

**Automatic depth optimization for a quantum approximate optimization algorithm**Yu Pan<sup>✉\*</sup> and Yifan Tong<sup>†</sup>*State Key Laboratory of Industrial Control Technology, Institute of Cyber-Systems and Control,  
College of Control Science and Engineering, Zhejiang University, Hangzhou 310027, China*Yi Yang<sup>‡</sup>*Department of Mathematics and Statistics, McGill University, Montreal, Quebec, Canada H3A 0B9*

(Received 24 October 2021; revised 11 January 2022; accepted 7 March 2022; published 21 March 2022)

Quantum approximate optimization algorithm (QAOA) is a hybrid algorithm whose control parameters are classically optimized. In addition to the variational parameters, the right choice of hyperparameter is crucial for improving the performance of any optimization model. Control depth, or the number of variational parameters, is considered as the most important hyperparameter for QAOA. In this paper we investigate the control depth selection with an automatic algorithm based on proximal gradient descent. The performances of the automatic algorithm are demonstrated on seven-node and ten-node max-cut problems, which show that the control depth can be significantly reduced during the iteration whereas achieving an sufficient level of optimization accuracy. With theoretical convergence guarantee, the proposed algorithm can be used as an efficient tool for choosing the appropriate control depth as a replacement of random search or empirical rules. Moreover, the reduction of control depth will induce a significant reduction in the number of quantum gates in circuit, which improves the applicability of QAOA on noisy intermediate-scale quantum devices.

DOI: [10.1103/PhysRevA.105.032433](https://doi.org/10.1103/PhysRevA.105.032433)**I. INTRODUCTION**

Quantum approximate optimization algorithm (QAOA) is considered as one of the most promising applications for near-term noisy intermediate scale quantum (NISQ) [1] computing device to demonstrate quantum advantage [2–4]. The goal of QAOA is to prepare a quantum state that yields an approximate solution to a classical optimization problem. QAOA is a hybrid quantum-classical algorithm, in the sense that the quantum state is generated and measured using quantum hardware whereas the control parameters are optimized by the classical algorithm to form a closed loop. In the standard QAOA, the quantum state is generated by executing  $p$  blocks of noncommuting quantum operations which consist of  $2p$ -variational control parameters in total. The control parameters can be optimized by gradient-free [2,5–7], gradient-based [6,8–10], and machine learning methods [11–15]. In this case, the control depth is  $2p$ , which is a hyperparameter that has a strong influence on the performance of the model.

Although it has been proven that QAOA converges to the optimal solution in the  $p \rightarrow \infty$  limit [2], an extremely large  $p$  is not physically realizable due to the noise effect and limited control capability of NISQ devices. In practice, the control depth has a finite value which is often preselected. However, although there have been studies on the dependence of QAOA performances upon the control or circuit depth, it is

still not clear how to determine an optimal depth with respect to any specific problem [16–20]. The current research (e.g., Refs. [2,8,20–23]) have found several empirical or analytical rules to select the control depth for certain problems, which cannot be formulated as automated algorithms. In addition, considering that the empirical or analytical selection rules are different for different problems or even different class of instances, a generic selection rule seems unlikely to exist. However, searching for a good control depth by hand tuning is computationally inefficient. In particular, none of the current works have investigated any automatic algorithm framework for balancing model accuracy and model complexity with an intermediate control depth, which is critical for the robust implementation of QAOA.

This paper presents an attempt to derive a generic and automatic algorithm for optimizing the control depth of QAOA, which is more efficient than random search and more generally applicable than existing empirical or analytical selection rules. Since the control depth is a hyperparameter of the model, the depth optimization can be taken as a model selection problem. Therefore, any model selection method [24,25] can be considered to solve this problem. In this paper, we employ the model selection method with  $l_1$  regularization (by additionally minimizing the  $l_1$  norm of a parameter vector) [26] for mainly two reasons. First, the regularized model can be optimized iteratively, which makes it very efficient. For example, regularization techniques, such as least absolute shrinkage and selection operator (LASSO) [26], and its variants are often used for the automatic model selection. LASSO imposes an  $l_1$  regularization on the parameters to be estimated in addition to the objective function, which can effectively

\*ypan@zju.edu.cn

†yftong@zju.edu.cn

‡yi.yang6@mcgill.ca

shrink the number of parameters of the model during iteration. Second, an  $l_1$ -regularized optimization problem can be solved by fast algorithms with optimality and convergence guarantee. In particular, the commonly used proximal gradient (PG) descent method can be used to solve a linear and convex problem with a basic convergence rate of  $O(1/k)$  [27] which can be further accelerated by a line search [28], where  $k$  is the current iteration step. Since the objective function in QAOA is possibly nonconvex [17], we also have to consider an extension of PG descent algorithm that is compatible with both convex and nonconvex problems [29,30]. Other regularization terms can also be used as candidates if certain assumptions on the correlations between the parameters are made [31].

The most common criterion for model selection is to achieve the balance between model accuracy and complexity [24,25]. It should be noted that for the noise-free case in this paper, the accuracy of the model will be monotonically increasing with the control depth. As a result, a criterion that balances the accuracy and control depth has to be proposed, otherwise, the best model would be the unregularized one if accuracy is the sole consideration. Fortunately, the necessity of complexity control naturally arises in the context of QAOA. To be more specific, the objective for QAOA can be stated as optimizing the approximation ratio until a desired value has been reached, which has been proven to be NP-hard with classical algorithms [17]. In other words, we only require QAOA to give an optimization result that is good enough instead of the best result. In this case, the optimization target can be expressed as

$$\frac{C^*}{C_{\max}} \geq r^*, \quad (1)$$

where  $C^*$  is the optimized value of the objective function and  $C_{\max}$  is the maximum value.  $C_{\max}$  can be obtained by brute force methods across different instances [32] or enforcing a frustration-free Hamiltonian whose ground-state energy is known in advance [33]. Here  $C^*/C_{\max}$  stands for the approximation ratio. According to (1), the approximation ratio measures the optimization accuracy. An approximate solution  $C^*$  is smaller than  $C_{\max}$ , and QAOA intends to produce a solution that is as close to  $C_{\max}$  as possible. Finding the minimum depth that satisfies (1) with a good-enough  $r^*$  naturally confines the control complexity and is crucial for the understanding of the fundamental performance limitations of QAOA [34,35], which can further contribute to a better design that mitigates the circuit complexity in the practical implementation.

In this paper, we provide a comprehensive study of the performance of PG descent methods for the control depth optimization of QAOA. Numerical results show that the performance of QAOA is heavily affected by the control depth, which can be significantly optimized during the iteration using a proper regularization strength. This paper is organized as follows. A short introduction on QAOA is given in Sec. II. The algorithms for control depth optimization are detailed in Sec. III. In Sec. IV, the max-cut problem and metrics for numerical experiment are defined. Section V presents the numerical results on seven-node and ten-node examples. The conclusion is summarized in Sec. VI.

## II. STANDARD QAOA

The Pauli matrices are defined by

$$\sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \quad \sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}. \quad (2)$$

The two computational basis states of a single qubit are denoted as  $|0\rangle = (0 \ 1)^T$  and  $|1\rangle = (1 \ 0)^T$ . Let  $H_o$  be the Hamiltonian that encodes the corresponding optimization problem, and  $H_c = \sum_{n=1}^N \sigma_x^{(n)}$  be the control Hamiltonian. The optimal solution to the classical problem, which is often a combinatorial problem, is encoded as the ground state of  $H_o$ . The standard QAOA is executed by alternatively applying the noncommuting  $H_o$  and  $H_c$  as

$$|\psi(\beta, \gamma)\rangle = e^{-iH_c\beta_p} e^{-iH_o\gamma_p} \dots e^{-iH_c\beta_1} e^{-iH_o\gamma_1} |s\rangle. \quad (3)$$

Here  $|s\rangle = |+\rangle_1 \dots |+\rangle_N$  is the initial  $N$ -qubit state that can be easily prepared with  $|+\rangle = 1/\sqrt{2}(|0\rangle + |1\rangle)$ . The parameters  $\beta = (\beta_1, \dots, \beta_p)$ ,  $\gamma = (\gamma_1, \dots, \gamma_p)$  stand for control angles constrained in the region  $(-\pi, \pi]$ , and  $2p$  is the control depth. A negative parameter corresponds to the angle of a backward rotation. The goal of QAOA can be formulated as

$$\min_{\beta, \gamma} f(\beta, \gamma) = \min_{\beta, \gamma} \langle \psi(\beta, \gamma) | H_o | \psi(\beta, \gamma) \rangle. \quad (4)$$

That is, the goal is to optimize  $\beta, \gamma$  such that the generated quantum state  $|\psi(\beta, \gamma)\rangle$  approximates the ground state of  $H_o$ . In QAOA, (3) is implemented using quantum hardware. In practice, each unitary operation defined in (3) is realized as a sequence of quantum gates. Moreover,  $f(\beta, \gamma)$  in (4) is the expectation value of  $H_o$  with respect to state  $|\psi(\beta, \gamma)\rangle$ , which can be calculated by repeatedly executing (3) and taking the average of the quantum measurement results with respect to  $H_o$ . In contrast, the optimization algorithm for the control parameters  $\beta, \gamma$  can be designed and implemented classically.

## III. ALGORITHMS FOR DEPTH SELECTION

In this paper, we adopt the  $l_1$  regularization technique for the automatic reduction of control depth. The regularized model is given by

$$\begin{aligned} \min_{\beta, \gamma} f(\beta, \gamma) + g(\beta, \gamma) &= \min_{\beta, \gamma} \langle \psi(\beta, \gamma) | H_o | \psi(\beta, \gamma) \rangle \\ &+ \lambda \sum_{j=1}^p (|\beta_j| + |\gamma_j|). \end{aligned} \quad (5)$$

The effect of the additional regularization term  $g(\beta, \gamma)$  is shrinking some of the control parameters to be exactly zero during optimization. Letting  $x = (\beta, \gamma)$ , (5) can be written in standard form as

$$\min_x \langle \psi(x) | H_o | \psi(x) \rangle + \lambda \|x\|_1, \quad (6)$$

where  $\|\cdot\|_1$  denotes the  $l_1$  norm of a vector. Note that  $f(\cdot)$  is differentiable but possibly nonconvex [17], whereas the regularization term  $g(\cdot)$  is convex but nondifferentiable.

The PG descent method updates the control parameters as

$$\begin{aligned} x_{k+1} &= \arg \min_z f(x_k) + \nabla f(x_k)^T (z - x_k) + \frac{1}{2\eta} \|z - x_k\|_2^2 \\ &+ g(z). \end{aligned} \quad (7)$$

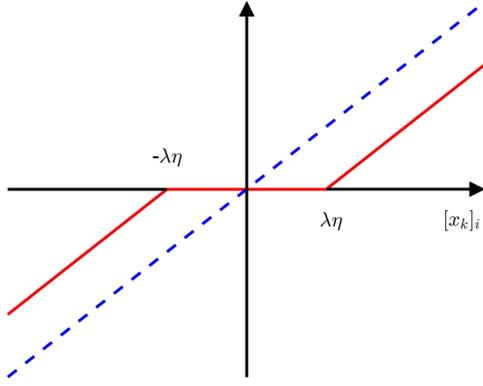


FIG. 1. The illustration of a soft-thresholding operation with the y-axis value determines the value of parameter after thresholding. The blue dotted line indicates the reference case without soft thresholding. During each iteration, the updated control parameter is then shifted towards zero by an amount of  $\lambda\eta$  (red line). If the absolute value of the control parameter is below the threshold, then it will be directly penalized to zero (red line).

The motivation is to minimize the quadratic approximation to  $f(\cdot)$  around  $x_k$  with  $\nabla^2 f(\cdot)$  replaced by  $\frac{1}{\eta}$  and leave the nondifferentiable  $g(\cdot)$  alone. Here  $\eta$  is the learning rate. For  $g(\cdot) = \lambda \|\cdot\|_1$ , the minimization problem (7) has the following closed-form solution,

$$x_{k+1} = S_{\lambda\eta}[x_k - \eta \nabla f(x_k)]. \tag{8}$$

Here  $S_{\lambda\eta}(x_k)$  is the soft-thresholding operator defined by

$$S_{\lambda\eta}([x_k]_i) = \begin{cases} [x_k]_i - \lambda\eta, & [x_k]_i > \lambda\eta \\ 0, & -\lambda\eta \leq [x_k]_i \leq \lambda\eta \\ [x_k]_i + \lambda\eta, & [x_k]_i < -\lambda\eta. \end{cases} \tag{9}$$

where  $[x_k]_i$  is the  $i$ th element of the vector  $x_k$  for  $i = 1, \dots, 2p$ . As shown in Fig. 1, once the condition  $|[x_k]_i| \leq \lambda\eta$  is met,  $[x_k]_i$  will be penalized to exactly zero, which will remove the corresponding control action and reduce the control depth. The gradient vector  $\nabla f(x_k)$  is approximated by

$$[\nabla f(x_k)]_i \approx \frac{f([x_k]_i + \epsilon) - f([x_k]_i - \epsilon)}{2\epsilon}, \quad \epsilon > 0, \tag{10}$$

in which  $f([x_k]_i \pm \epsilon)$  can be estimated using either classical simulation or quantum hardware. It should be noted that the stop criterion is to continue the iterations until a good-enough threshold value has been reached and, thus, whether the expectations are obtained by classical simulation or quantum measurement would not matter. Particularly, the gradient descent update did not use any information of the underlying dynamical model. Therefore, the only consequence brought by realistic noise is that the achieved approximation ratio may be lower than the noise-free case, and gradient estimates may be noisy.

Generally speaking, the convergence rate of  $O(1/k)$  cannot be guaranteed in the absence of the convexity assumption on  $f(\cdot)$ . Nevertheless, since PG descent is a first-order optimization method just as the stochastic gradient descent, its convergence curve is still expected to be stable. In addition, as can be seen from (8), PG descent is simply a reweighting of the conventional gradient descent, which is unlikely to get

unstable. Due to these reasons, the stability of the convergence curves can be clearly observed in the numerical results of this paper.

**Algorithm 1.** Nonconvex accelerated proximal gradient (APG) algorithm for QAOA.

---

```

Require:  $f(\cdot)$  is  $L$ -Lipschitz smooth.
Initialize:  $\eta \in (0, 1/L)$ ;  $x_0 = x_1$ ;  $\lambda, \text{tol} > 0$ ; a positive integer  $q$ .
1: for  $k = 1, \dots, K$  do
2:  $y_k = x_k + \frac{k-1}{k+2}(x_k - x_{k-1})$ ;  $\backslash\backslash$  line search for possible
   acceleration
3:  $F_k = \max_{t \in \{\max(1, k-q), \dots, k\}} [f(x_t) + g(x_t)]$ ;
4: if  $f(y_k) + g(y_k) \leq F_k$  then
5:  $\backslash\backslash$  check if  $y_k$  is a good extrapolation
6:  $v_k = y_k$ ;
7: else
8:  $v_k = x_k$ ;
9: end if
10:  $\nabla f(v_k)$  is calculated using classical simulation or quantum
   hardware;
11:  $x_{k+1} = S_{\lambda\eta}[v_k - \eta \nabla f(v_k)]$ ;  $\backslash\backslash$  PG update
12:  $f(x_{k+1})$  is calculated using classical simulation or quantum
   hardware;
13: if  $|f(x_{k+1}) + g(x_{k+1}) - F_k| < \text{tol}$  then
14: return  $x_k$ ;  $\backslash\backslash$  early stopping if convergence has been achieved
15: end if
16: end for
17: return  $x_{K+1}$ ;

```

---

We also consider a second algorithm which does not require  $f(\cdot)$  to be convex and has a proved convergence rate. This algorithm is an efficient APG descent method proposed in Ref. [30]. We can use this unified approach for both convex and nonconvex problems to construct the classical optimization part in QAOA. The details are shown in Algorithm 1. It should be noted that the extrapolation defined by step 2 can accelerate the convergence if  $f(\cdot)$  is locally convex. Steps 3–9 are designed to check if the acceleration works. In step 3, choosing  $q > 0$  allows  $y_k$  to occasionally increase the objective function  $f(\cdot)$  in order to escape from local minima [30,36]. In particular,  $q$  is set to be 2 in the numerical experiments of this paper. Although the only difference between APG and PG appears to be the acceleration part, the convergence rate of APG is shown to be  $O(1/k)$  for nonconvex  $f(\cdot)$  [30].

The regularization parameter  $\lambda$  in (6) is used to control the model complexity. A large  $\lambda$  forces most of the elements in  $x$  to be zero, whereas a very small  $\lambda$  results in a model that is close to an unregularized one. For this reason,  $\lambda$  has to be carefully chosen to encourage a simple control model with sparse  $x$  whereas achieving a sufficient level of optimization accuracy. The common practice is to conduct a grid search to find the optimal  $\lambda$ . In order to mitigate the resource requirement and running time of the algorithm, the candidate hyperparameters can be examined in the following order:

$$\lambda_1 > \lambda_2 > \dots > \lambda_M. \tag{11}$$

That is, the algorithms start with a relatively large  $\lambda$ , which will produce the most sparse control model. Then the

procedures are repeated by decreasing  $\lambda$  until a predetermined optimization accuracy has been achieved. Since the regularization strength cannot be too large or too small, only a few  $\lambda$ s have to be tested before an acceptable result has been obtained, irrespective of problem size. In contrast, the number of experiments needed for finding an optimized depth with random search scales as  $O(p)$ . Therefore, the automatic algorithm is very efficient when dealing with large-scale problems where we have  $M \ll p$ .

It should be noted that for the noise-free QAOA considered in this paper, the optimization accuracy always tends to increase with the control depth [2,37], which means we can always improve the optimization accuracy by decreasing  $\lambda$ . In that case, the model will become overly complex, or the control depth will become too large. As we have stated, the model complexity can be controlled by selecting the largest  $\lambda$  that satisfies the criterion (1).

#### IV. MAX CUT AND APPROXIMATION RATIO

As a proof-of-principle demonstration, the performance of depth optimization is demonstrated on the max-cut problem [3]. Consider an  $N$ -node nondirected and weighted graph  $G = (V, E)$ . Max cut seeks for a partition of  $V$  into two subsets  $S_1$  and  $S_2$  such that the sum of weights of edges connecting nodes in the two disjoint subsets is maximized. By assigning  $Z_i = 1, i \in S_1$  and  $Z_j = -1, j \in S_2$ , the max-cut problem can be formulated as a binary optimization,

$$C_{\max} = \max_{i,j} \frac{1}{2} \sum_{(i,j) \in E} \omega_{ij} (1 - Z_i Z_j). \quad (12)$$

It is known that finding an approximate solution of binary string  $Z^*$  such that

$$\frac{C(Z^*)}{C_{\max}} \geq r^* \quad (13)$$

is NP-hard, with  $r^*$  being a minimum ratio [17,38]. In the QAOA setup, the problem-based Hamiltonian for (12) is given by

$$H_{\mathcal{O}} = \frac{1}{2} \sum_{(i,j) \in E} \omega_{ij} (I - \sigma_{z_i} \sigma_{z_j}). \quad (14)$$

The approximation ratio for benchmarking the performance of the QAOA can be defined as

$$r' = \frac{\max_{\beta, \gamma} \langle \psi(\beta, \gamma) | H_{\mathcal{O}} | \psi(\beta, \gamma) \rangle}{C_{\max}}. \quad (15)$$

Particularly, maximization the expectation of  $H_{\mathcal{O}}$  is equivalent to the minimization of expectation of the following

Hamiltonian:

$$H_o = \sum_{(i,j) \in E} \omega_{ij} \sigma_{z_i} \sigma_{z_j}, \quad (16)$$

for which the approximation ratio can be defined by

$$r = 1 - \frac{\min_{\beta, \gamma} \langle \psi(\beta, \gamma) | H_o | \psi(\beta, \gamma) \rangle - C_{\min}}{C_{\max} - C_{\min}} \in [0, 1], \quad (17)$$

with  $C_{\min}$  being the theoretical best result. According to (17), the approximation ratio  $r$  measures the optimization accuracy. In particular,  $r = 1$  implies that a perfect solution has been obtained by QAOA.

#### V. NUMERICAL RESULTS

Two randomly generated instances of max cut for the numerical illustration are given in Table I. The seven-node instance has 9 weighted edges, and the ten-node instance has 14 weighted edges. The code and data for the experiments are available in the Supplemental Material [39].

##### A. Seven-node max cut

First, we demonstrate the performance of depth optimization on the seven-node max cut problem. We let  $\eta = 0.006$ . The elements of  $\beta$  and  $\gamma$  are all initialized as 0.3 with an initial control depth of 14. The minimum acceptable approximation ratio is set to be 0.9, which is drawn in red dashed line in Fig. 2. The regularization parameter has four values  $\{2, 1.2, 0.72, 0.432\}$ , which are tested in decreasing order. Note that the effective range of regularization parameter can be estimated based on the range of the objective function and the initial values of the control parameters as the regularization term has to be comparable with the objective function for the optimization to be successful.

As can be seen in Fig. 2, the depth selection results are consistent for APG and PG with the same  $\lambda$ . More specifically, the final control depths are basically the same with the same  $\lambda$  after 200 iterations. Also, it has been shown that if  $\lambda$  is too large, the algorithms cannot reach the minimum acceptable ratio. The final accuracy is mainly determined by the control depth. That is, the approximation ratios achieved for different algorithms are very close under the same depth, which highlights the importance of depth optimization.

In particular, acceleration in convergence can be observed in Fig. 2(a). APG converges in less than 50 iterations, whereas PG may need more than 100 iterations to converge. APG achieves the minimum acceptable ratio in about 40 iterations with  $\lambda = 0.72$ , and the control depth has been shrunk to 11 at that point. Note that the removal of one control parameter

TABLE I. Instances of seven-node and ten-node max cut.

	Vertex	Weight												
Seven-node max cut	$Z_0 Z_4$	0.73	$Z_0 Z_6$	0.50	$Z_1 Z_5$	0.36	$Z_2 Z_6$	0.58	$Z_3 Z_6$	0.43				
	$Z_0 Z_5$	0.33	$Z_1 Z_4$	0.69	$Z_2 Z_5$	0.88	$Z_3 Z_5$	0.67						
Ten-node max cut	$Z_0 Z_1$	0.21	$Z_0 Z_6$	0.67	$Z_1 Z_5$	0.34	$Z_1 Z_9$	0.89	$Z_2 Z_9$	0.92	$Z_3 Z_8$	0.77	$Z_4 Z_9$	0.68
	$Z_0 Z_5$	0.41	$Z_0 Z_8$	0.82	$Z_1 Z_6$	0.77	$Z_2 Z_6$	0.45	$Z_3 Z_7$	0.81	$Z_4 Z_5$	0.35	$Z_8 Z_9$	0.15

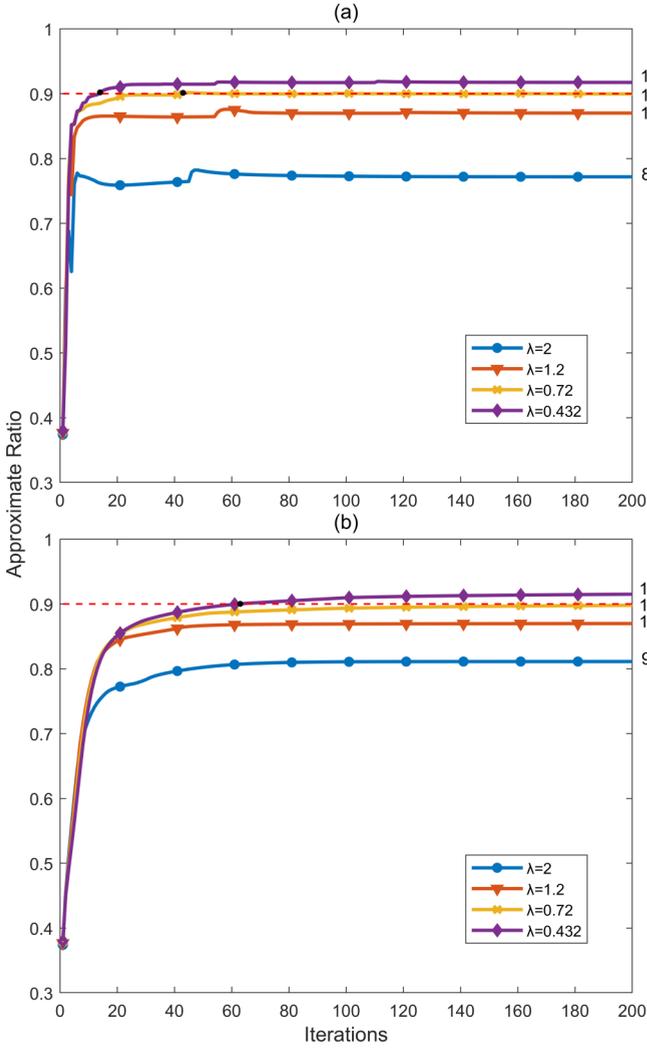


FIG. 2. Approximation ratios versus number of iterations for different  $\lambda$ s. (a) APG algorithm. (b) PG algorithm. The black dots mark for the approximation ratios have reached the minimum acceptable value of 0.9. The numbers on the right of the figure refer to the final depths after 200 iterations.

may reduce the number of required control operations by two. For example, if  $\beta_j$  in the control sequence,

$$H_o \xrightarrow{\gamma_j} H_c \xrightarrow{\beta_j} H_o \xrightarrow{\gamma_{j+1}} \quad (18)$$

has been penalized to 0, then the two control operations generated by  $H_o$  will be combined into one operation with a duration of  $\gamma_j + \gamma_{j+1}$ . As a result, the reduction of control depth from 14 to 11 implies a more significant reduction in the number of control operations. As shown in Fig. 3, although the final control depth is 11, the number of required control operations is only 8. In this case, the actual percentage of reduction in control operations is 6/14 that exceeds 40%. Since each unitary control operation is realized as a composition of quantum gates, this reduction in control depth will induce a significant reduction in the number of quantum gates or the depth of quantum circuit for practical implementation of QAOA.

The standard practice for model selection is to use the regularized model to solve for the best hyperparameter, and

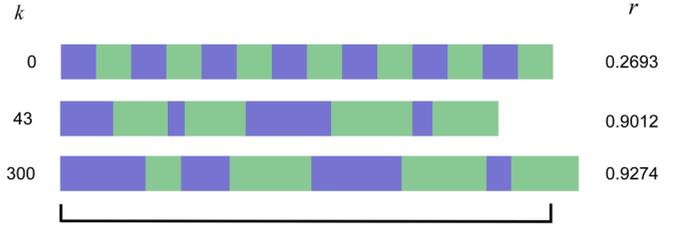


FIG. 3. The optimized control sequence with  $\lambda = 0.72$ . The first 43 iterations are calculated using the regularized model and APG algorithm. After the approximation ratio has reached 0.9, the regularization term has been removed. Conventional gradient descent is used to continue the optimization process until 300 iterations with a fixed control depth. The length of each operation is the absolute value of the parameter.

then continue the optimization process without the regularization term to achieve the best accuracy. We have followed this practice and the optimized control sequences are depicted in Figs. 3 and 4. In Fig. 4, the APG and PG algorithms are iterated for 100 times, and then the control depth is fixed for further optimization. It should be noted that the three experiments have achieved basically the same approximation ratio after 300 iterations, which is around 0.9274. This may be because the selected control depth is the same for all three experiments. Started with different values of parameters after the regularization, the ultimate accuracy is still determined by the control depth. In this case, this phenomenon can be taken as evidence to show the importance of optimizing the control depth. Next, we conduct another experiment to show that the performances of APG and PG are not sensitive to small changes in the initial parameters. To do this, the mean and standard derivation of the convergence curves of approximation ratios are calculated with respect to randomly sampled

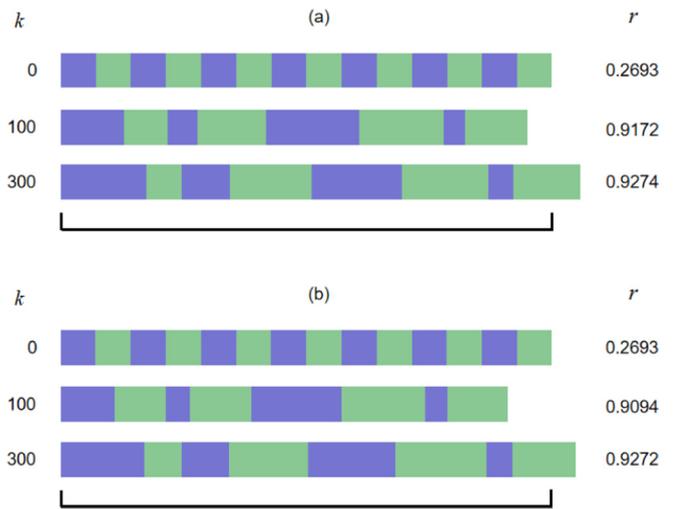


FIG. 4. The optimized control sequences for (a) APG algorithm and (b) PG algorithm with  $\lambda = 0.432$ . The first 100 iterations are calculated using the regularized model, and after that the regularization term has been removed. Conventional gradient descent is used to continue the optimization process until 300 iterations with a fixed control depth.

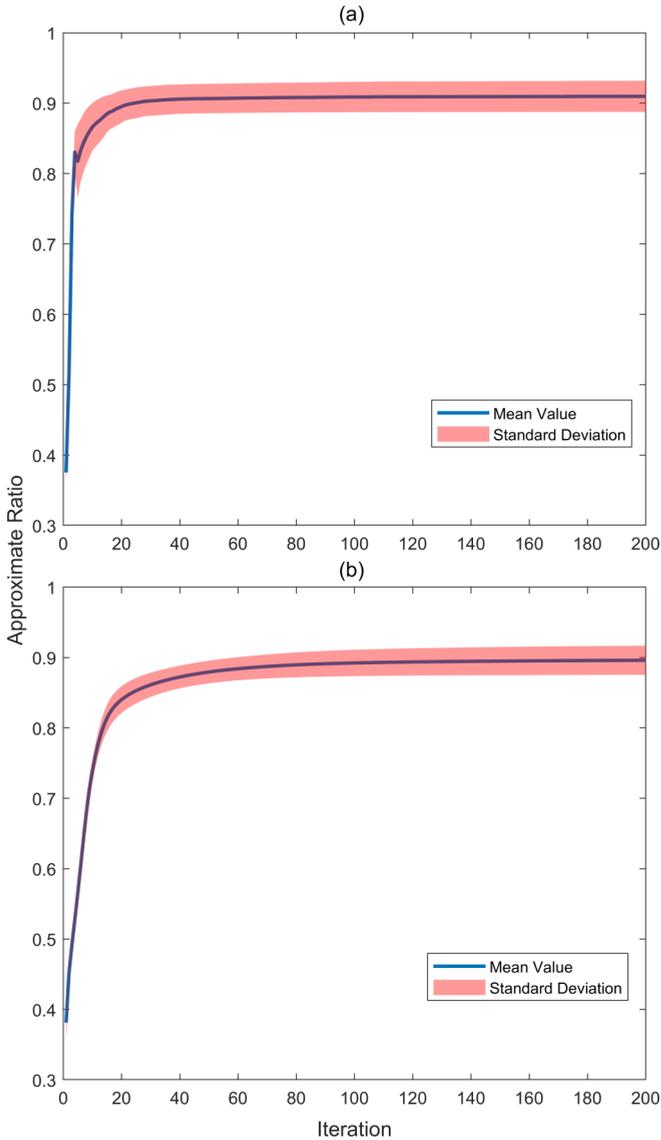


FIG. 5. Mean and standard deviation of the approximation ratios are plotted versus number of iterations for (a) APG algorithm and (b) PG algorithm. The initial values for the 14 control parameters are independently sampled from a uniform distribution over an interval [0.27,0.33].

initial values over the interval [0.27,0.33]. As shown in Fig. 5, 100 samples are drawn from a uniform distribution for APG and PG, respectively. The numerical results have confirmed that the convergence of both algorithms are stable against random initialization of parameters.

In Fig. 6, we test the performance of depth optimization by varying the initial control depth, from 20 to 40. For  $\lambda = 2$  and 1.2, the regularization term dominates and the optimized control depths are all similar for different initial values. Even if a large initial value is chosen, the control depth will be automatically reduced to a small value during the iteration, otherwise, the regularization term will become too large to optimize. In this case, both the optimized depth and the approximation ratio are steady. For  $\lambda = 0.72$  and 0.432, there is a more subtle trade-off between the regularization term and the

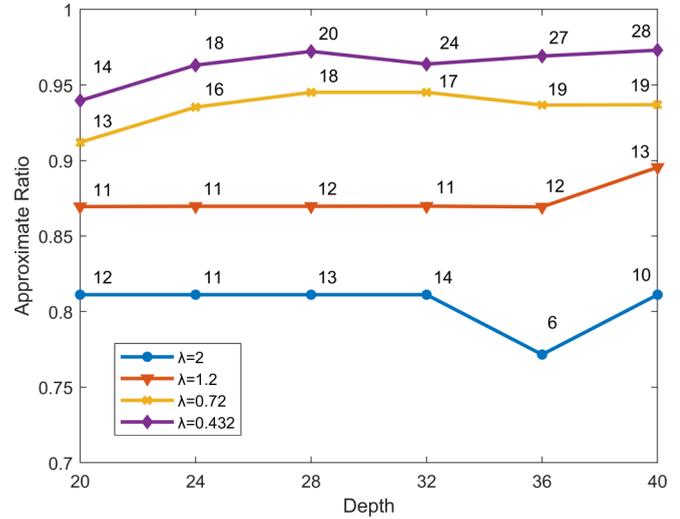


FIG. 6. The approximation ratios versus initial control depth for different  $\lambda$ . The approximation ratio for each initial depth is calculated after 200 iterations of PG algorithm. The numbers along the curve are the final depths after the optimization.

objective function. In particular, the control depths are allowed to increase slowly to promote the optimization accuracy. This observation is consistent with the existing conclusion that the optimization accuracy will monotonically increase with the control depth for noise-free QAOA. Specifically, the selected control depth doubled as we increase its initial value from 20 to 40 for  $\lambda = 0.432$ . In this case, the regularization strength is too small to penalize the parameters.

We conduct numerical experiments with initial depths 20, 28 and compare the results of regularized and unregularized cases. The detailed results are shown in Table II. In this paper, the performance is measured by the model complexity when approximation ratio reaches a desired threshold value. As can be seen from Table II, the control length, depth and iterations for the regularized models to reach 0.9 are always smaller than the unregularized model, which means that the unregularized model is always the worst and has much room to improve. Even after 200 iterations, the regularized model

TABLE II. The results are obtained when the approximation ratios reach 0.9 and after 200 iterations, respectively.

	Initial depth	$\lambda = 0$	$\lambda = 0.432$	$\lambda = 0.72$
20	Length ( $r = 0.9$ )	5.36382	5.21552	4.9451
	Depth ( $r = 0.9$ )	20	19	15
	Iterations ( $r = 0.9$ )	32	27	46
	Length (200 iterations)	6.39597	5.11906	4.99829
	Depth (200 iterations)	20	14	13
	$r$ (200 iterations)	0.9753	0.9396	0.9120
28	Length ( $r = 0.9$ )	8.95274	7.02808	6.71423
	Depth ( $r = 0.9$ )	28	28	26
	Iterations ( $r = 0.9$ )	33	25	30
	Length (200 iterations)	9.13553	6.78664	6.51269
	Depth (200 iterations)	28	20	18
	$r$ (200 iterations)	0.9882	0.9722	0.9451

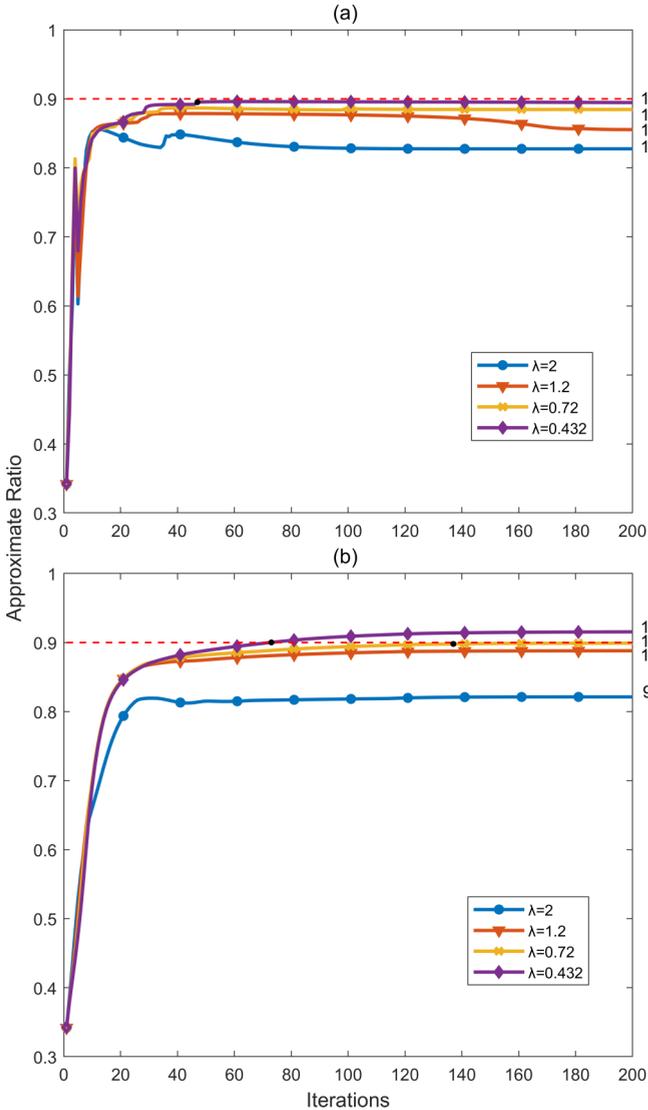


FIG. 7. Approximation ratios versus number of iterations for different  $\lambda$ . (a) APG algorithm. (b) PG algorithm. The black dots mark for the first time the approximation ratios have reached the minimum acceptable value 0.9. The numbers on the right of the figure are the final depths after 200 iterations.

with  $\lambda = 0.432$  still has similar accuracy as the unregularized one, whereas the control length and depth are significantly reduced under the regularization.

**B. Ten-node max cut**

For the ten-node instance, the initial control depth is set to be 20. The optimization landscape of the ten-node instance is more complex than that of the seven-node instance, and, thus, we choose a smaller learning rate  $\eta = 0.0035$  to stabilize the convergence. The initial values for the control parameters are still 0.3. The convergence curves for different  $\lambda$  are plotted in Fig. 7. It can be observed that the performance of APG is unstable at the initial stage. In contrast, the convergence of PG is always stable during the iteration. At the initial stage, APG seeks for an acceleration by exploring the optimization

landscape with extrapolations. Since bad extrapolations are allowed ( $q = 2$ ) in the experiments, the convergence curves can be severely perturbed at initial stage that may cause the objective function to fall into bad local minima. Moreover, the ultimate approximation ratio achieved with APG is slightly worse than PG. This may suggest one has to be cautious with APG in high-dimensional setting due to its indeterministic behavior in acceleration, although it has a proved convergence rate of  $O(1/k)$ .

**VI. CONCLUSION**

In this paper, we have proposed an automatic and fast algorithm for depth optimization of QAOA based on proximal gradient descent. On one hand, this algorithm can be used as an efficient tool to study the ultimate performances and limitations of QAOA by optimizing the hyperparameter. On the other hand, the automatic algorithm can be directly integrated with quantum hardware to accelerate the optimization process and reduce the control complexity. In particular, the proper control depth can be found with  $M$  experiments, where  $M$  is the number of candidate regularization parameters which is small and irrespective of problem size. This results in a significant reduction in the number of experiments for determining the best control depth and approximation ratio as the number of experiments for finding the best control depth by random search scales with  $O(p)$ , which could be extremely time consuming for large-scale problems.

There are two interesting directions for future work. First, due to the inevitable noise in NISQ devices, the optimal control depth could be naturally constrained. As the errors induced by the noise accumulate over time and increase with the control depth, the performance of QAOA may start to decline after the control depth reaches an optimal value. This effect can also be studied by adding noise terms to the ideal quantum dynamical model. For example, it has been found by numerical simulations in Refs. [40,41] that the performance of QAOA may not be monotonically increasing with higher depth when subjected to typical quantum noises. This phenomenon has also been observed in a recent experiment [42]. In this case, automatic approach may be able to determine the optimal control depth without using the condition (1) for complexity control. Second, a dataset can be generated for the depth optimization if robust control is considered. For example, by assuming system uncertainties, a large number of samples can be generated for optimizing the control with respect to a single control task [43,44]. Clearly, the algorithm proposed in this paper can be extended to realize the robust depth optimization for QAOA. In this case, cross validation can be used to determine the critical value  $\lambda$  that maximizes the model generalizability and robustness with a train-test split.

**ACKNOWLEDGMENT**

This research was supported by the National Natural Science Foundation of China under Grants No. 62173296 and No. 61703364.

- [1] J. Preskill, Quantum Computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
- [2] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, Technical Report MIT-CTP/4610, MIT, 2014.
- [3] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm applied to a bounded occurrence constraint problem, Technical Report MIT-CTP/4628, MIT, 2014.
- [4] E. Farhi and A. W. Harrow, Quantum supremacy through the quantum approximate optimization algorithm, Technical Report MIT-CTP/4771, MIT, 2016.
- [5] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, The theory of variational hybrid quantum-classical algorithms, *New J. Phys.* **18**, 023023 (2016).
- [6] G. G. Guerreschi and M. Smelyanskiy, Practical optimization for hybrid quantum-classical algorithms, [arXiv:1701.01450](https://arxiv.org/abs/1701.01450).
- [7] G. Verdon, M. Broughton, and J. Biamonte, A quantum algorithm to train neural networks using low-depth circuits, [arXiv:1712.05304](https://arxiv.org/abs/1712.05304).
- [8] Z. Wang, S. Hadfield, Z. Jiang, and E. G. Rieffel, Quantum approximate optimization algorithm for maxcut: A fermionic view, *Phys. Rev. A* **97**, 022304 (2018).
- [9] R. Sweke, F. Wilde, J. Meyer, M. Schuld, P. K. Faehrmann, B. Meynard-Piganeau, and J. Eisert, Stochastic gradient descent for hybrid quantum-classical optimization, *Quantum* **4**, 314 (2020).
- [10] D. Liang, L. Li, and S. Leichenauer, Investigating quantum approximate optimization algorithms under bang-bang protocols, *Phys. Rev. Research* **2**, 033402 (2020).
- [11] J. S. Otterbach, R. Manenti, N. Alidoust, A. Bestwick, M. Block, B. Bloom, S. Caldwell, N. Didier, E. Schuyler Fried, S. Hong, P. Karalekas, C. B. Osborn, A. Papageorge, E. C. Peterson, G. Prawiroatmodjo, N. Rubin, Colm A. Ryan, D. Scarabelli, M. Scheer, E. A. Sete *et al.*, Unsupervised machine learning on a hybrid quantum computer, [arXiv:1712.05771](https://arxiv.org/abs/1712.05771).
- [12] R.-B. Wu, X. Cao, P. Xie, and Y.-x. Liu, End-to-End Quantum Machine Learning Implemented with Controlled Quantum Dynamics, *Phys. Rev. Appl.* **14**, 064020 (2020).
- [13] S. Khairy, R. Shaydulin, L. Cincio, Y. Alexeev, and P. Balaprakash, Learning to optimize variational quantum circuits to solve combinatorial problems, *Proceedings of the AAAI Conference on Artificial Intelligence*, 34 (AIAA, Reston, VA, 2020).
- [14] J. Yao, M. Bukov, and L. Lin, Policy gradient based quantum approximate optimization algorithm, *Proceedings of Machine Learning Research* **107**, 605 (2020).
- [15] M. M. Wauters, E. Panizon, G. B. Mbeng, and G. E. Santoro, Reinforcement-learning-assisted quantum optimization, *Phys. Rev. Research* **2**, 033446 (2020).
- [16] M. Y. Niu, S. Lu, and I. L. Chuang, Optimizing QAOA: Success probability and runtime dependence on circuit depth, [arXiv:1905.12134](https://arxiv.org/abs/1905.12134).
- [17] L. Zhou, S.-T. Wang, S. Choi, H. Pichler, and M. D. Lukin, Quantum Approximate Optimization Algorithm: Performance, Mechanism, and Implementation on Near-Term Devices, *Phys. Rev. X* **10**, 021067 (2020).
- [18] J. Larkin, M. Jonsson, D. Justice, and G. G. Guerreschi, Evaluation of QAOA based on the approximation ratio of individual samples, [arXiv:2006.04831](https://arxiv.org/abs/2006.04831).
- [19] R. Majumdar, D. Madan, D. Bhoumik, D. Vinayagamurthy, S. Raghunathan, and S. Sur-Kolay, Optimizing ansatz design in QAOA for Max-cut, [arXiv:2106.02812](https://arxiv.org/abs/2106.02812).
- [20] T. S. Humble, R. Herrman J. Ostrowski, and G. Siopsis, Lower bounds on circuit depth of the quantum approximate optimization algorithm, *Quant. Info. Proc.* **20**, 59 (2021).
- [21] Z. Jiang, E. G. Rieffel, and Z. Wang, Near-optimal quantum circuit for grover's unstructured search using a transverse field, *Phys. Rev. A* **95**, 062317 (2017).
- [22] M. Hastings, Classical and quantum bounded depth approximation algorithms, *Quantum Inf. Comput.* **19**, 1116 (2018).
- [23] S. Bravyi, A. Kliesch, R. Koenig, and E. Tang, Obstacles to Variational Quantum Optimization from Symmetry Protection, *Phys. Rev. Lett.* **125**, 260505 (2020).
- [24] G. James, D. Witten, T. Hastie, and R. Tibshirani, *An Introduction to Statistical Learning* (Springer-Verlag, New York, 2013).
- [25] U. von Luxburg and B. Scholkopf, in *Statistical Learning Theory: Models, Concepts, and Results*, edited by Dov M. Gabbay, S. Hartmann, and J. Woods, Inductive Logic, Handbook of the History of Logic (North-Holland, Amsterdam, 2011), Vol. 10, pp. 651–706.
- [26] R. Tibshirani, Regression shrinkage and selection via the lasso, *J. Royal Stat. Soc.: Ser. B (Method.)* **58**, 267 (1996).
- [27] P. L. Combettes and J.-C. Pesquet, *Proximal Splitting Methods in Signal Processing* (Springer, New York, 2011), pp. 185–212.
- [28] A. Beck and M. Teboulle, A fast iterative shrinkage-thresholding algorithm for linear inverse problems, *SIAM J. Imaging Sci.* **2**, 183 (2009).
- [29] H. Li and Z. Lin, Accelerated proximal gradient methods for nonconvex programming, in *Proceedings of the 28th International Conference on Neural Information Processing Systems* (ACM, New York, 2015), Vol. 1, pp. 379–387.
- [30] Q. Yao, J. T. Kwok, F. Gao, W. Chen, and T.-Y. Liu, Efficient inexact proximal gradient algorithm for nonconvex problems, *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17* (ACM, New York, 2017), pp. 3308–3314.
- [31] H. Zou and T. Hastie, Regularization and variable selection via the elastic net, *J. R. Statist. Soc.: Ser. B (Statistical Methodology)* **67**, 301 (2005).
- [32] S. Khairy, R. Shaydulin, L. Cincio, Y. Alexeev, and P. Balaprakash, Learning to optimize variational quantum circuits to solve combinatorial problems, *Proceedings of the AAAI Conference on Artificial Intelligence* (AAAI, Palo Alto, CA, 2020), Vol. 2367–2375.
- [33] F. Verstraete, M. M. Wolf, and J. I. Cirac, Quantum computation and quantum-state engineering driven by dissipation, *Nat. Phys.* **5**, 633 (2009).
- [34] V. Akshay, H. Philathong, M. E. S. Morales, and J. D. Biamonte, Reachability Deficits in Quantum Approximate Optimization, *Phys. Rev. Lett.* **124**, 090504 (2020).
- [35] E. Campos, D. Rabinovich, V. Akshay, and J. Biamonte, Training saturation in layerwise quantum approximate optimization, *Phys. Rev. A* **104**, L030401 (2021).
- [36] L. Grippo and M. Sciandrone, Nonmonotone globalization techniques for the barzilai-borwein gradient method, *Comput. Opt. and Appl.* **23**, 143 (2002).

- [37] A. Bengtsson, P. Vikstål, C. Warren, M. Svensson, X. Gu, A. F. Kockum, P. Krantz, C. Križan, D. Shiri, I.-M. Svensson, G. Tancredi, G. Johansson, P. Delsing, G. Ferrini, and J. Bylander, Improved Success Probability with Greater Circuit Depth for the Quantum Approximate Optimization Algorithm, *Phys. Rev. Appl.* **14**, 034010 (2020).
- [38] J. Håstad, Some optimal inapproximability results, *J. Assoc. Comput. Mach.* **48**, 798 (2001).
- [39] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevA.105.032433> for code and data for generating figures and tables in this paper.
- [40] J. Marshall, F. Wudarski, S. Hadfield, and T. Hogg, Characterizing local noise in QAOA circuits, *IOP SciNotes* **1**, 025208 (2020).
- [41] M. Alam, A. Ash-Saki, and S. Ghosh, Design-space exploration of quantum approximate optimization algorithm under noise, *2020 IEEE Custom Integrated Circuits Conference (CICC)* (IEEE, Piscataway, NJ, 2020), pp. 1–4.
- [42] M. P. Harrigan, K. J. Sung, M. Neeley *et al.*, Quantum approximate optimization of non-planar graph problems on a planar superconducting processor, *Nat. Phys.* **17**, 332 (2021).
- [43] D. Dong, X. Xing, H. Ma, C. Chen, Z. Liu, and H. Rabitz, Learning-based quantum robust control: Algorithm, applications, and experiments, *IEEE Trans. Cybernetics* **50**, 3581 (2020).
- [44] Y. Dong, X. Meng, L. Lin, R. Kosut, and K. B. Whaley, Robust control optimization for quantum approximate optimization algorithms, *IFAC-PapersOnLine* **53**, 242 (2020); 21st IFAC World Congress.