

PAPER

Using models to improve optimizers for variational quantum algorithms

To cite this article: Kevin J Sung *et al* 2020 *Quantum Sci. Technol.* **5** 044008

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



IOP Publishing
A VIRTUAL CONFERENCE
QUANTUM 2020
19–22 OCTOBER

CO-ORGANIZERS:

 中国物理学会
Chinese Physical Society

 中国科学技术大学
University of Science and Technology of China

REGISTER TODAY

Quantum Science and Technology



PAPER

Using models to improve optimizers for variational quantum algorithms

RECEIVED
25 May 2020REVISED
7 August 2020ACCEPTED FOR PUBLICATION
9 September 2020PUBLISHED
24 September 2020Kevin J Sung^{1,2,*} , Jiahao Yao³, Matthew P Harrigan¹, Nicholas C Rubin¹ , Zhang Jiang¹, Lin Lin^{3,4}, Ryan Babbush¹ and Jarrod R McClean^{1,*}¹ Google Research, Venice, CA² Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI, United States of America³ Department of Mathematics, University of California, Berkeley, CA, United States of America⁴ Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley, CA, United States of America

* Author to whom any correspondence should be addressed.

E-mail: kevjsung@umich.edu and jmcclean@google.com**Keywords:** quantum algorithms, variational quantum algorithms, optimizationSupplementary material for this article is available [online](#)

Abstract

Variational quantum algorithms are a leading candidate for early applications on noisy intermediate-scale quantum computers. These algorithms depend on a classical optimization outer-loop that minimizes some function of a parameterized quantum circuit. In practice, finite sampling error and gate errors make this a stochastic optimization with unique challenges that must be addressed at the level of the optimizer. The sharp trade-off between precision and sampling time in conjunction with experimental constraints necessitates the development of new optimization strategies to minimize overall wall clock time in this setting. In this work, we introduce two optimization methods and numerically compare their performance with common methods in use today. The methods are surrogate model-based algorithms designed to improve reuse of collected data. They do so by utilizing a least-squares quadratic fit of sampled function values within a moving trusted region to estimate the gradient or a policy gradient. To make fair comparisons between optimization methods, we develop experimentally relevant cost models designed to balance efficiency in testing and accuracy with respect to cloud quantum computing systems. The results here underscore the need to both use relevant cost models and optimize hyperparameters of existing optimization methods for competitive performance. The methods introduced here have several practical advantages in realistic experimental settings, and we have used one of them successfully in a separately published experiment on Google's Sycamore device.

1. Introduction

With recent developments in quantum hardware, including the ability to perform select tasks faster than classical supercomputers [1], the push toward practical applications on these devices has intensified. Variational quantum algorithms are among the top candidates for early applications on noisy intermediate-scale quantum computers [2–4]. These algorithms can be used to approximate ground energies of Hamiltonians or find approximate solutions to discrete optimization problems. A main component of these algorithms is the minimization of some function of a parameterized quantum state, where that function is measured using the quantum computer. Commonly, the function is the expectation value of a Hamiltonian, determined by the problem of interest. The presence of sampling error and gate errors makes the function stochastic, and the stochasticity due to sampling error is fundamental to measuring the values on a quantum device. The output of this stochastic function is fed to a classical optimizer, and it is those optimizers and constraints presented by real devices that we will focus on here.

As the classical optimizers are at the core of variational quantum algorithms, their performance can determine the resources required to solve a problem. Non-linear optimization of continuous functions of the type that exist in variational quantum algorithms are commonplace in fields like machine learning, but quantum systems offer unique trade-offs that must be considered to improve efficiency. Given the current focus on these algorithms and the core role played by the optimizer, there have been a number of works evaluating the performance of optimizers for different problems and contexts. For example, at least two experimental implementations of variational algorithms [2, 5] used the Nelder–Mead simplex algorithm [6] to optimize the objective function. Other experimental implementations [7–11] used algorithms including simultaneous perturbation stochastic approximation (SPSA) [12], Bayesian optimization [13], particle swarm optimization [14], dividing rectangles [15], and gradient descent. In addition, there have been a number of numerical investigations of optimization in the context of variational quantum algorithms. Several of these studies introduce novel heuristics and test them numerically on example problems [16–21]. Other work [22–29] has compared the performance of methods including Nelder–Mead, limited-memory Broyden–Fletcher–Goldfarb–Shanno [30], constrained optimization by linear approximation [31], Powell’s method [32], SPSA, RBFOpt [33], stable noisy optimization by branch and fit [34], bound optimization by quadratic approximation [35], mesh adaptive direct search [36], implicit filtering [37], policy-gradient-based reinforcement learning [38], and natural gradient [28].

There is a considerable body of work in evaluating optimizers for use in variational algorithms, but not all of these works use cost metrics relevant to quantum experiments. For example, it is common to evaluate a suite of optimizers based on number of optimizer iterations required for convergence to a local optima, using noiseless function evaluations. However, the inherent quantum nature of the sampling procedure implies that the first iteration could have taken an unbounded amount of experimental time in such a setup (noiseless evaluation), and hence conclusions based on such studies may not be applicable to experiments. A meaningful comparison of these methods must treat the stochastic nature of the objective function and related costs in terms of experimental time to solution to properly compare methods. While some past works do account for the effect of stochastic noise [20, 21, 25, 26], in this work we additionally incorporate other experimental parameters into our cost models. In developing our models, we focus on the case of superconducting quantum computers accessed through the internet, though our models can be easily modified for other architectures. We account for parameters such as the sampling rate of the quantum processor and the latency induced by communicating over the internet. The proper choice of optimizer ultimately depends on the details of the experiment constraints.

In consideration of constraints we did not find satisfied in other methods, we introduce two surrogate model-based optimization algorithms we call model gradient descent (MGD) and model policy gradient (MPG) and numerically compare their performance against commonly used methods. In particular, we target the tendency for local methods to under-utilize the existing history of function evaluations. We have successfully used MGD in an experimental implementation of the quantum approximate optimization algorithm [39] on a superconducting qubit processor [40]. We perform systematic tuning of optimizer hyperparameters before comparison for all methods, and measure performance using estimates of actual wall clock time needed in a realistic experimental setting. An important, though unsurprising, implication of our results is that hyperparameter tuning under the correct cost models is crucial for performance in practice.

The outline of this work is as follows. In section 2 we set up the example problems we study and describe in more depth the problem of developing efficient cost models to allow comparison of methods. In section 3 we describe the optimizers we study and how we tuned their hyperparameters. After this setup, we compare the performance of optimizers numerically in section 4 using our developed cost models. At a glance, our results highlight the importance of different cost model features, how constraints influence the optimal choice of optimizer, and the importance of hyperparameter optimization. Stochastic optimizers with hyperparameters permitting varying levels of noise in the objective are found to be generally more robust and efficient. Finally, we end with some concluding thoughts in section 5.

2. Problems studied and cost models

2.1. Problems studied

As the performance of an optimizer can be intimately tied to the problem studied, it is important to look at a range of problems in evaluating their relative performance. As two of the most common areas studied in variational quantum algorithms are combinatorial optimization and ground state preparation of fermionic systems, we select these for our sample problems. Here we aim to clarify the details of the systems, circuit ansatz, and initial parameters modeled in our numerical tests.

While multi-modality of cost functions is an important consideration in variational quantum algorithms, it turns out that even optimization within a single convex basin can be challenging enough to warrant independent investigation due to constraints imposed by the quantum device. To this end, we assume throughout that we have knowledge of an initial guess which is in the convex vicinity of an optimum and our goal is simply to converge to that local optimum. Several strategies have been proposed for choosing such an initial guess in contexts including optimization and chemistry [16, 17, 23, 41].

2.1.1. Max-Cut on 3-regular graphs

The maximum cut problem (Max-Cut) is widely studied and known to be NP-hard. It has been used in several previous experimental implementations of variational quantum algorithms [8, 40] and hence allows for straightforward performance comparisons. The problem is specified by an undirected graph on n vertices and the goal is to label each vertex with either $+1$ or -1 in order to maximize the number of edges whose vertices have different labels. This cost function is represented by the Hamiltonian

$$C = \sum_{\langle ij \rangle} \frac{1}{2} (I - Z_i Z_j), \quad (1)$$

where Z_j is the standard Pauli Z operator applied to qubit j which is node j on the graph, and $\langle i, j \rangle$ ranges over the edges of the graph. The goal is to find a computational basis state that *maximizes* the Hamiltonian.

We use the quantum approximate optimization algorithm (QAOA) [39] ansatz used to approximately solve the Max-Cut problem on random 3-regular graphs. The QAOA ansatz depends on the number of rounds, $p > 0$, and is parameterized by $2p$ real numbers $\gamma = (\gamma_1, \dots, \gamma_p)$ and $\beta = (\beta_1, \dots, \beta_p)$. The ansatz is

$$|\gamma, \beta\rangle = U_B(\beta_p) U_C(\gamma_p) \dots U_B(\beta_1) U_C(\gamma_1) |+\rangle^{\otimes n}, \quad (2)$$

where

$$U_C(\gamma) = e^{-i\gamma C}, \quad U_B(\beta) = e^{-i\beta B}, \quad B = \sum_{i=1}^n X_i, \quad (3)$$

and $|+\rangle^{\otimes n}$ is the uniform superposition of all 2^n computational basis states.

For our numerics, we focus on a randomly chosen instance to minimize the number of uncontrolled variables. Moreover, for QAOA focusing on a single instance is justified because the optimization landscape has been shown to concentrate for different randomly chosen instances [41]. To obtain an initial guess for this problem, we classically computed a locally optimal parameter vector and then perturbed it with a uniformly random vector of length 0.1. At $p = 1$ the optimal parameter vector had a length of 0.462, and at $p = 5$, 1.285.

In our numerics we report the approximation ratio

$$\frac{\langle \gamma, \beta | C | \gamma, \beta \rangle}{C_{\max}}, \quad (4)$$

where $C_{\max} = \max_z \langle z | C | z \rangle$. The goal is to maximize this value, which falls in the range $[0, 1]$.

2.1.2. Sherrington–Kirkpatrick model

Another model we consider is the Sherrington–Kirkpatrick (SK) model [42], which is a canonical example of a frustrated spin glass. It has been used in at least one previous experimental implementation of variational algorithms [40]. The Hamiltonian is given by

$$H = \sum_{i < j} J_{ij} Z_i Z_j, \quad (5)$$

where J_{ij} is selected uniformly at random from $\{-1, 1\}$. We use the QAOA ansatz to approximate the solution of this problem, by minimizing the expected cost.

Again, for our numerics we focus on a single randomly generated instance, where generality of performance is supported by concentration results in QAOA. As an initial guess for this problem, we classically computed a locally optimal parameter vector and then perturbed it with a uniformly random vector of length 0.1. At $p = 1$ the optimal parameter vector had a length of 0.452, and at $p = 5$, 1.044.

For comparison between problems, we normalize energy values E to new values E' by the formula

$$E' = \frac{E - E_{\max}}{E_{\min} - E_{\max}}, \quad (6)$$

where E_{\min} and E_{\max} are the lowest and highest eigenvalues of the Hamiltonian, respectively. Thus we are in fact maximizing this normalized energy value, which falls in the range $[0, 1]$.

2.1.3. Hubbard model

We study the task of approximating the ground state energy of the two-dimensional Hubbard model [43], a widely studied model that has resisted exact solution for decades in large size limits. It is believed to be relevant to understanding high-temperature superconductivity [44]. The Hamiltonian of the Hubbard model is

$$H = -t \sum_{\langle i,j \rangle, \sigma} \left(a_{i,\sigma}^\dagger a_{j,\sigma} + a_{j,\sigma}^\dagger a_{i,\sigma} \right) + U \sum_i a_{i,\uparrow}^\dagger a_{i,\uparrow} a_{i,\downarrow}^\dagger a_{i,\downarrow} \quad (7)$$

$$= T + V \quad (8)$$

$$= T_h + T_v + V, \quad (9)$$

where the $a_{i,\sigma}$ are fermionic annihilation operators, $\langle i,j \rangle$ ranges over edges in the lattice, $\sigma \in \{\uparrow, \downarrow\}$ is a spin degree of freedom, and we have split the sum into the hopping term T and interaction term V . T is further decomposed into sub-terms T_h and T_v corresponding to horizontal and vertical edges, respectively. We set $t = 1$ and $U = 4$ for our numerical experiments, which corresponds to a regime of modest correlation ill-suited for mean-field methods.

We use a ‘Hamiltonian variational’ ansatz similar to the one in reference [16]. It is inspired by the idea of state preparation via adiabatic evolution. Similar to QAOA, our ansatz has a basic circuit repeated p times, but for flexibility it is varied non-uniformly with respect to hopping. The basic circuit has three parameters which we call θ_h , θ_v , and θ_U , and it approximates a unitary of the form

$$\exp[-i(\theta_h T_h + \theta_v T_v + \theta_U V)]. \quad (10)$$

The approximation is achieved using a second-order Trotter step based on the fermionic swap network [45], in which a swap network is used to apply the terms of the Hamiltonian and then the same network is applied but in reverse order. Because the swap network can be implemented with only linear qubit connectivity, this ansatz is amenable to implementation on near-term superconducting qubit hardware. The ansatz is similar to the one used in reference [16] but corresponds to a different ordering of terms. In total there are $3p$ parameters.

We study the model at half-filling. Our numerics are performed on the 2×2 system, which under standard encodings corresponds to an 8 qubit system. For our initial state we use a ground state of the hopping term that is precisely described in appendix C (<https://stacks.iop.org/QST/5/044008/mmedia>). This state is easy to prepare on a quantum computer and is expected to be adiabatically connected to the ground state of H for modest values of t/U . For our initial guess, we set the parameters so that the ansatz circuit consists of a sequence of second-order Trotter steps approximating the dynamics of the time-dependent Hamiltonian $H(t) = T + (t/A)V$ for $t \in [0, A]$, where $A = 0.1 \cdot Up$. This choice is motivated by the idea of state preparation via adiabatic evolution.

As with the SK model, we normalize energy values E to new values E' by the formula

$$E' = \frac{E - E_{\max}}{E_{\min} - E_{\max}}, \quad (11)$$

where E_{\min} and E_{\max} are the lowest and highest eigenvalues of the Hamiltonian, respectively. Thus we are in fact maximizing this normalized energy value, which falls in the range $[0, 1]$.

2.2. Cost models

An essential element of developing and improving optimizers for variational algorithms is an accurate cost model that respects the quantum nature of the problem and imperfections of the device. Studies that restrict evaluation of optimizers to abstract ‘number of iterations’ using perfect function queries can yield faulty conclusions and hide the implication that a single function evaluation to that precision could have taken years or more. A core challenge is the stochastic nature of the function evaluation and shot limited precision in the estimates. Moreover, imperfections in the device and implementation can complicate matters. Unfortunately, without a quantum device, precise simulation of the impact of noise can be prohibitively expensive, and so a balance must be struck between accuracy and cost effectiveness of the simulations to maximize applicability. Here we detail how we construct our models to strike this balance.

We restrict our interest to minimizing the expected energy of a Hamiltonian H with efficient Pauli expansions $H = \sum_j \alpha_j P_j$ (in the case of the Hubbard model (7), the Jordan–Wigner transformation [46] is applied to obtain the Pauli expansion), so the objective function is

$$f(\theta) = \langle \theta | H | \theta \rangle, \quad (12)$$

where $|\theta\rangle$ represents the ansatz state with parameters θ . Most of the optimizers that we present results for use queries to the objective function without any additional kinds of queries, but we also present results for stochastic gradient descent, which queries the gradient.

2.2.1. Objective function queries

The exact estimator used to query the objective function on the quantum device can take a wide variety of forms depending on factors in the device and the problem of interest. At a glance, however, a query to the objective function is often answered by measuring the expectation values of the terms P_j and using the coefficients α_j to form an estimate of $f(\theta)$. When simulated in the most accurate way, the measurement of each individual term implies a variance on the estimate which is state-dependent, and functions like a Bernoulli random variable. Moreover, the variance of that measurement can be influenced by parallel measurements being performed, even when they commute [3]. Trade-offs in the influences of these factors have inspired recent research in developing more efficient estimators with a given number of samples [47–51]. However, perfect emulation of these proposals can be prohibitively expensive, even in classical simulation of small systems, and hence it is desirable to develop models of the process that strike a good balance between accuracy and simulation cost so that the full variational process can be simulated on a range of systems.

In the cases of Max-Cut and the SK model, the Hamiltonian is diagonal and all of its terms can be measured simultaneously in one shot. In our numerical experiments, we simulated these measurements directly. However, for non-diagonal Hamiltonians such as the Hubbard model, we take a different strategy.

As there are many terms in the sum, which are typically evaluated by repeated and independent measurement, a Gaussian random function query turns out to be a good and extremely cost effective model. That is, in our simulations a query to the objective function is modeled as

$$f(\theta) = \langle \theta | H | \theta \rangle + \mathcal{N}(0, \lambda^2/M). \quad (13)$$

$\langle \theta | H | \theta \rangle$ is evaluated exactly, $\mathcal{N}(\mu, \sigma^2)$ is a normal random variable with mean μ and variance σ^2 , and M is the number of repeated experiment repetitions. Note that even in the presence of hardware errors, the expectation value of the Hamiltonian would be the sum of many independent random variables, so this would still be a good model. Here, we estimate the variance is using a known lower bound for common measurement strategies, previously derived for the general case

$$\lambda^2 = \left(\sum_j |\alpha_j| \right)^2, \quad (14)$$

which empirically we have observed to be loose when compared with exact models, but qualitatively matches the behavior and overestimates the number of measurements by a factor of 2 in many cases. We note that a wealth of other strategies have been developed to shrink the effective variance for a fixed number of queries M [47–51], but we do not consider them in detail here. Since the bound we use is a worst-case bound that is independent of the quantum state, our cost estimates are likely to be conservative.

The dependence of the variance of the estimate on the number of samples represents a key trade-off we consider in many algorithms here, as some optimizers can tolerate heavier amounts of noise than others, and hence we take the number of shots at each iterate to be an important hyperparameter. In our numerical experiments on the Hubbard model, we simulated queries by computing the exact expectation value and then artificially adding noise drawn from a normal distribution, using this bound to determine the variance of the distribution for a specified number of measurement shots.

2.2.2. Gradient queries

For optimizers that use analytic gradient queries, we assume that queries to the gradient of the objective function are answered by applying the ‘parameter-shift rule’ [52–54]. This is a method of obtaining an unbiased estimator of the gradient without using ancilla qubits, and applies to ansatz of the form

$$|\theta\rangle = \exp(-i\theta_p A_p) \dots \exp(-i\theta_1 A_1) |\psi\rangle \quad (15)$$

where for our purposes each A_j is a Hermitian sum of commuting Pauli matrices. The technique exploits the fact that if A_j has two eigenvalues $\pm r$, then $\frac{\partial f}{\partial \theta_j}(\theta) = r(f(\theta^+) - f(\theta^-))$ where θ^+ is θ but with the j th coordinate equal to $\theta_j + \frac{\pi}{4r}$ and θ^- is θ but with the j th coordinate equal to $\theta_j - \frac{\pi}{4r}$. If some parameters are constrained to be the same, then the derivative is obtained by summing the results of this expression for each parameter; the number of objective function queries needed is then two times the number of those parameters. If $A_j = \sum_k P_k$ for commuting Pauli operators P_k , then we decompose

$\exp(-i\theta_j A_j) = \prod_k \exp(-i\theta_j P_k)$ and then apply the previous rule. Thus, the cost of evaluating the partial derivative is proportional to the number of terms in the sum, in a loose way. In practice, this sum is evaluated stochastically with a probability depending on the weight of the term in the sum [55].

2.2.3. Wall clock time

Ultimately, one is interested in minimizing the amount of time it takes to run a complete experiment to some fixed precision. The models we develop here are meant to capture this in a cost efficient way, without using a wildly inaccurate proxy like mere ‘number of optimizer iterations’. To this end, we not only consider the sampling noise, but also constraints like latency concerns inherent to real experiments.

To estimate the running time of an experiment we develop a model based on superconducting qubits [56, 57]. We also assume the user is executing the experiment through a cloud computing service, potentially introducing network latency. We consider three scenarios regarding network latency: zero latency, corresponding to the optimizer running completely on the server side; circuit batching, in which the user is allowed to send multiple circuits to the service in one batch; and finally no circuit batching, where the user is only allowed to send one circuit at a time.

The total running time of an experiment is equal to the number of queries made times the amount of time it takes to satisfy a single query. The time needed to satisfy a single query can be split into the time T_{sample} used in sampling circuits on the quantum processor, the time T_{switch} representing the overhead in switching between different circuits, and T_{cloud} representing the latency in communicating over the Internet. We have $T_{\text{sample}} = M/s$ where M is the number of measurements made to satisfy the query and s is the sampling rate of the processor; $T_{\text{switch}} = r \times c$ where r is the overhead in readying the quantum processor to execute a circuit and c is the number of different circuits executed; and $T_{\text{cloud}} = \ell \times c/b$ where ℓ is the network round-trip time for communicating with the cloud server and b is the number of circuits sent to the server in a single round of communication. We use the values $s = 10^5$ Hz and $r = 0.1$ s. This sampling rate has not yet been achieved experimentally but is plausible assuming an order of magnitude or two improvement in current capabilities is possible; a recent experiment achieved a sampling rate of about 5×10^3 Hz [1]. When including network latency, we set $\ell = 4.0$ s; this value is based on our own experience executing experiments through an internal cloud interface. The value of b depends on the details of the algorithm. We ignore as negligible the time taken by the classical optimization algorithm to select parameters for querying, as the optimizers here use relatively simple classical updates.

3. Optimization strategies

3.1. Choice of optimizers

A wide range of optimizers now exist for continuous, non-linear optimizations, with different strengths and weaknesses. One key element for consideration is the stochastic nature of our objective function and its relation to the number of measurements made for each function evaluation. Some optimizers were designed with noiseless (up to reasonable precision limits) function evaluations in mind, and are relatively unstable with respect to even small amounts of noise. While one could insist on a number of measurements that renders the function evaluations essentially exact, this incurs a huge overhead per iteration. We group algorithms into two categories, distinguished by whether they have inherent hyperparameters that allow them to adjust their resilience to noise. If an algorithm in practice requires that the input be given to a fixed precision in order to be stable, we term it deterministic. If it has a hyperparameter that naturally allows it to accept more or less noise, we call it stochastic.

The difference between the two classes can be subtle, and depend on the details of implementation. For example, a gradient descent implementation that makes use of an exact line search can accidentally rule out good regions of space from small wobbles in a query value, and is hence deterministic. However, if that sample implementation substitutes a fixed step with a learning rate, it is not only more robust to noise, but that learning rate can be adjusted to match noise levels in the objective queries. Hence we term that a stochastic optimizer. Considering the costs of each with external hyperparameters (e.g. number of measurements) and internal hyperparameters (e.g. learning rate) tuned for optimal performance will show us these trade-offs.

Overall, we investigated six different optimizers. Four of these have been studied in past work, and the last two are surrogate model-based optimizers that we introduce here. Surrogate model-based optimizers construct a model of the objective function using previously evaluated points and use the model to determine what points to evaluate next. They are popular choices for the optimization of objective functions that are expensive to evaluate or noisy (or both) [58, 59].

Listed briefly, the optimizers we study here are:

- Deterministic algorithms:
 - * The Nelder–Mead simplex method [6]. This method has been used in previous theoretical [22, 23] and experimental [2, 5] works on variational algorithms. We used the implementation from SciPy [60].
 - * Bounded optimization by quadratic approximation (BOBYQA) [35]. This is a surrogate model based algorithm that uses an interpolating quadratic model to approximate the objective function, and has been studied in a previous work on variational algorithms [26]. We used the implementation from the Python package Py-BOBYQA [59].
- Stochastic algorithms:
 - * SPSA [12]. This method has also been used in previous theoretical [24] and experimental [7] works on variational algorithms. We used our own implementation.
 - * Stochastic gradient descent using analytic gradient measurements obtained via the ‘parameter-shift rule’ [52–54].
 - * MGD. This is a surrogate model-based algorithm we introduce here that uses a least-squares quadratic model to estimate the gradient of the objective function. We give pseudocode in appendix A.
 - * MPG. Building on the vanilla policy gradient method [25], this method additionally introduces a least-squares quadratic model to reduce the variance in the estimation of the policy gradient. We give pseudocode in appendix A.

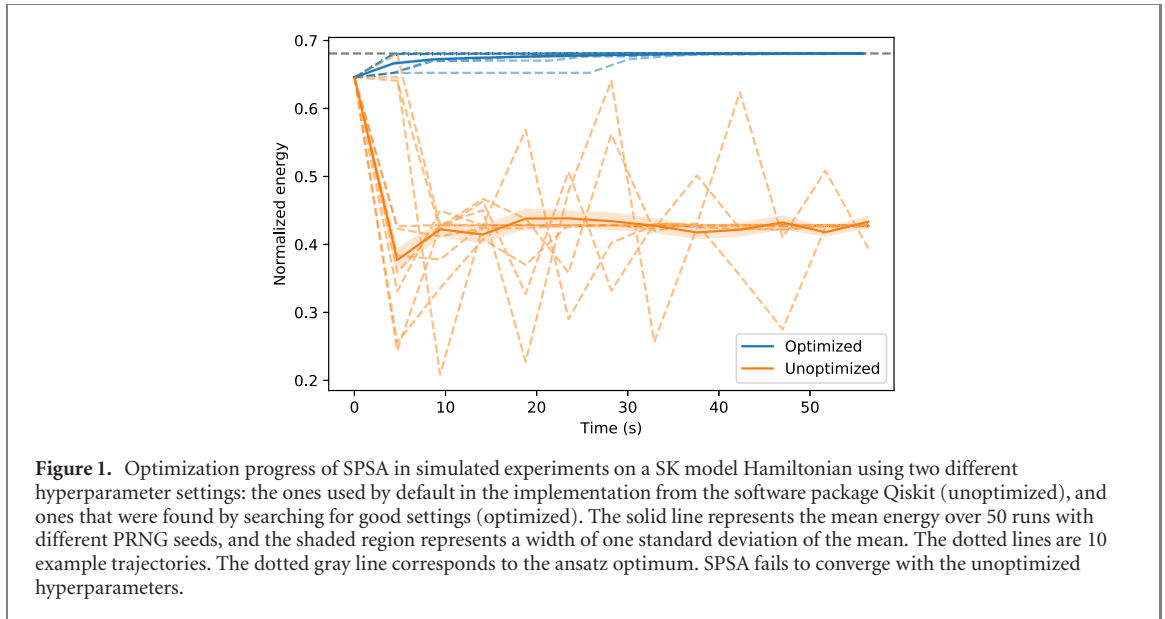
3.2. MGD and policy gradient

In this section we describe and motivate the design choices of our new algorithms, MGD and MPG, which are described in pseudocode in algorithms 1 and 2. These are surrogate model-based methods which use least-squares regression to fit quadratic models of the objective function. A key expense in variational quantum algorithms is the evaluation of the function at different points, which is costly due to the underlying variance. Hence, it would be beneficial to reuse the history of point evaluations, rather than to discard them at each iteration. For local optimizations where iterates proceed gradually, it seems intuitive that this should be possible. Eventually, if one collected enough points in a small enough region, it should be possible to construct a surrogate model that is more accurate than raw function evaluations at a fixed number of measurements.

As a combination of this motivation and simplicity, we use a least-squares fit to a quadratic function. However, it is also clear that if the region of sampled points is too large, the function may not be well approximated by a quadratic, hence we use a trusted region of sample points, which may be new or reused from previous iterates.

In each iteration, the algorithms sample a number of points randomly from the vicinity of the current iterate. They fit quadratic models to these points and other previously evaluated points within the vicinity. Finally, MGD uses the gradient of this quadratic model as an approximation to the true gradient and performs gradient descent; MPG queries the model to evaluate a large batch of data points and performs policy gradient optimization. The reason we did not use standard trust-region solution techniques after building the quadratic model is that we found empirically that the eigenvalues of the Hessian of the quadratic model built upon stochastic function evaluations may be slightly negative, which dictates in a standard trust region solution method that the solution is on the exterior of the trust region. This constant jumping to the exterior of the trust region represented a sort of fundamental inefficiency under stochastic functions. In contrast, the gradient or policy gradient of the model, while stochastic, represented a reliable estimator that, in conjunction with techniques like a fixed learning rate, combined the increased accuracy of additional samples with the robustness of a stochastic gradient descent.

To enhance the performance and stability of the methods, we introduced several hyperparameters to our algorithms. In particular, as algorithms approach an optimum, decreasing the radius of the neighborhood from which points are sampled is expected to give a more accurate estimate of the function value and its gradient. Thus, we introduce a hyperparameter ξ for MGD which controls the rate at which the radius decreases. As for MPG, we introduce the fixed sample radius ratio δ_r with respect to the maximal sample radius of the policy. The selected sample radius adaptively shrinks along with the maximal sample radius as the policy gradually becomes more confident. It may also be advantageous to decrease the learning rate of both algorithms. Thus, we introduce hyperparameters α and A which control the rate of this decrease. The parameters ξ and α are exponents for geometric decay, which is a standard way to scale parameters like learning rates throughout an optimization algorithm, used in methods such as SPSA. The details of how these parameters enter can be found in the pseudocode of the algorithms.



3.3. Hyperparameter selection

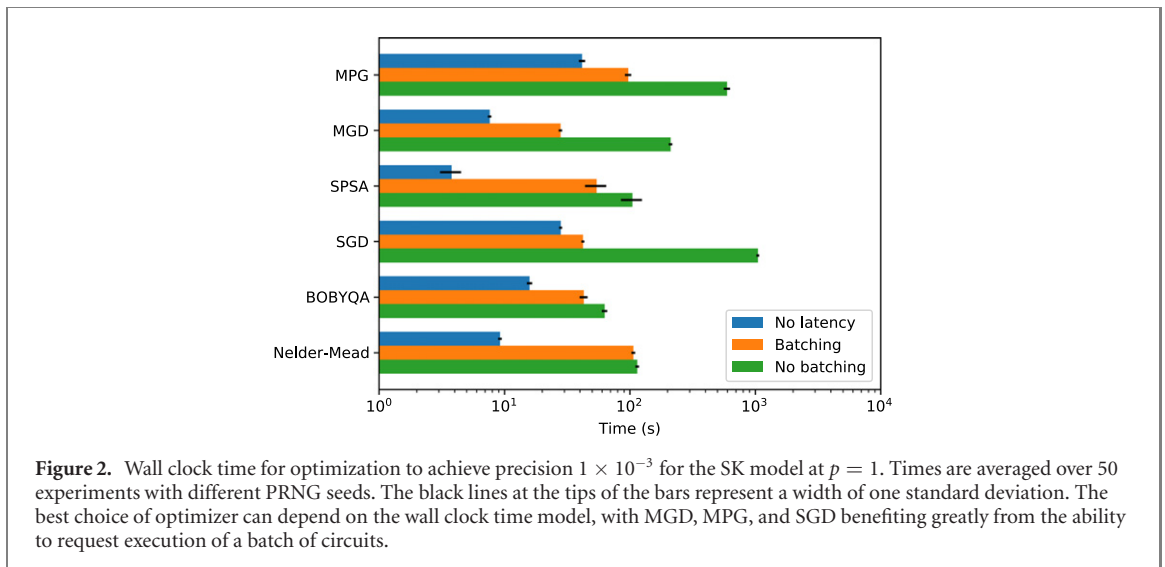
Each optimizer we considered here has a number of hyperparameters, and empirically we noted that the choice of these hyperparameters had a great impact on performance. Strikingly, some optimizers that failed completely with out of the box settings became competitive choices with even slight adjustments. Recalling that many of the optimizers we consider are inherently deterministic, one important hyperparameter external to all methods is the number of measurement shots per energy evaluation.

We tuned hyperparameters by grid search, and separately for each problem class and ansatz depth considered. For each combination of hyperparameters considered in the search, we performed an optimization run using the wall clock time model that includes network latency and circuit batching. The optimal hyperparameters were those that minimized time to convergence with a precision target of 10^{-3} . To avoid effects of overfitting, we restricted consideration to single realizations, where other runs are not further optimized within a problem class. Note that the details of hyperparameter selection has a significant effect on the performance of the algorithms. For example, choosing a more lenient precision requirement while still minimizing time to solution leads to different performance characteristics on other problems. See appendix B for more details, including descriptions of the hyperparameters.

As a simple demonstration of the importance of hyperparameter selection, we considered the performance on a simple test case with two different hyperparameter settings. Figure 1 shows the optimization progress of SPSA in simulated experiments on a SK model Hamiltonian with $n = 8$ and $p = 1$, using two different hyperparameter settings: the ones used by default in the implementation from the software package Qiskit [61], and ones that we optimized for minimal time to solution with a fixed precision cutoff. Depicted is the normalized energy versus wall clock time, using the wall clock time model that includes network latency and circuit batching. With tuned hyperparameters, SPSA converges to the solution rapidly, and without tuning it quite obviously does not. The erratic trajectory when using the unoptimized default parameters can be attributed to the fact that the initial learning rate of the algorithm is set to a value over 100 times larger than the optimized value. Hence, while SPSA is a powerful stochastic method capable of dealing with variable function noise, hyperparameter tuning must be actively used to make a proper comparison. Not taking advantage of this capability has led previous studies to underestimate the performance of SPSA or outright conclude that it is not effective for these problems [18, 24]. This demonstrates the importance of tuning hyperparameters in making a fair comparison between optimization algorithms, and throughout this study we tune all methods under consideration.

4. Results

To increase the applicability of our results to experiment, we consider both ideal and faulty operation of a quantum device. In the first case, in order to isolate challenges pertaining only to sampling noise, we assume an ideally functioning quantum computer, so that the only source of stochasticity in the objective function is finite sampling effects. In the other case, we modeled the effect of gate rotation error as follows: each time the optimizer queries the point θ , the objective function is evaluated at the point $\theta + \varepsilon$ instead,



where each component of ϵ is chosen from the normal distribution with mean 0 and standard deviation ϵ (for some gate error level ϵ). Since this error model does not straightforwardly translate to the calculation of gradients for SGD, we did not perform simulations of gate error with SGD. This error model is a simplified model of coherent control error, an important source of errors on actual hardware [1], and which is especially pertinent to the case of quantum computers accessed through cloud services which are used often but calibrated only periodically.

Each simulation we perform is characterized by four attributes: the problem (3-regular Max-Cut, SK, or Hubbard), the ansatz depth p , the choice of optimizer, and gate error level ϵ (possibly 0). For each set of attributes considered, we performed 50 statistically independent simulations. For each numerical simulation we performed, we estimate the wall-clock time of actually performing the experiment on a quantum computer accessed through a cloud service using the various cost models described in section 2.2, and set a limit to the total amount of time allowed. We are interested in how quickly a given optimization algorithm converges to the optimal energy to within a target precision. By ‘optimal energy’ we mean the energy of the ansatz state at the nearest local optimum as determined from a classical optimization of the noiseless objective function.

4.1. The case of $p = 1$ and no gate errors

First, we present the results of simulations with $p = 1$ and no gate errors. Figure 2 shows the wall clock time for different optimizers to achieve precision 10^{-3} for the SK model at $n = 8$ and $p = 1$. We define this time to be the earliest time at which the current and all future evaluated points have an approximation ratio or normalized energy close to the optimal value to within 10^{-3} . We show the results for the three different wall-clock models described in section 2.2: no network latency, network latency present but with circuit batching, and network latency present with no circuit batching. Note that Nelder–Mead converged in only 44 out of 50 runs; the other algorithms converged in all of them.

These results show that the proper choice of optimizer depends on the situation. SPSA performed the best under the wall clock time model with no latency, but was outperformed by MGD, SGD, and BOBYQA under the model that included latency and circuit batching. Under the model that included latency but did not have circuit batching, BOBYQA performed the best.

The importance of the wall clock time model, and in particular the effect of network latency, is evident. In the presence of network latency, MPG, MGD and SGD benefit much more from circuit batching than the other algorithms do. Both algorithms work by obtaining an estimate of the objective function gradient in each iteration. Circuit batching provides a benefit because multiple different circuits are needed to estimate the gradient, and these circuits can be sent over the network in one batch, reducing total network latency costs. SPSA also estimates the gradient, but it only uses 2 different circuits for that purpose. In contrast, the hyperparameters of MGD and MPG were chosen so that they both used 10, while SGD used 72. Indeed, the plot shows SGD benefiting from batching to a greater degree than MGD.

As an illustration of the ability of the various optimizers to tolerate different amounts of variance in the objective function, we note that the optimal hyperparameters dictates that SGD uses 1000 measurement shots per evaluations, MGD and SPSA use 5000, MPG uses 20 000, Nelder–Mead uses 25 000, and BOBYQA uses 125 000. This makes clear our distinction between deterministic and stochastic optimizers. While one

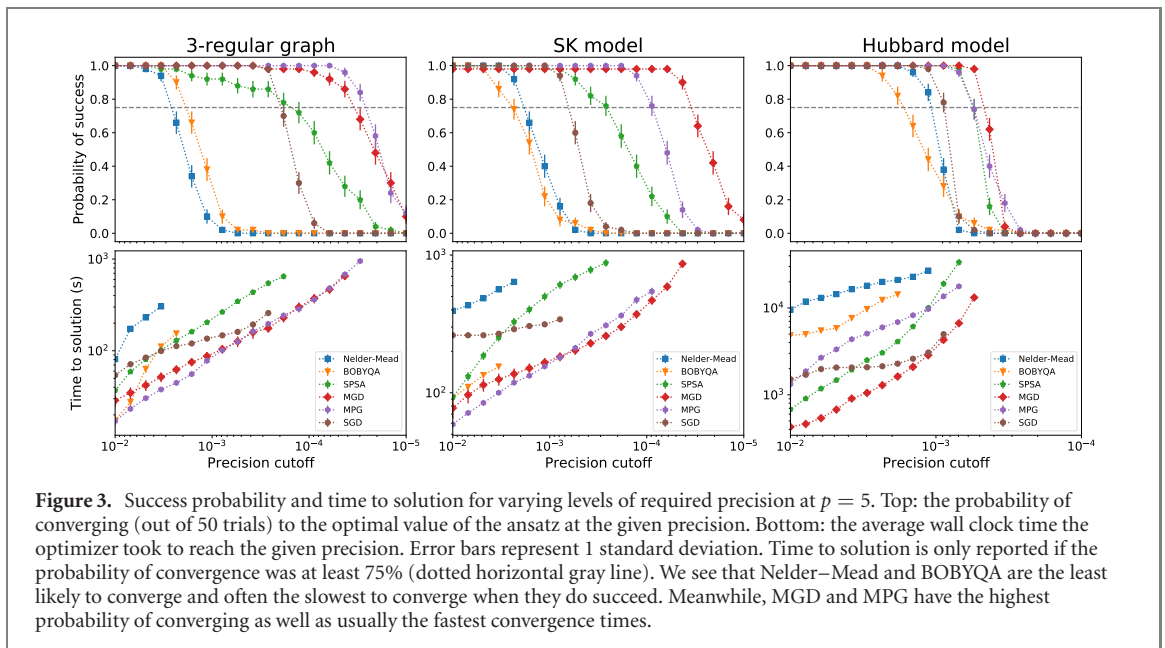


Figure 3. Success probability and time to solution for varying levels of required precision at $p = 5$. Top: the probability of converging (out of 50 trials) to the optimal value of the ansatz at the given precision. Bottom: the average wall clock time the optimizer took to reach the given precision. Error bars represent 1 standard deviation. Time to solution is only reported if the probability of convergence was at least 75% (dotted horizontal gray line). We see that Nelder–Mead and BOBYQA are the least likely to converge and often the slowest to converge when they do succeed. Meanwhile, MGD and MPG have the highest probability of converging as well as usually the fastest convergence times.

can find external hyperparameter settings that allow Nelder–Mead and BOBYQA to succeed, the lack of internal hyperparameters for noise tolerance means the number of measurements grows wildly. In contrast, stochastic methods like MPG, MGD and SPSA can find balanced settings using far fewer measurements per point while remaining stable. In larger systems beyond the scope of simulation, it may not be easy to *a priori* determine the required measurements to make a deterministic method stable, and hence the flexibility of naturally stochastic methods is likely to be preferred. For all cases, however, some amount of hyperparameter tuning is a necessity for good performance.

4.2. The case of $p = 5$ and no gate errors

At $p = 5$ there are a greater number of parameters to optimize. For the QAOA problems there are now 10 parameters, and for the Hubbard model there are 15. Here we fixed the wall clock time model to the one that includes network latency and circuit batching, and plot the performance of the optimizers as a function of the desired level of precision of convergence to the ansatz optimum. We present the results in figure 3. The optimizers did not always converge within the time limit we allowed (1500 s for the QAOA problems and 24 h for the Hubbard model). The top row depicts the probability of convergence to the desired precision, out of 50 runs. The bottom row depicts the average wall clock time for convergence, with data plotted only if the probability of convergence was at least 75%.

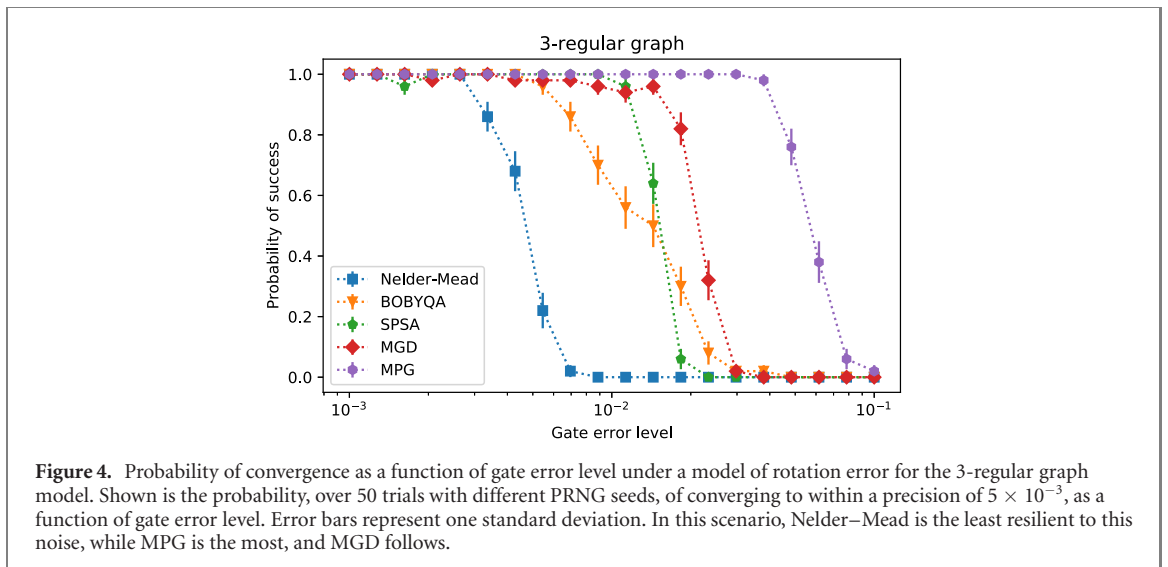
These simulations show that not only were Nelder–Mead and BOBYQA the least likely to converge; they were also often the slowest to converge when they did succeed. Meanwhile, MGD, MPG, and SPSA converged even at high levels of precision, with MGD and MPG consistently converging the most quickly in this regime. This is again a symptom of the fragility of using deterministic optimizers in a stochastic setting. Outside the regime of precise tuning, methods like Nelder–Mead and BOBYQA become unstable, whereas even outside the regime of tuning, methods like MGD, MPG, and SPSA are able to succeed.

Note that the plots would look different if we had tuned the hyperparameters with a different strategy. For example, we tuned the hyperparameters to minimize the time to convergence to a precision of 10^{-3} . If we had instead used a less precise cutoff, such as 10^{-2} , then we would expect the optimizers to converge faster to less precise cutoffs, but perhaps more slowly or less robustly to higher precision cutoffs at smaller ones. At a glance in these figures, one can see remnants of the hyperparameter selection cutoff. In appendix B we highlight this effect with an example.

4.3. The impact of rotation errors at $p = 5$

Finally, to understand the impact of gate error in addition to simple sampling noise, at $p = 5$ we consider gate rotation errors as well. As described above, the model of gate rotation error that we used does not simply translate to SGD, so we do not include results for it. Again, we fixed the wall clock time model to the one that includes network latency and circuit batching. In running the optimization algorithms, we used the hyperparameters that were optimized for the case of no gate errors.

Figure 4 shows the probability of convergence to a precision of 5×10^{-3} for the various optimizers as a function of the gate error level ε , for the 3-regular graph model. The results show that in this scenario,



Nelder–Mead is the least resilient to this type of noise, while MPG performs the best and MGD follows. The reason why MPG is particularly robust to noise is because it learns a stochastic policy. Its probability-based optimization minimizes the objective function in the expectation sense and thus manages to handle various levels of uncertainty. SPSA also showed good noise resilience in other scenarios; see section D in the appendix for data for the other models.

Note that for a given gate error level, algorithmic improvements can increase the success probability with respect to the ideal solution only up to a certain point. That is, beyond a certain level of noise, the device cannot produce a more precise solution, and hence this is not a failing of the optimizer but rather represents a device limitation. We do not differentiate between these circumstances in the presented data, but merely note that it is a consideration when defining probability of success.

5. Conclusion

Variational quantum algorithms are a promising candidate for execution on near-term quantum computers, and a number of experimental demonstrations of these algorithms have already been performed. These algorithms rely on a classical optimization subroutine, and hence the efficiency of these algorithms can be limited by the performance of these optimizers. Here, we saw that to accurately assess the performance of these optimizers, it is crucial to develop a good cost model, and tune available hyperparameters to operational specifications.

Given the unique considerations of quantum systems, we developed two new surrogate model-based optimizers, MGD and MPG, to fill some of the gaps of previous methods. We numerically compared their performance with other popular alternatives, and found it advantageous in several realistic settings. We also probed how the cost model and presence of errors can significantly impact the choice of optimizer in a practical setting.

Now that quantum computers are coming online, accessing superconducting qubits through a cloud interface is an important scenario to consider. The latency of communicating over the Internet can cause large increases in running times, but this can be mitigated by circuit batching, though the cost savings depends on the optimizer.

We also observed that inherently stochastic optimizers, such as MPG, MGD and SPSA, were more robust to variations in problems or setting once properly tuned. This extended to situations where finite gate or circuit noise was present. In contrast, while it was sometimes possible to make deterministic optimizers competitive through careful tuning, these tunings were fragile with respect to small variations in the problem or the introduction of noise. Overall, MPG and MGD’s tolerance of noise, ability to take advantage of circuit batching, and good overall performance make them good candidates for actual experiments, but the best optimizer can depend on the processor’s wall-clock model, level of noise, number of parameters, or the specific circuit ansatz.

In this work, we have shown how practical considerations can significantly affect the calculus of choosing an optimizer for running variational algorithms. Future work will develop more accurate noise and cost models, and further development of optimizers can take these unique considerations into account.

Acknowledgments

We thank Eddie Farhi and Bill Huggins for helpful discussions. This work was partially supported by the Department of Energy under Grant No. DE-SC0017867 (JY, LL) and a Google Quantum Research Award (LL).

Code Availability

Implementations of MGD and MPG are available at <https://github.com/quantumlib/ReCirq>.

ORCID iDs

Kevin J Sung  <https://orcid.org/0000-0001-6459-6374>

Nicholas C Rubin  <https://orcid.org/0000-0003-3963-1830>

References

- [1] Arute F *et al* 2019 Quantum supremacy using a programmable superconducting processor *Nature* **574** 505–10
- [2] 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** 1–7
- [3] 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
- [4] Preskill J 2018 Quantum computing in the NISQ era and beyond *Quantum* **2** 79
- [5] Hempel C *et al* 2018 Quantum chemistry calculations on a trapped-ion quantum simulator *Phys. Rev. X* **8** 031022
- [6] Nelder J A and Mead R 1965 A simplex method for function minimization *Comput. J.* **7** 308–13
- [7] 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** 242–6
- [8] Otterbach J S *et al* 2017 Unsupervised machine learning on a hybrid quantum computer (arXiv:1712.05771 [quant-ph])
- [9] Colles J I, Ramasesh V V, Dahlen D, Blok M S, Kimchi-Schwartz M E, McClean J R, Carter J, de Jong W A and Siddiqi I 2018 Computation of molecular spectra on a quantum processor with an error-resilient algorithm *Phys. Rev. X* **8** 011021
- [10] Kokail C *et al* 2019 Self-verifying variational quantum simulation of lattice models *Nature* **569** 355–60
- [11] Pagano G *et al* 2019 Quantum approximate optimization of the long-range Ising model with a trapped-ion quantum simulator (arXiv:1906.02700 [quant-ph])
- [12] Spall J C 1992 Multivariate stochastic approximation using a simultaneous perturbation gradient approximation *IEEE Trans. Autom. Control* **37** 332–41
- [13] Shahriari B, Swersky K, Wang Z, Adams R P and de Freitas N 2016 Taking the human out of the loop: a review of Bayesian optimization *Proc. IEEE* **104** 148–75
- [14] Parsopoulos K E and Vrahatis M N 2002 Recent approaches to global optimization problems through particle swarm optimization *Nat. Comput.* **1** 235–306
- [15] Jones D R, Perttunen C D and Stuckman B E 1993 Lipschitzian optimization without the Lipschitz constant *J. Optim. Theory Appl.* **79** 157–81
- [16] Wecker D, Hastings M B and Troyer M 2015 Progress towards practical quantum variational algorithms *Phys. Rev. A* **92** 042303
- [17] Zhou L, Wang S-T, Choi S, Pichler H and Lukin M D 2018 Quantum approximate optimization algorithm: performance, mechanism, and implementation on near-term devices (arXiv:1812.01041 [quant-ph])
- [18] Nakanishi K M, Fujii K and Todo S 2019 Sequential minimal optimization for quantum-classical hybrid algorithms (arXiv:1903.12166 [quant-ph])
- [19] Parrish R M, Iosue J T, Ozaeta A and McMahon P L 2019 A Jacobi diagonalization and Anderson acceleration algorithm for variational quantum algorithm parameter optimization (arXiv:1904.03206 [quant-ph])
- [20] Kübler J M, Arrasmith A, Cincio L and Coles P J 2019 An adaptive optimizer for measurement-frugal variational algorithms (arXiv:1909.09083 [quant-ph])
- [21] Arrasmith A, Cincio L, Somma R D and Coles P J 2020 Operator sampling for shot-frugal optimization in variational algorithms (arXiv:2004.06252 [quant-ph])
- [22] Guerreschi G G and Smelyanskiy M 2017 Practical optimization for hybrid quantum-classical algorithms (arXiv:1701.01450 [quant-ph])
- [23] Romero J, Babbush R, McClean J, Hempel C, Love P and Aspuru-Guzik A 2018 Strategies for quantum computing molecular energies using the unitary coupled cluster ansatz (arXiv:1701.02691 [quant-ph])
- [24] Nannicini G 2019 Performance of hybrid quantum-classical variational heuristics for combinatorial optimization *Phys. Rev. E* **99** 013304
- [25] Yao J, Bukov M and Lin L 2020 Policy gradient based quantum approximate optimization algorithm (arXiv:2002.01068 [quant-ph])
- [26] Lavrijsen W, Tudor A, Mäijller J, Iancu C and de Jong W 2020 Classical optimizers for noisy intermediate-scale quantum devices (arXiv:2004.03004 [quant-ph])
- [27] Leng Z, Mundada P, Ghadimi S and Houck A 2019 Robust and efficient algorithms for high-dimensional black-box quantum optimization (arXiv:1910.03591 [quant-ph])
- [28] Stokes J, Izaac J, Killoran N and Carleo G 2020 Quantum natural gradient *Quantum* **4** 269
- [29] Wierichs D, Gogolin C and Kastoryano M 2020 Avoiding local minima in variational quantum eigensolvers with the natural gradient optimizer (arXiv:2004.14666 [quant-ph])

- [30] Byrd R H, Lu P, Nocedal J and Zhu C 1995 A limited memory algorithm for bound constrained optimization *SIAM J. Sci. Comput.* **16** 1190–208
- [31] Powell M J D 1994 A direct search optimization method that models the objective and constraint functions by linear interpolation *Advances in Optimization and Numerical Analysis* ed S Gomez and J-P Hennart (Berlin: Springer) pp 51–67
- [32] Powell M J D 1964 An efficient method for finding the minimum of a function of several variables without calculating derivatives *Comput. J.* **7** 155–62
- [33] Costa A and Nannicini G 2018 RBFOpt: an open-source library for black-box optimization with costly function evaluations *Math. Prog. Comp.* **10** 597–629
- [34] Huyer W and Neumaier A 2008 SNOBFIT—stable noisy optimization by branch and fit *ACM Trans. Math. Softw.* **35** 9
- [35] Powell M J D 2009 The BOBYQA algorithm for bound constrained optimization without derivatives
- [36] Le Digabel S 2011 Algorithm 909: Nomad: nonlinear optimization with the mads algorithm *ACM Trans. Math. Softw.* **37** 44
- [37] Kelley C T 2011 *Implicit Filtering* (Philadelphia, PA: SIAM)
- [38] Williams R J 1992 Simple statistical gradient-following algorithms for connectionist reinforcement learning *Mach. Learn.* **8** 229–56
- [39] Farhi E, Goldstone J and Gutmann S 2014 A quantum approximate optimization algorithm (arXiv:1411.4028 [quant-ph])
- [40] Arute F et al 2020 Quantum approximate optimization of non-planar graph problems on a planar superconducting processor (arXiv:2004.04197 [quant-ph])
- [41] Brandao F G S L, Broughton M, Farhi E, Gutmann S and Neven H 2018 For fixed control parameters the quantum approximate optimization algorithm’s objective function value concentrates for typical instances (arXiv:1812.04170 [quant-ph])
- [42] Sherrington D and Kirkpatrick S 1975 Solvable model of a spin-glass *Phys. Rev. Lett.* **35** 1792–6
- [43] Hubbard J 1963 Electron correlations in narrow energy bands *Proc. R. Soc. A* **276** 238–57
- [44] Dagotto E 1994 Correlated electrons in high-temperature superconductors *Rev. Mod. Phys.* **66** 763–840
- [45] Kivlichan I D, McClean J, Wiebe N, Gidney C, Aspuru-Guzik A, Kin-Lic Chan G and Babbush R 2018 Quantum simulation of electronic structure with linear depth and connectivity *Phys. Rev. Lett.* **120** 110501
- [46] Jordan P and Wigner E 1928 über das Paulische äquivalenzverbot *Z. Phys.* **47** 631–51
- [47] Jena A, Genin S and Mosca M 2019 Pauli partitioning with respect to gate sets (arXiv:1907.07859 [quant-ph])
- [48] Izmaylov A F, Yen T-C and Ryabinkin I G 2019 Revising the measurement process in the variational quantum eigensolver: is it possible to reduce the number of separately measured operators? *Chem. Sci.* **10** 3746–55
- [49] Huggins W J, McClean J, Rubin N, Jiang Z, Wiebe N, Birgitta Whaley K and Babbush R 2019 Efficient and noise resilient measurements for quantum chemistry on near-term quantum computers (arXiv:1907.13117 [quant-ph])
- [50] Izmaylov A F, Yen T-C, Lang R A and Verteletskyi V 2020 Unitary partitioning approach to the measurement problem in the variational quantum eigensolver method *J. Chem. Theory Comput.* **16** 190–5
- [51] Verteletskyi V, Yen T-C and Izmaylov A F 2020 Measurement optimization in the variational quantum eigensolver using a minimum clique cover *J. Chem. Phys.* **152** 124114
- [52] Mitarai K, Negoro M, Kitagawa M and Fujii K 2018 Quantum circuit learning *Phys. Rev. A* **98** 032309
- [53] Schuld M, Bergholm V, Gogolin C, Izaac J and Killoran N 2019 Evaluating analytic gradients on quantum hardware *Phys. Rev. A* **99** 032331
- [54] Crooks G E 2019 Gradients of parameterized quantum gates using the parameter-shift rule and gate decomposition (arXiv:1905.13311 [quant-ph])
- [55] Harrow A and Napp J 2019 Low-depth gradient measurements can improve convergence in variational hybrid quantum-classical algorithms (arXiv:1901.05374)
- [56] Barends R et al 2014 Superconducting quantum circuits at the surface code threshold for fault tolerance *Nature* **508** 500
- [57] Córcoles A D, Magesan E, Srinivasan S J, Cross A W, Steffen M, Gambetta J M and Chow J M 2015 Demonstration of a quantum error detection code using a square lattice of four superconducting qubits *Nat. Commun.* **6** 6979
- [58] Jones D R 2001 A taxonomy of global optimization methods based on response surfaces *Journal of Global Optimization* **21** 345–83
- [59] Cartis C, Fiala J, Marteau B and Roberts L 2019 Improving the flexibility and robustness of model-based derivative-free optimization solvers *ACM Trans. Math. Softw.* **45** 32
- [60] Virtanen P et al 2020 SciPy 1.0: fundamental algorithms for scientific computing in Python *Nat. Methods* **17** 261–72
- [61] Abraham H et al 2019 Qiskit: an open-source framework for quantum computing doi: 10.5281/zenodo.2562110
- [62] Kingma D P and Ba J 2014 Adam: a method for stochastic optimization (arXiv:1412.6980 [cs.LG])