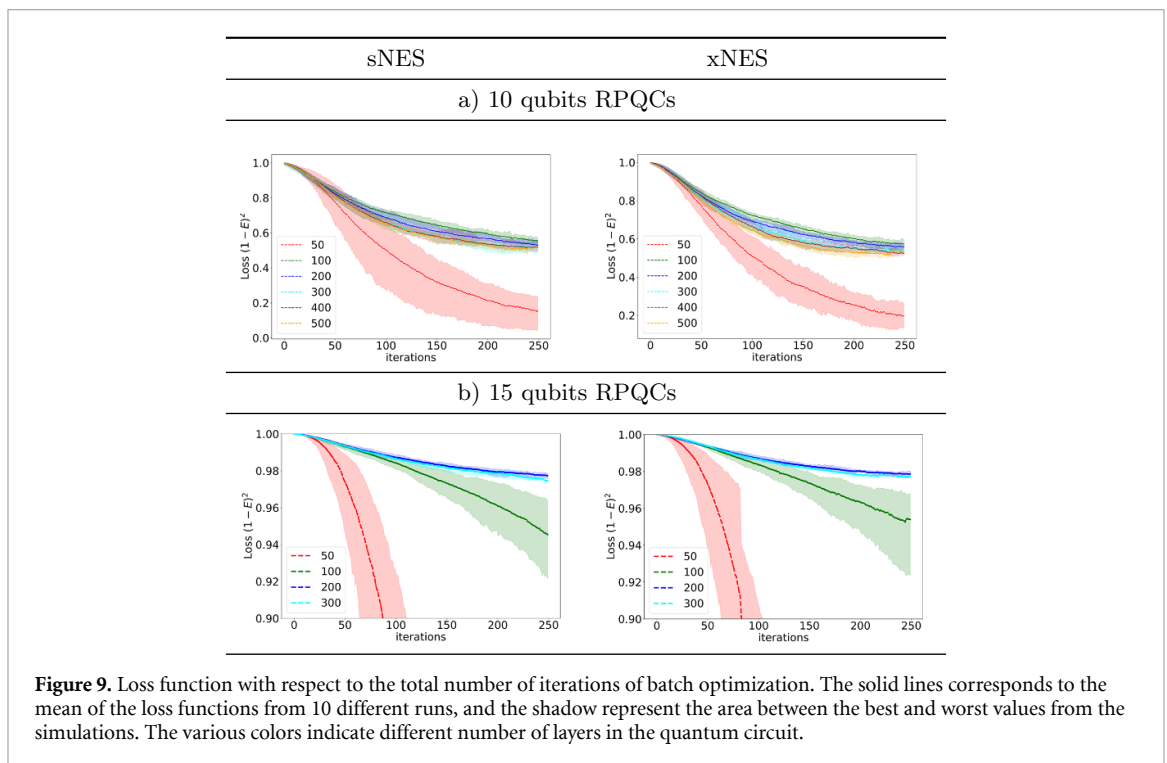
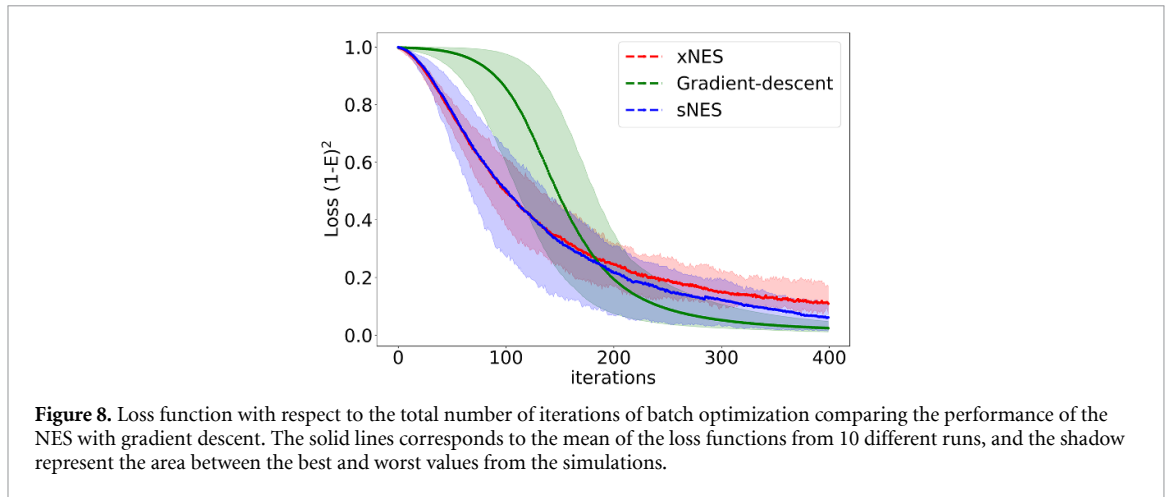
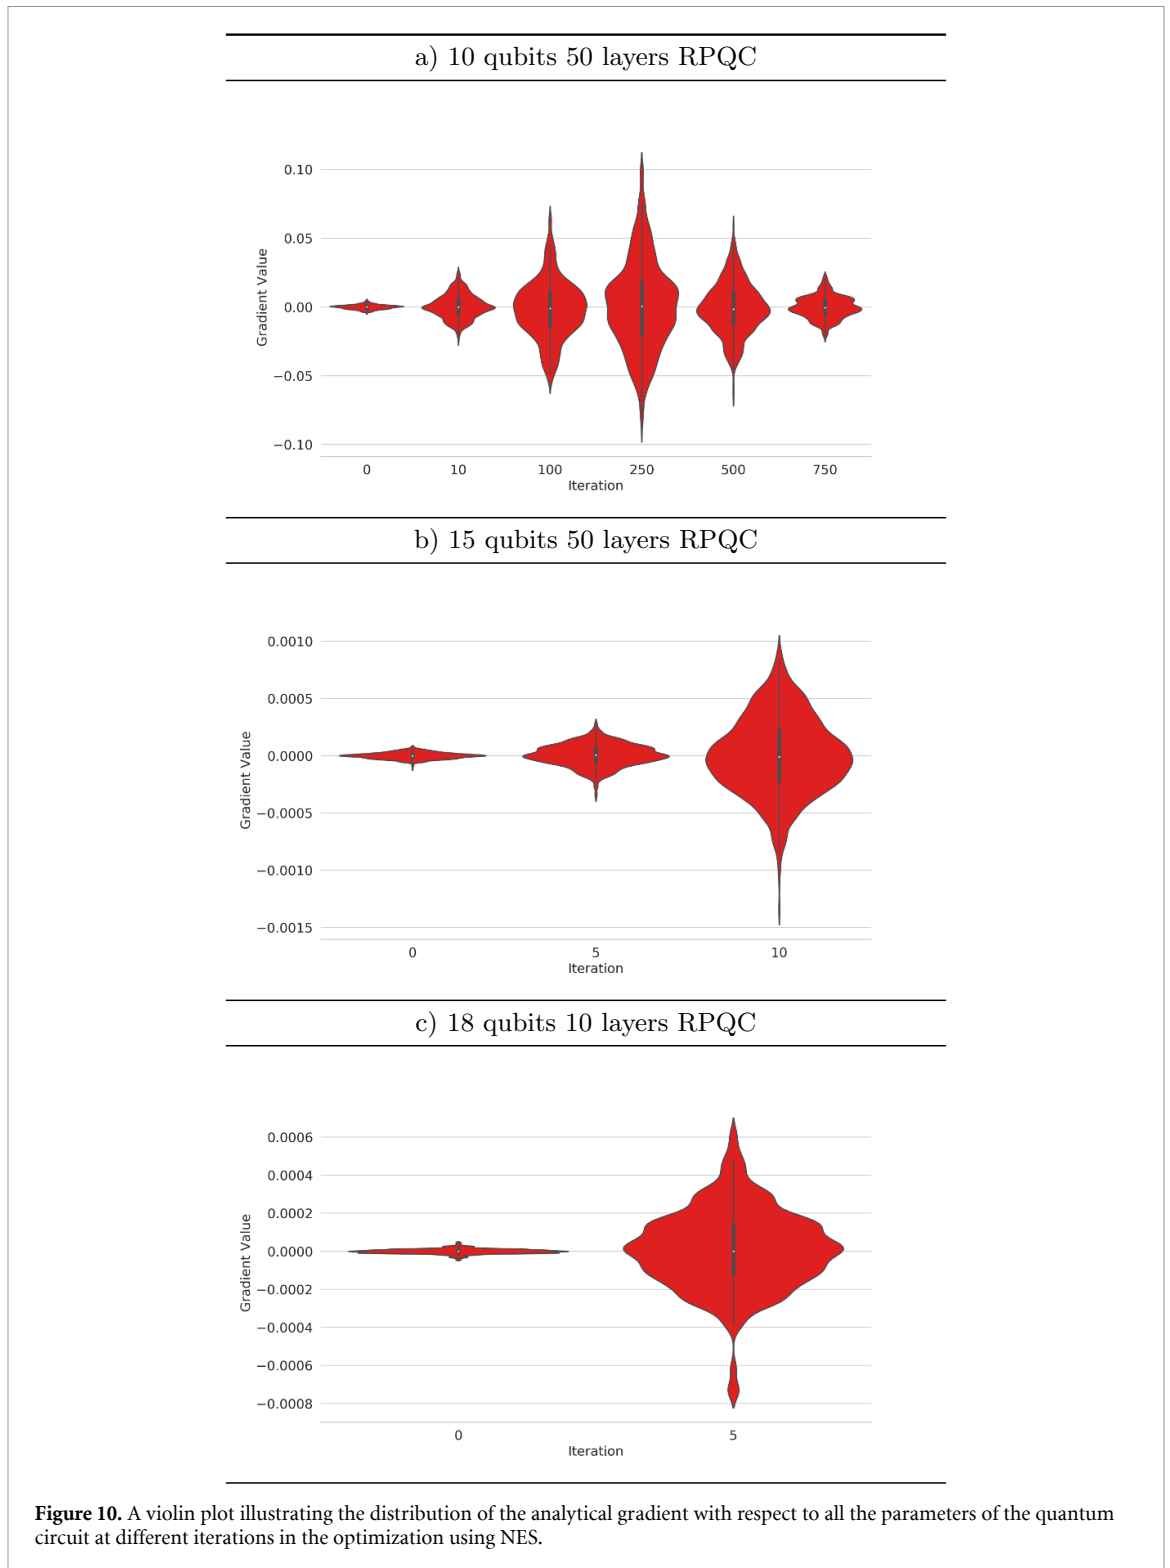


figure 7. It can be seen from the plot that the curves are identical for all the strategies, with the layer-wise batch optimization strategy a bit deviant because of the different batch size. We also carried out different simulation with RPQCs of larger depths using the different partition strategies and found that depending on the problem one might get faster convergence using one of the strategies, thus we recommend investigating the use of different strategies when dealing with larger RPQCs. Other techniques such as the parameter efficient circuit training techniques [49] can be also used together with NES to further improve the optimization procedure.

We also analysed the performance of NES with gradient descent with the batch optimization strategy. We use RPQCs with 10 qubits and 50 layers, and use batch size of 50 parameters and 16 walkers for the simulations. The results from the simulation is plotted in figure 8. The optimization with NES performs better than gradient descent initially, as seen previously in figure 4.

Finally we ran simulations to carry out the optimization of RPQCs with varying number of qubits and layers using both xNES and sNES by randomly dividing the parameters in batches of size 50 and using 16 walkers. We show the results for 10-qubit and 15-qubits circuits with different number of layers in figure 9. The plots show that we can perform optimization of significantly deep quantum circuits. As the circuits depth increases the rate of convergence decreases, and one might require a high number of circuit runs for the final result. Faster convergence of deep RPQCs might be achieved by selecting specific values of different parameters such as the initial learning rates, initial values of covariance matrix, number of walkers and the batch size. We leave such hyperparameter tuning for a future study.



4.5. Hybrid strategy

As shown above, the convergence with NES is sometimes slower than gradient-based methods in regions with non-vanishing gradients. In this section we explore a hybrid strategy that uses sNES in areas with vanishing gradients and regular gradient descent when a search direction has been found. For this, we simulated circuits with 10 qubits 50 layers, 15 qubits 50 layers, and 18 qubits 10 layers RPQCs with 500, 750 and 180 parameters respectively. We provide numerical evidence of the viability of this strategy by tracking the value of the analytical gradient with respect to all the parameters at different iterations in the optimization with NES, and plot the resulting distribution of the analytical gradients using violin plots in figure 10.

We tracked the analytical gradient distribution for the full optimization run for the 10 qubit 50 layers RPQC plotted in figure 10(a), and as expected, the distributions follow the pattern where it spreads initially to larger values and then narrows down about zero near the end of the optimization. This indicates that if

one can increase the spread of the gradients at the beginning of the optimization one might be able to use gradient-based methods to complete the process. Thus, we generate the distribution of the analytical gradients after every five iteration of optimization by sNES, and plot the distributions for 15 and 18 qubits RPQCs in figures 10(b) and (c). The distribution of the analytical gradient in figures 10(b) and (c) start to spread to higher values already after five iterations of optimization with sNES, which indicates that we can then use gradient-based methods in these cases for the full optimization. The choice of the number of optimization step will depend more on the problem, and one can use sNES until the magnitudes of the analytical gradients are large enough. These empirical results indicate that using a surrogate gradient as the one in sNES, one might be able to navigate through areas of vanishing gradients and then use gradient-based methods for finding the optimal solution.

5. Conclusion and outlook

We presented a strategy of representing parameters of quantum circuits as multinormal distributions for PQCs and use NES for optimizing randomly initialized PQCs. We have shown that the search direction used in NES can be amplified for randomly initialized PQCs, and that with appropriate number of walkers one can use them for optimization in all cases. We numerically analyse the performance of NES for PQCs, and show that they achieve nearly similar accuracies when compared to gradient-based methods. While sNES works pretty good for deep PQCs, xNES performs nearly well for shallow PQCs and outperforms sNES for problems where the parameters are correlated. We observe that the convergence curves for both NES and gradient-based methods are very similar, however, the number of circuit evaluations decrease significantly for NES and the initial rate of convergence seems to be higher too.

We also extend the optimization problem to deep quantum circuits by using of batch optimization strategy and show that one can use NES with batch optimization for deep PQCs. Based on various numerical simulation we observe that we can increase the rate of convergence by choosing appropriate values of different hyper-parameters, learning rate, number of walkers, and batch size. We also show that random partition of parameters into batches works comparably well as structured partitioning.

In the end we proposed a hybrid strategy, where we show that using NES we can amplify the gradients significantly and then use gradient-based methods for faster convergence. Our results indicate that the NES can be used as alternative to gradient-based methods in regions of vanishing gradients and can be a useful tool for exploring deep quantum circuits for different tasks.

Our study only explores the basic use of NES for optimization of randomly initialized PQCs, we believe that they can be used with different gradient-based methods to improve the overall optimization of deep quantum circuits. Other possibilities like, using other distributions to sample the parameters from, or combination of strategies to partition parameters into batches and adjusting the hyper-parameters for the problem in hand, can be used to further improve the training. These different directions are something we leave for future studies, as they require additional research.

Data availability statement

The data that support the findings of this study are available upon reasonable request from the authors.

Acknowledgments

The authors thank Sukin Sim and Jakob Kottmann for valuable suggestions and comments on the manuscript. A A and A A G acknowledge support by the U.S. office of naval research. M D and A A G acknowledge support by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Quantum Algorithms Teams Program. A A G acknowledges the generous support from Google, Inc. in the form of a Google Focused Award. A A G is grateful for the funding from the Vannevar Bush Faculty Fellowship Program, and the Canada Research Chairs Program. A A G thanks Anders G Frøseth for his generous support. Simulations were performed on the Niagara supercomputer at the SciNet HPC Consortium [50, 51]. SciNet is funded by: the Canada Foundation for Innovation; the Government of Ontario; Ontario Research Fund—Research Excellence; and the University of Toronto.

ORCID iDs

Abhinav Anand  <https://orcid.org/0000-0002-8081-2310>

Matthias Degroote  <https://orcid.org/0000-0002-8850-7708>

Alán Aspuru-Guzik  <https://orcid.org/0000-0002-8277-4434>

References

- [1] Arute F *et al* 2019 Quantum supremacy using a programmable superconducting processor *Nature* **574** 505
- [2] Jurcevic P *et al* 2020 Demonstration of quantum volume 64 on a superconducting quantum computing system (arXiv:2008.08571)
- [3] Pino J *et al* 2020 Demonstration of the QCCD trapped-ion quantum computer architecture (arXiv:2003.01293)
- [4] Preskill J 2018 Quantum computing in the NISQ era and beyond *Quantum* **2** 79
- [5] McClean J R, Romero J, Babbush R and Aspuru-Guzik A 2016 The theory of variational hybrid quantum-classical algorithms *New J. Phys.* **18** 023023
- [6] Peruzzo A, McClean J, Shadbolt P, Yung M-H, Zhou X-Q, Love P J, Aspuru-Guzik A and O'Brien J L 2014 A variational eigenvalue solver on a photonic quantum processor *Nat. Commun.* **5** 4213
- [7] Cao Y *et al* 2019 Quantum chemistry in the age of quantum computing *Chem. Rev.* **119** 10856
- [8] McArdle S, Endo S, Aspuru-Guzik A, Benjamin S C and Yuan X 2020 Quantum computational chemistry *Rev. Mod. Phys.* **92** 015003
- [9] Romero J, Olson J P and Aspuru-Guzik A 2017 Quantum autoencoders for efficient compression of quantum data *Quantum Sci. Technol.* **2** 045001
- [10] Wan K H, Dahlsten O, Kristjánsson H, Gardner R and Kim M 2017 Quantum generalisation of feedforward neural networks *npj Quantum Inf.* **3** 1
- [11] 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
- [12] Schuld M, Bocharov A, Svore K M and Wiebe N 2020 Circuit-centric quantum classifiers *Phys. Rev. A* **101** 032308
- [13] Grant E, Benedetti M, Cao S, Hallam A, Lockhart J, Stojevic V, Green A G and Severini S 2018 Hierarchical quantum classifiers *npj Quantum Inf.* **4** 1
- [14] Pérez-Salinas A, Cervera-Lierta A, Gil-Fuster E and Latorre J I 2020 Data re-uploading for a universal quantum classifier *Quantum* **4** 226
- [15] Romero J and Aspuru-Guzik A 2019 Variational quantum generators: generative adversarial quantum machine learning for continuous distributions (arXiv:1901.00848)
- [16] Zhu D *et al* 2019 Training of quantum circuits on a hybrid quantum computer *Sci. Adv.* **5** eaaw9918
- [17] Lloyd S and Weedbrook C 2018 Quantum generative adversarial learning *Phys. Rev. Lett.* **121** 040502
- [18] Dallaire-Demers P-L and Killoran N 2018 Quantum generative adversarial networks *Phys. Rev. A* **98** 012324
- [19] Hu L *et al* 2019 Quantum generative adversarial learning in a superconducting quantum circuit *Sci. Adv.* **5** eaav2761
- [20] Anand A, Romero J, Degroote M and Aspuru-Guzik A 2020 Experimental demonstration of a quantum generative adversarial network for continuous distributions (arXiv:2006.01976)
- [21] Farhi E, Goldstone J and Gutmann S 2014 A quantum approximate optimization algorithm (arXiv:1411.4028)
- [22] Kottmann J S, Krenn M, Kyaw T H, Alperin-Lea S and Aspuru-Guzik A 2020 Quantum computer-aided design of quantum optics hardware (arXiv:2006.03075)
- [23] Kyaw T H, Menke T, Sim S, Sawaya N P, Oliver W D, Guerreschi G G and Aspuru-Guzik A 2020 Quantum computer-aided design: digital quantum simulation of quantum processors (arXiv:2006.03070)
- [24] Bharti K *et al* 2021 Noisy intermediate-scale quantum (NISQ) algorithms (arXiv:2101.08448)
- [25] McClean J R, Boixo S, Smelyanskiy V N, Babbush R and Neven H 2018 Barren plateaus in quantum neural network training landscapes *Nat. Commun.* **9** 1
- [26] Cerezo M, Sone A, Volkoff T, Cincio L and Coles P J 2020 Cost-function-dependent barren plateaus in shallow quantum neural networks (arXiv:2001.00550)
- [27] Wang S, Fontana E, Cerezo M, Sharma K, Sone A, Cincio L and Coles P J 2020 Noise-induced barren plateaus in variational quantum algorithms (arXiv:2007.14384)
- [28] Marrero C O, Kieferová M and Wiebe N 2020 Entanglement induced barren plateaus (arXiv:2010.15968)
- [29] Uvarov A and Biamonte J 2020 On barren plateaus and cost function locality in variational quantum algorithms (arXiv:2011.10530)
- [30] Arrasmith A, Cerezo M, Czarnik P, Cincio L and Coles P J 2020 Effect of barren plateaus on gradient-free optimization (arXiv:2011.12245)
- [31] Verdon G, Broughton M, McClean J R, Sung K J, Babbush R, Jiang Z, Neven H and Mohseni M 2019 Learning to learn with quantum neural networks via classical neural networks (arXiv:1907.05415)
- [32] Volkoff T and Coles P J 2020 Large gradients via correlation in random parameterized quantum circuits (arXiv:2005.12200)
- [33] Rechenberg I 1978 Evolutionsstrategien *Simulationsmethoden in der Medizin und Biologie* (Berlin: Springer) pp 83–114
- [34] Schwefel H-P 1977 *Numerische Optimierung von Computer-Modellen Mittels der Evolutionsstrategie: mit Einer Vergleichenden Einführung in die Hill-Climbing- und Zufallsstrategie* vol 1 (Basel: Birkhäuser) (<https://doi.org/10.1007/978-3-0348-5927-1>)
- [35] Wierstra D, Schaul T, Peters J and Schmidhuber J 2008 Natural evolution strategies 2008 *IEEE Congress Evolutionary Computation (IEEE World Congress Computational Intelligence)* (IEEE) pp 3381–7
- [36] Glasmachers T, Schaul T, Yi S, Wierstra D and Schmidhuber J 2010 Exponential natural evolution strategies *Proc. 12th Annual Conf. Genetic and Evolutionary Computation* pp 393–400
- [37] Zhao T, Carleo G, Stokes J and Veerapaneni S 2020 Natural evolution strategies and variational Monte Carlo *Mach. Learn.: Sci. Technol.* **2** 02LT01
- [38] Coopmans L, Luo D, Kells G, Clark B K and Carrasquilla J 2020 Protocol discovery for the quantum control of majoranas by differential programming and natural evolution strategies (arXiv:2008.09128)
- [39] Wierstra D, Schaul T, Glasmachers T, Sun Y, Peters J and Schmidhuber J 2014 Natural evolution strategies *J. Mach. Learn. Res.* **15** 949
- [40] Hansen N and Ostermeier A 2001 Completely derandomized self-adaptation in evolution strategies *Evol. Comput.* **9** 159
- [41] Salimans T, Ho J, Chen X, Sidor S and Sutskever I 2017 Evolution strategies as a scalable alternative to reinforcement learning (arXiv:1703.03864)
- [42] Spall J C 1992 Multivariate stochastic approximation using a simultaneous perturbation gradient approximation *IEEE Trans. Autom. Control* **37** 332
- [43] Amari S-I 1998 Natural gradient works efficiently in learning *Neural Comput.* **10** 251
- [44] Amari S-I and Douglas S C 1998 Why natural gradient? *Proc. 1998 IEEE Int. Conf. Acoustics, Speech and Signal Processing, ICASSP'98 (Cat. No. 98CH36181)* vol 2 (IEEE) pp 1213–16

- [45] Kottmann J et al 2021 Tequila: a platform for rapid development of quantum algorithms *Quantum Sci. Technol.* **6** 024009
- [46] Suzuki Y et al 2020 Qulacs: a fast and versatile quantum circuit simulator for research purpose (arXiv:2011.13524)
- [47] Cervera-Lierta A, Kottmann J S and Aspuru-Guzik A 2020 The meta-variational quantum eigensolver (meta-VQE): learning energy profiles of parameterized Hamiltonians for quantum simulation (arXiv:2009.13545)
- [48] Skolik A, McClean J R, Mohseni M, van der Smagt P and Leib M 2020 Layerwise learning for quantum neural networks (arXiv:2006.14904)
- [49] Sim S, Fontalvo J R, Gonthier J F and Kunitsa A A 2021 Adaptive pruning-based optimization of parameterized quantum circuits *Quantum Sci. Technol.* **6** 025019
- [50] Loken C et al 2010 SciNet: lessons learned from building a power-efficient top-20 system and data centre *J. Phys.: Conf. Ser.* **256** 012026
- [51] Ponce M et al 2019 Deploying a top-100 supercomputer for large parallel workloads: the Niagara supercomputer *Proc. Practice and Experience in Advanced Research Computing on Rise of the Machines (Learning) (PEARC'19)* (New York: Association for Computing Machinery)