

<https://doi.org/10.1038/s42005-024-01577-x>

Quantum approximate optimization via learning-based adaptive optimization

Check for updates

Lixue Cheng^{1,2}✉, Yu-Qin Chen^{1,2}, Shi-Xin Zhang¹✉ & Shengyu Zhang¹✉

Combinatorial optimization problems are ubiquitous and computationally hard to solve in general. Quantum approximate optimization algorithm (QAOA), one of the most representative quantum-classical hybrid algorithms, is designed to solve combinatorial optimization problems by transforming the discrete optimization problem into a classical optimization problem over continuous circuit parameters. QAOA objective landscape is notorious for pervasive local minima, and its viability significantly relies on the efficacy of the classical optimizer. In this work, we design double adaptive-region Bayesian optimization (DARBO) for QAOA. Our numerical results demonstrate that the algorithm greatly outperforms conventional optimizers in terms of speed, accuracy, and stability. We also address the issues of measurement efficiency and the suppression of quantum noise by conducting the full optimization loop on a superconducting quantum processor as a proof of concept. This work helps to unlock the full power of QAOA and paves the way toward achieving quantum advantage in practical classical tasks.

Combinatorial optimization, which involves identifying an optimal solution from a finite set of candidates, has a wide range of applications across various fields, such as logistics, finance, physics, and machine learning. However, the problem in many typical scenarios is NP-hard since the set of feasible solutions is discrete and expands exponentially with the growing problem size without any structure that seems to admit polynomial-time algorithms. As a representative NP-hard problem, MAX-CUT aims to find a bipartition of the input graph's vertices, such that the number of edges (or total edge weights) between the two subsets reaches the maximum. Classical approaches such as greedy algorithms and AI methods by graph neural networks, despite remarkable attempts and progresses^{1–3}, are generally inefficient to address combinatorial optimization problems such as MAX-CUT due to their NP-hard nature. In the recent two decades, quantum computing approaches have emerged as a new toolbox for tackling these difficult but crucial problems, including quantum annealing^{4–8} and quantum approximate optimization algorithm (QAOA)^{9–13}, from both theoretical and experimental perspectives. In this article, we focus on the latter paradigm, which is fully compatible with the universal gate-based quantum circuit model and is considered to be one of the most promising algorithms in the noisy intermediate-scale quantum (NISQ) era for potential quantum advantages.

In the QAOA paradigm, the exponential solution space is encoded in the Hilbert space of the output wavefunction of a parameterized quantum circuit. By this, the classical optimization problems in the discrete domain

are relaxed to a continuous domain composed of circuit variational parameters via QAOA as a proxy. However, the classical optimization over the continuous circuit parameter domain is still challenging (the worst case is NP-hard¹⁴) since the energy landscape of the QAOA ansatz is filled with local minima and a large amount of independent optimization processes is required to identify a near-optimal solution^{10,15}. In addition, barren plateaus can also emerge in the QAOA landscape with increasing qubit number or circuit depth^{16–19}. To overcome these optimization difficulties, various learning-based^{20–28} or heuristic-based approaches^{10,29–33} have previously been explored. These methods either rely on optimization data previously obtained or require a huge number of circuit evaluation budgets in total by progressively searching solutions of QAOA with different depths. A universally efficient and effective optimization approach suitable for real quantum processors without prior knowledge remains elusive.

In this work, we design a gradient-free classical optimizer dubbed Double Adaptive-Region Bayesian Optimization (DARBO), which exploits and explores the QAOA landscape with a Gaussian process (GP) surrogate model and iteratively suggests the most possible optimized parameter set restricted by two auto-adaptive regions, i.e., an adaptive trust region and an adaptive search region. The performance of DARBO for QAOA and ultimately for combinatorial optimization problems in terms of speed, stability, and accuracy is superior to existing methods. Furthermore, DARBO is robust against measurement shot noise and quantum noise. We demonstrated its effectiveness in extensive numerical simulations as well as a proof

¹Tencent Quantum Laboratory, Shenzhen 518057, China. ²These authors contributed equally: Lixue Cheng, Yu-Qin Chen.

✉e-mail: sherrylixuecheng@gmail.com; shixinzhang@tencent.com; shengyuzhang@tencent.com

of concept demonstration of the quantum-classical optimization pipeline, where QAOA is implemented and evaluated on a real superconducting quantum processor using five qubits with integrated quantum error mitigation (QEM) techniques.

Results

QAOA framework with DARBO

MAX-CUT problem and a large family of combinatorial optimization problems can be embedded into a more general formalism as quadratic unconstrained binary optimization (QUBO)³⁴, where the objective function $C(\mathbf{z})$ to optimize is in the form as follows in terms of binary-valued variables \mathbf{z} :

$$C(\mathbf{z}) = \sum_{ij} w_{ij} z_i z_j, \quad (1)$$

where w_{ij} can be regarded as the edge weights defined on a given graph.

The QAOA framework is designed as a quantum-enhanced method to solve these QUBO problems. The quantum circuit ansatz for QAOA consists of the repetitive applications of two parameterized unitary operators. We denote the number of repetitions as p , and the number of qubits (binary freedoms in QUBO) as n . The quantum program ansatz is constructed as:

$$|\psi(\gamma, \beta)\rangle = U(\gamma, \beta)|0^n\rangle = \prod_{k=1}^p \left(e^{-i\beta_k \sum_i X_i} e^{-i\gamma_k \sum_{ij} w_{ij} Z_i Z_j} \right) \prod_i H_i |0^n\rangle,$$

where X and Z are Pauli matrices on each site and H are Hadamard gate. The trainable parameters γ and β both consist p real-valued components, and the outer classical training loop adjusts these parameters so that the objective $C(\gamma, \beta) = \langle \psi(\gamma, \beta) | \sum_{ij} w_{ij} Z_i Z_j | \psi(\gamma, \beta) \rangle$ is minimized. Therefore, by utilizing the QAOA framework, the optimization over discrete binary \mathbf{z} variables is reduced to the optimization over continuous variables of β and γ .

However, the continuous optimization problem still faces lots of pressing challenges. In the QAOA framework, gradient descent optimizers commonly utilized in the deep learning community do not work well. The common ansatz consisting of a large number of parameters can enjoy the benefits of over-parameterization and their global minima are easier to locate^{35,36}. However, QAOA ansatz has a small number of parameters and thus a large number of local minima, which often destroys the effort of conventional optimizers to identify the global minimum. Besides, barren plateaus that render the gradient variance vanishing exponentially can occur similarly as the generic cases in variational quantum algorithms. More

importantly, gradient evaluations on real quantum chips are too noisy and costly to use for a classical optimizer.

Bayesian optimization (BO) is a class of black-box and gradient-free classical optimization approaches that can effectively optimize expensive black-box functions and tolerate stochastic noise in function evaluations. The method typically creates a surrogate for the unknown objective, and quantifies and manages the uncertainty using a Bayesian learning framework^{37,38}. Although conventional BO has become a highly competitive technique for solving optimizing problems with a small number of parameters, it usually does not scale well to problems with high dimensions³⁷⁻⁴⁰. Aside from the plentiful local minima in the exponentially large search space, another challenge with BO is that the surrogate function fitting with very few samples can hardly be globally accurate.

To overcome the above issues and enable efficient QAOA executions on real quantum chips, we propose DARBO as a powerful classical optimizer for QAOA. The schematic quantum-classical hybrid workflow for DARBO-enabled QAOA is shown in Fig. 1. The advantages of DARBO are both from its Bayesian optimization nature and the two adaptive regions utilized in the algorithm. The idea of including an adaptive trust region is directly borrowed from Turbo³⁸ and is inspired by a class of trust region methods from stochastic optimization⁴¹. These methods utilize a simple surrogate model inside a trust region centered on the current best solution. For instance, COBYLA⁴² method used as a baseline in this work models the objective function locally with a linear model. As a deterministic approach, COBYLA is not good at handling noisy observations. By integrating with GP surrogate models within an adaptive trust region, DARBO inherits the robustness to noise and rigorous reasoning about uncertainties that global BO enjoys as well as the benefits that local surrogate model enables. In addition, the introduction of an adaptive search region makes DARBO more robust to different initial guesses by moving queries in some iterations to a more restricted region. The search efficiency increases when the search space is reduced by the adaptive search region, giving DARBO a higher chance of finding the global minimum rather than local minima.

In this study, the end-to-end performance of QAOA with DARBO is evaluated and benchmarked on the basis of analytical exact simulation, numerical simulation with measurement shot noise, and quantum hardware experiments (with both measurement shot noise and quantum noise). Overall, DARBO outperforms other common optimizers by a large margin in terms of (1) efficiency: the number of circuit evaluations to reach a given accuracy is the least, (2) stability: the fluctuation of the converged objective value across different initializations and graph instances is the least, (3) accuracy: the final converged approximation ratio is the best, and (4) noise robustness: the performance advantage is getting larger when noise presents, which is unavoidably the case on quantum processors.

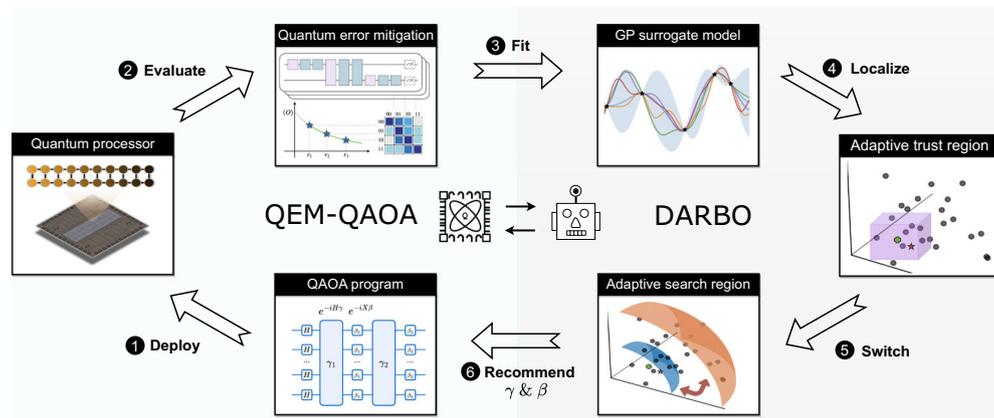


Fig. 1 | The proof-of-concept workflow for error mitigated QAOA on the superconducting quantum processor with DARBO. We compile and deploy the 5-qubit QAOA program for given objective functions on a 20-qubit real superconducting quantum processor and evaluate the objective value with quantum error

mitigation methods. DARBO treats the QEM-QAOA as a black-box, and optimizes the circuit parameters by fitting the surrogate model with constraints. The constraints are provided by the two adaptive regions, which are responsible for surrogate model building and acquisition function sampling, respectively.

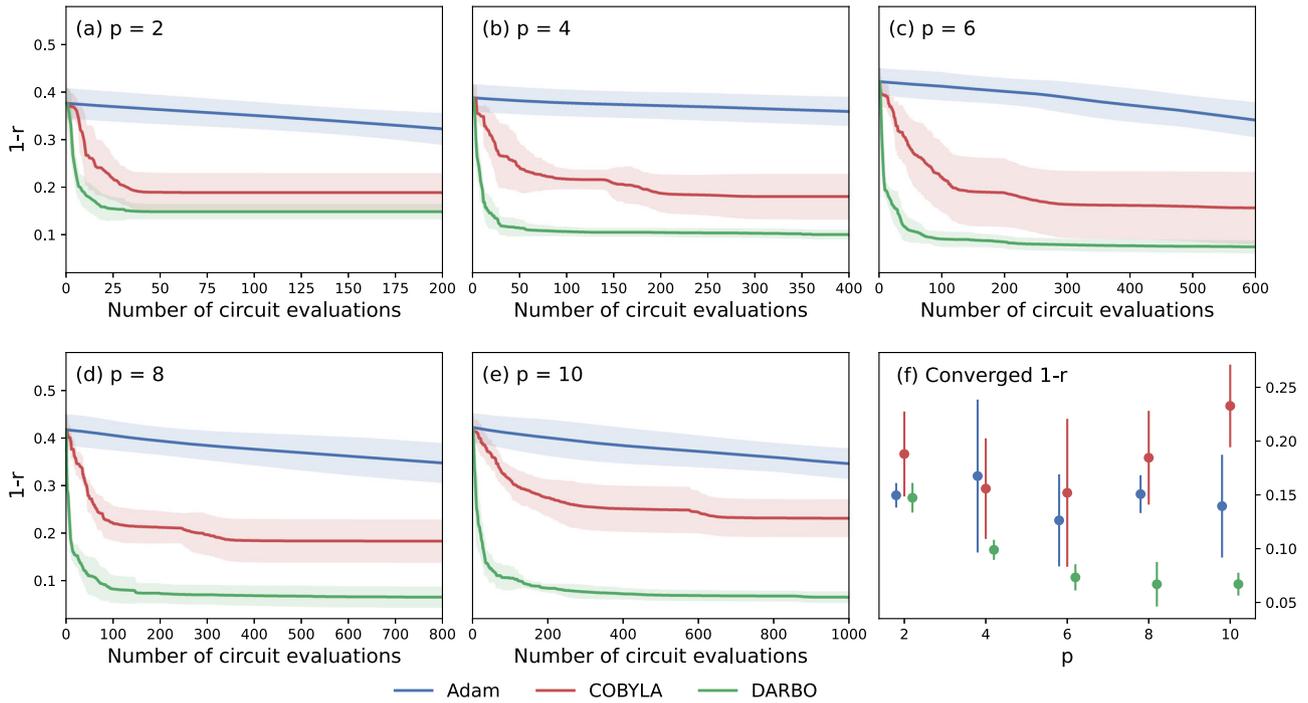


Fig. 2 | QAOA optimization for MAX-CUT problem on w3R graphs (exact simulations). All the optimizations are performed on $n = 16$ w3R graphs. **a–e** The optimization trajectories from different optimizers, i.e., Adam (in blue color), COBYLA (in red color), and DARBO (in green color) are plotted versus the number of circuit evaluations. Results on different circuit depths from $p = 2$ to $p = 10$ are reported, respectively. **f** The final converged approximation gap $1 - r$ after sufficient

numbers of optimization iterations. For each circuit evaluation, we query the exact expectation of the objective function via numerical simulation. Each line is averaged over five w3R graph instances, where the shaded range shows the standard deviation of the results across different graph instances. For each graph instance, the best optimization result among the 20 independent optimization trials is reported. The error bar in **(f)** shows the standard deviation across different graph instances.

We use approximation ratio r as a metric to measure the end-to-end performance of QAOA. In MAX-CUT problem, r is defined as the ratio of obtained cut value derived from the objective expectation over the exact max cut value ($r = 1$ indicates a perfect solution for the problem), and $1 - r$ is the (relative) approximation gap. Throughout this work, we investigate the MAX-CUT problem on the w3R graph which is a family of regular weighted graphs whose vertices all have degrees of three.

Analytically exact simulation

We firstly report the results on optimization trajectories and final converged objective values for analytically exact quantum simulation, which computes the objective function by directly evaluating the expectations from the wavefunction. Figure 2 shows the results for a collection of different optimizers, including Adam⁴³, COBYLA⁴², and DARBO. The results are collected from five different $n = 16$ w3R graph instances, and for each graph and each optimization method, the best record over 20 independent optimization trials is reported. The results show the superior efficiency of DARBO over other common optimizers. For all different p values, the approximation gaps $1 - r$ of Adam and COBYLA are 1.02 ~ 2.08 times and 1.28 ~ 3.47 times larger than those of DARBO, respectively. In addition to the efficiency and accuracy, the DARBO performance is also more stable in terms of different problem graph instances. The optimization settings, five graph instances and the results of more choices of common optimizers are shown in Supplementary Notes 1–3, respectively. We also report the fidelity of three chosen optimizers on different sizes of the graphs in the Supplementary Note 3, which essentially gives the probability of obtaining the exact solution state from measuring the QAOA circuit.

It is worth noting that stability is of great importance for optimizing over the QAOA parameter landscape, as the great number of local minima requires generically exponential independent optimization trials to reach a global minimum¹⁰. As a result, with the increasing depth p of QAOA, the

optimization problem becomes harder, and we see that Adam and COBYLA do not even exhibit a monotonic growth of the accuracy as p increases, despite the fact that the cut size given by the optimal $|\psi(\gamma, \beta)\rangle$ (over the choice of $2p$ parameters (γ, β)) is clearly non-decreasing with p . The reason of the performance drop is due to a lack of training stability and the sensitivity on initial parameter choices for conventional local optimization methods. DARBO, on the contrary, does give better results with large depth p . In other words, one important optimization advantage brought by DARBO is its capability of finding the near optimal parameters with a small number of independent optimization trials, which are not enough to locate the global optimal parameters for conventional optimizers. Some advanced optimization methods such as FOURIER heuristics reported in ref.¹⁰ can be good at locating global minimum but depend on progressive optimization on lower-depth QAOA, leading to a much larger total amount of required circuit evaluations.

Numerical simulation with finite measurement shots

With the introduction of noise, DARBO shows more advantageous results compared to other optimizers, including Adam⁴³, COBYLA⁴² and SPSA⁴⁴. In Fig. 3, we show the optimization results of QAOA on five different $n = 16$ w3R graph instances with $p = 10$ with different measurement shots at each iteration, and for each graph and each optimization method, the best record over 20 independent optimization processes is reported. For results on QAOA of different p , see Supplementary Note 4. Since QAOA for MAX-CUT has a commutable objective function, the budgets of the measurement shots are all taken on the computational basis. It is natural that with more measurement shots, better accuracy could be achieved for the objective evaluation. Taking m as the given number of measurement shots, we evaluate the final QUBO objective by reconstructing from measurement bitstrings represented by binary valued $z_{ij} = \pm 1$ where $i \leq m$ runs over different shots and $j \leq n$ runs over different qubits. The objective value C is

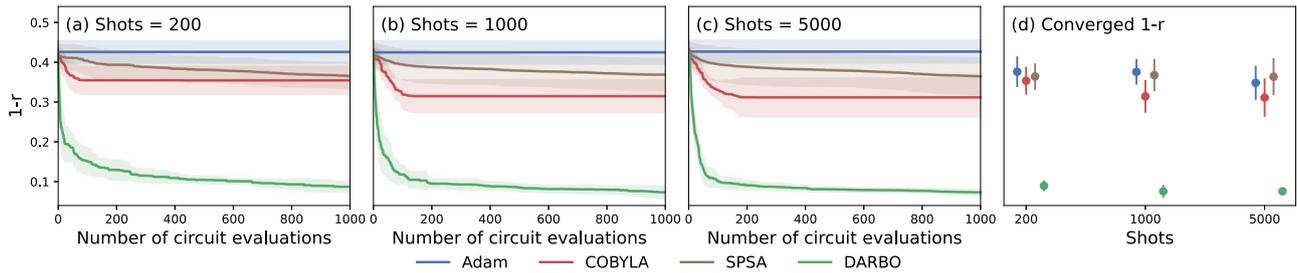


Fig. 3 | QAOA optimization for MAX-CUT problem on w3R graph (simulation with measurement shot noise). All the optimizations are performed on $n = 16$ w3R graphs with $p = 10$. **a–c** The optimization trajectories from different optimizers, i.e., Adam (in blue color), COBYLA (in red color), SPSA (in brown color), and DARBO (in green color) in terms of the number of circuit evaluations. Results for different shot numbers from shots = 200 to shots = 5000 are reported, respectively. **d** The final converged approximation ratio $1 - r$ after sufficient numbers of optimization

iterations. For each circuit evaluation, we collect the number of shot measurements to further reconstruct the loss expectation value. Each line is averaged over five w3R graph instances where the shaded range shows the standard deviation of the results across different graph instances. For each graph instance, the best optimization result from 20 independent optimization trials is kept. The error bar in **(d)** shows the standard deviation across different graph instances.

estimated as:

$$C = \frac{1}{m} \sum_{i=1}^m \sum_{j=1}^n \sum_{k=1}^n w_{jk} z_{ij} z_{ik}. \tag{2}$$

This value is a random variable with a Gaussian distribution whose mean value is determined by the analytical expectation value of the objective function, and the standard deviation is controlled by the number of total shots m with $\frac{1}{\sqrt{m}}$ scaling.

For DARBO, we can reach a satisfying optimized objective value with a small number of circuit evaluations, e.g., 200 measurement shots at each round are more than enough. The optimization efficiency gets further improved with an increased number of measurement shots. In Fig. 3, the approximation gaps $1 - r$ of Adam, COBYLA, and SPSA are 4.20 ~ 4.59, 3.95 ~ 4.10, and 4.07 ~ 4.79 times larger than those of DARBO, respectively. Besides, the solution quality from DARBO across different problem instances is impressively stable. The efficiency, accuracy, and stability of DARBO are all much better than those of other optimizers evaluated in our experiments. In addition, the performance gap between DARBO and other optimizers is getting larger compared to the infinite measurement shots (analytical exact) case, which reflects the noise robustness and adaptiveness of our proposed optimizer.

As a gradient-free optimization approach, Bayesian optimization has the advantages of robustness to noise and rigorous reasoning about uncertainty. Naïve BO methods have already been utilized to optimize variational quantum algorithms^{45–48}, but they often suffer from the low efficiency. Benefiting from the adaptive trust region and local Gaussian process (GP) surrogate model, DARBO has better potential to optimize noisy problems³⁸. Instead of directly using the currently best-observed solution x^* , we use the observation with the smallest posterior mean under the surrogate model, and therefore the noise affecting x^* has limited effects on the optimization. DARBO is specifically suitable for optimization with shot noises induced by statistical uncertainty of finite measurement shots, since its GP surrogate model assumes that the observations are Gaussian-distributed random variables⁴⁹ which is consistent with the case for measurement results with finite shot noises.

Experiments on superconducting quantum hardware

Finally, we run QAOA equipped with DARBO on real quantum hardware to demonstrate its performance. Quantum error, in addition to shot noise, has a huge impact on optimization performance on real quantum hardware. It has been studied that quantum noises would in general flatten the objective function landscape and induce barren plateaus in variational quantum algorithms¹⁸. Here we investigate the effect of quantum noise on the performance of DARBO for QAOA, and at the same time analyze how

the common error mitigation strategies^{50–53} can help in the DARBO process and achieve better end-to-end results.

The target problem is to optimize a five-variable QUBO (see Supplementary Note 5 for the problem definition in detail). The experimental results are shown in Fig. 4. We carry out the optimization on (1) raw objective value directly evaluated from measurement results on real hardware, (2) mitigation objective value evaluated from measurement results on real hardware with integrated quantum error mitigation techniques including layout benchmarking, readout error mitigation, and zero-noise extrapolation, see “Method” for more details, and (3) numerical exact value without quantum noise as a comparison. The optimization results are shown in terms of objective optimization history and success ratio from sampling the final QAOA circuit. Although the expectation value is conveniently taken as the objective value for the optimization process, the success ratio is another important representative metric to straightforwardly evaluate the final performance of QAOA for the QUBO objective since the true objective value can be directly reconstructed by the bitstring measured.

We noticed that DARBO conducted even on raw measurements can improve the cut estimation from the initial value, although it is not good enough compared to the ideal one due to the large influence of quantum noise. The optimization results combined with QEM are much better than the raw evaluations, both in terms of objective evaluation and success ratio from sampling the final QAOA circuit. Moreover, the performance gap between optimization on the mitigation value from experiments and that on the ideal value from numerics becomes larger for larger p , which is consistent with the fact that deeper circuits bring larger quantum noise. Still, we show that a deeper QAOA with $p = 2$ gives a better approximation of the QUBO objective than a shallow QAOA with $p = 1$, achieving a better trade-off between expressiveness and the accumulated noise.

The raw data collected directly from real quantum hardware contain both quantum noises and measure shot noises, which are essentially the bias and variances from the perspective of machine learning. In this QEM-QAOA + DARBO framework, QEM helps to reduce the effects of bias on the hardware (gate noises, readout noises, decoherence noises and so on) by mitigating these errors, and DARBO avoids the negative influence of variances from repetitive measurements (shot noises). Therefore, these two key components together make the proposed framework a powerful optimization protocol for combinatorial optimization problems.

Discussion

With a better exploration of the QAOA landscape, the optimization routine based on Bayesian optimization shows weak initial parameter dependence and a better probability of escaping the local minimum. Although, in this work, the dimension of the parameter spaces is still relatively low, an interesting future direction is to generalize similar BO methods from the QAOA setup to other variational quantum algorithms, which has a larger

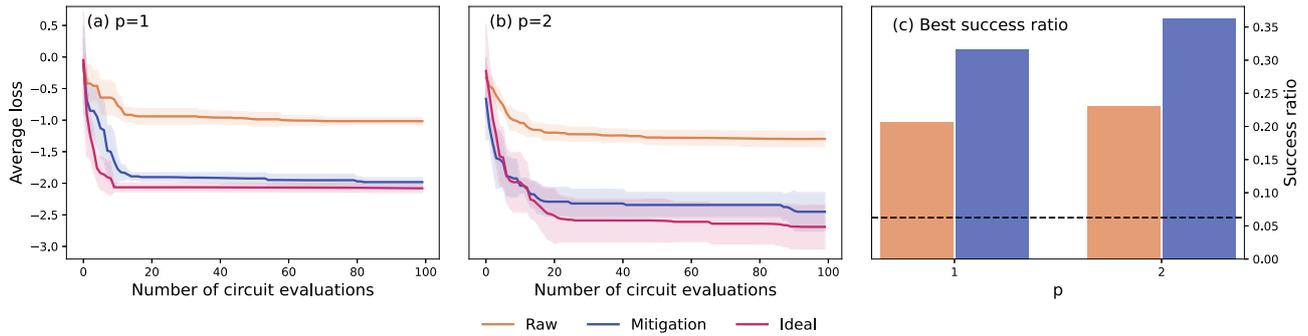


Fig. 4 | Quantum optimization of a five-variable QUBO problem on real quantum hardware. Measurement shots = 10,000. **a, b** show results from two circuit depths $p = 1$ and $p = 2$ QAOA, respectively. The line is the average optimization trajectory of five independent optimization trials, while the shaded area represents the standard deviation across five independent optimization trials. Averaged loss refers to the expectation value of the problem QUBO Hamiltonian. Raw (in orange color): at each step, we obtain the loss expectation directly from measurement results

on the real quantum processor. Mitigation (in blue color): at each step, we obtain the loss expectation from measurement results integrated with QEM techniques. Ideal (in red color): at each step, we obtain the loss from numerical simulation. **c** The success ratio when we run inference on the trained QAOA program, i.e., the probability that we can obtain a correct bitstring answer for the problem on real quantum hardware. The dashed line is the random guess baseline with a probability of $1/16$. We report the best success ratio of the five optimization trials.

number of parameters. Recently, several advanced BO variants have been proposed to increase the optimization efficiency and robustness in high-dimensional problems^{40,54,55} and in problems with noisy observations^{39,56–58}. These approaches show superior optimization performances in challenging benchmarks with large parameter sizes and the presence of noises. For instance, an advanced BO approach could efficiently optimize a higher dimensional problem ($D = 385$)⁵⁴, and accurately find the best experimental settings for the real-world problems in chemistry⁵⁹, material sciences⁶⁰, and biology⁶¹. These cases are potentially relevant for optimization in variational quantum eigensolver, quantum machine learning, and quantum architecture search scenarios.

We also note that the double-adaptive region idea in BO is a general framework. The detail settings in the DARBO approach could be designed differently for different optimization problems. As a future direction, DARBO algorithm could be extended to include more than two adaptive search regions, and the ranges of these regions themselves could also be adapted during the optimizations.

To successfully scale the QAOA program on real quantum hardware with meaningful accuracy, more pruning and compiling techniques for QAOA deployment, as well as more error mitigation techniques, can be utilized in future works. For example, by differentiable quantum architecture search-based compiling⁶², we can greatly reduce the number of two-qubit quantum gates required with even better approximation performance. There is also QAOA tailored error mitigation algorithm⁶³ that trades qubit space for accuracy.

In summary, we proposed an optimizer—DARBO suitable for exploring the variational quantum algorithm landscapes and applied it to the QAOA framework for solving combinatorial problems. The end-to-end performance for combinatorial problems is greatly improved in both numerical simulation and experiments on quantum processors. These promising results imply a potential quantum advantage in the future when scaling up the QAOA on quantum hardware, and give a constructive and generic method to better exploit this advantage.

Methods

Double adaptive-region Bayesian optimization

In this work, the QAOA problem is formulated as a maximization problem with an objective function of $-\mathcal{L}$ with D total number of parameters to be optimized:

$$\max_{x \in \mathcal{X}} -\mathcal{L}(y, \beta), \quad (3)$$

Generally, initialized with one randomly selected point from $[0, 1]^D$, a Bayesian optimization (BO) algorithm optimizes a hidden objective

function $y = y(x)$ over a search space \mathcal{X} by sequentially requesting $y(x)$ on points $x \in \mathcal{X}$, usually with a single point in each iteration^{37,64}. At each iteration i , a Bayesian statistical surrogate model s regressing the objective function is constructed using all currently available data $(x_1, y_1), \dots, (x_{i-1}, y_{i-1})$. The next point x_i to be observed is determined by optimizing a chosen acquisition function, which balances exploitation and exploration and quantifies the utility associated with sampling each $x \in \mathcal{X}$. This newly requested data (x_i, y_i) is then updated into the available dataset. This BO procedure continues until the predetermined maximum number of iterations (1000 in this work) is reached or the convergence criteria are satisfied. The general BO approach is available in the ODBO package⁶¹.

Gaussian process surrogate model. In this study, we use the Gaussian process (GP)⁶⁵ as the surrogate model^{40,55}. With a given set of available observations (\mathbf{X}, \mathbf{y}) , GP provides a prediction for each point x' as a Gaussian distributed $y' \sim \mathcal{N}(\mu(x'), \sigma^2(x'))$, where $\mu(x')$ (Eq. (4)) is the predictive mean and $\sigma(x')$ is the corresponding uncertainty (Eq. (5))

$$\mu(x') = k(x', \mathbf{X})^T \mathbf{K}^{-1} \mathbf{y} \quad (4)$$

$$\sigma^2(x') = k(x', x') - k(x', \mathbf{X})^T \mathbf{K}^{-1} k(x', \mathbf{X}), \quad (5)$$

where k is a Matérn5/2 kernel function (Eq. (6)) with a parameter set $\theta = \{\sigma_v, l\}$ to be optimized, and $\mathbf{K} = k(\mathbf{X}, \mathbf{X}) + \sigma_n^2 \mathbf{I}$ with a white noise term σ_n ⁶⁵. In this study, the variance parameter σ_v and isotropic lengthscales l constructed by automatic relevance determination in the kernel k are optimized by Adam⁴³ implemented in GPyTorch⁶⁶:

$$k(x, x') = \sigma_v^2 \left(1 + \sqrt{5}r + \frac{5}{3}r^2 \right) \exp(-\sqrt{5}r), \quad (6)$$

where $r = \|x - x'\|_2 / l$.

The DARBO algorithm is inspired by one of the most efficient BO algorithms, trust region Bayesian optimization algorithm (TuRBO)³⁸, which performs global optimization by conducting BO locally to avoid exploring highly uncertain regions in the search space. TuRBO was developed to mainly resolve the issues of high-dimensionality and heterogeneity of the problem and has been demonstrated to obtain remarkable accuracy on a range of datasets³⁸. We have applied TuRBO to QAOA problem and identified its performance advantages. DARBO inherits the advantages of TuRBO with an additional abstraction of the adaptive search region; therefore, it further enhances the optimization efficiency of QAOA problems.

Adaptive trust region. In DARBO, instead of directly querying the next best point for the quantum circuit, we first determine the two adaptive regions, starting from the adaptive trust region. At optimization iteration i , the adaptive trust region (TR) is a hyper-rectangle centered with the i th base side length $L_{\min} \leq L_i \leq L_{\max}$ at the current best solution x^{*38} . In our case, the minimum allowed length $L_{\min} = 2^{-10}$, and the maximum allowed length $L_{\max} = 3.2$. To obtain a robust and accurate surrogate for more efficient acquisitions, the GP surrogate model is regressed locally within the trust region, s.t. points far away from the current best solution cannot affect the regression quality. If the most recent queried point is better than the current best solution, we count the query in this iteration as a success. Otherwise, we count it as a failure. To guarantee that it is small enough to ensure the accuracy of the local surrogate model and big enough to include the actual best solution, the TR (trust region) length L_i is automatically updated with the proceeding of BO cycles as follows:

$$L_0 = 1.6,$$

$$L_i = \begin{cases} \min(L_{\max}, 2L_{i-1}), & \text{if } t_s \geq \tau_s \\ L_{i-1}/2, & \text{if } t_f \geq \tau_f \end{cases} \quad (i \geq 1)$$

where τ_s and τ_f are the threshold hyperparameters for the number of the maximum consecutive successes and that of the maximum consecutive failures, respectively, and t_s and t_f are the actual numbers of consecutive successes and failures in the current BO procedure. We set $\tau_s = 3$ and $\tau_f = 10$ in this study. If L_i reaches the minimum allowed L_{\min} before the end of the execution, we rescale L_i as $L_i = L_i \times 16$. The introduction of TR could not only enjoy the traditional benefits of robustness to noisy observations and rigorous uncertainty estimations in BO, but also allow for heterogeneous modeling of the objective function without suffering from over-exploration.

Adaptive search region. We also maintain a second adaptive region, the adaptive search region, with the proceeding of the optimization. The region is automatically determined by the switch counter c_s , which counts the consecutive searching failure number in the current search region. Once c_s reaches the maximum allowed consecutive failure hyperparameter $\kappa_s = 4$, the adaptive search region switches to the other pre-determined searching region. This also indicates that exploitation within this current region might be currently exhausted. The adaptive search regions serve as constraints for the parameters x . Only the points in the current searching region will be considered as possible candidates to be queried, and the switch counter allows BO search with different constraints. Inspired by the conclusion from ref. ¹⁰ that the parameter space can be reduced in given graph ensembles the two adaptive search regions are determined as $A = [-\pi/2, \pi/2]^D$ (the restricted search space) and $B = [-\pi, \pi]^D$ (the full search space) in our study.

Note that the two adaptive regions take different roles in the DARBO algorithm. The adaptive trust region provides a more precise surrogate model around the best solution by limiting the training points to be fitted in GP, while the adaptive search region constrains the candidate parameter sets temporarily by switching between the restricted search space and the full search space. In this work, to search more efficiently, we further restrict the acquisition function to select new candidate points that lie in the overlap between the TR and the adaptive search region, as in the default implementation of the ODBO package⁶¹. For the cases where there is no overlap between the adaptive trust region and the adaptive search region, we reset the trust region to be equal to the current adaptive search region.

Upper confidence bound acquisition function. In order to query the next best point, acquisition functions that balance exploitation and exploration using the posterior distributions from GP (Eq. eq:mu and eq:sigma) are required. The point with the highest acquisition value is the candidate point to be queried from the quantum circuit. In this study, we only evaluate the points within the adaptive search region using upper

confidence bound (UCB)⁶⁷ acquisition function in Eq. (7):

$$\alpha_{\text{UCB}}(x) = \mu(x) + \beta\sigma(x), \quad (7)$$

where $\beta = 0.2$ is a predefined hyperparameter to control the degree of exploration, and μ and σ are the predictive mean and uncertainty from local GP modeled with points in the adaptive trust region.

Quantum error mitigation

Besides the quantum algorithm, another key to operating experiments on quantum devices is the investigation and mitigation of quantum errors. We utilize a number of error mitigation methods in order to obtain desirable results for our QAOA program.

Layout benchmarking. The qubit quality and the single- and two-qubit gate fidelity vary across different quantum devices and vary over time. Device error can be initially attenuated by selecting qubits with better quality and links that host two-qubit gates with a lower error rate. These metrics can be benchmarked and collected by calibration experiments, including T1/T2 characterization and randomized benchmarking. In particular, we chose two-qubit gates that are directly connected on hardware to avoid additional swap manipulations introduced in quantum compiling.

In order to further determine the circuit structures, especially the applying order of those two-qubit couplings (all two-qubit couplings commute with each other in QAOA for QUBO objectives), we run multiple reference circuits by permuting those two-qubit gates under the same set of parameters and identify the optimal circuit structure that shows the highest fidelity with the ideal state. These trial experiments provide valuable insights on the circuit structures with overall low noise effects that balance the influence of crosstalk and circuit depth. The key tradeoff in layout benchmark is that: on the one hand, for compact two-qubit gate layout, the overall circuit depth is short, while there are more two-qubit gates applied at the same time which may induce larger cross-talk effect. On the other hand, the two-qubit layout can be placed in a rather sparse fashion, which has less cross-talk effects but takes longer physical evolution time. Therefore, we can explore different two-qubit layouts to minimize the overall noise effect. In our implementation, we use brute-force search. For systems with larger sizes, greedy search or more advanced reinforcement learning methods can be explored for better scalability, which is an interesting future direction.

Readout error mitigation. The imperfect measurement operation on a quantum circuit can result in readout errors that bias the original quantum state to certain bit strings. The readout error on the device used in the experiments is around 10^{-1} . We mitigate the readout error by several steps: (1) learn how the readout is biased by measuring states that produce fixed bitstring outputs, (2) encode all deviations in a confusion matrix, and (3) invert the confusion matrix and apply it to raw counts of bitstrings to correct the measurement bias. The size of the confusion matrix is 2^n where n is the number of measured qubits. For the error learning process, we tried both local learning and global learning modes. The local learning process characterizes the readout bias on each single qubit independently (involving 2 calibration circuits in the minimal case), while the global learning process models the readout bias of the Hilbert space expanded on all the qubits (involving 2^n calibration circuits) by capturing the readout correlation between qubits. We find that the local learning is good enough in our experiments as the readout correlation is negligible on the device we used.

Zero-noise extrapolation. Zero-noise extrapolation (ZNE) is one of the most widely used error mitigation methods that can be applied without detailed knowledge of the underlying noise model and exhibits significant improvement in the results evaluated on quantum devices. The main idea of ZNE is to obtain expectation values at several different error rates and extrapolate to the zero noise limit according to those noisy expectation

values. Suppose that two-qubit gates contribute the most of the errors, we conduct experiments on different error rates [1, 3, 5] by locally folding those two-qubit gates $[U, UU^\dagger U, UU^\dagger UU^\dagger U]$ to avoid circuit depth that challenges the coherence time. As for the experiments in the main text, for $p = 1$ ($p = 2$), we adopt linear (quadratic polynomial) extrapolation to estimate the mitigated results. All the expectation values used in ZNE are firstly mitigated by readout error mitigation.

All the numerical simulations and quantum hardware experiments including error mitigation in this work are implemented and managed using TensorCircuit⁶⁸—a high-performance and full-featured quantum software framework for the NISQ era.

Data availability

All graphs and results presented in this study are shared on a github repository (<https://github.com/sherrylixuecheng/EMQAOA-DARBO>). The additional figures for test results are shown in the Supplementary Information and the full statistics of the optimized losses and r values are included in a separate excel (Supplementary Data 1).

Code availability

The entire DARBO framework is available on github (<https://github.com/sherrylixuecheng/EMQAOA-DARBO>) with example jupyter notebooks and all the testing results. The example codes to perform QAOA evaluations and DARBO optimizations are also available on TensorCircuit (<https://github.com/tencent-quantum-lab/tensorcircuit>)⁶⁸, and ODBO (<https://github.com/tencent-quantum-lab/ODBO>)⁶¹.

Received: 20 June 2023; Accepted: 25 February 2024;

Published online: 06 March 2024

References

- Schuetz, M. J., Brubaker, J. K. & Katzgraber, H. G. Combinatorial optimization with physics-inspired graph neural networks. *Nat. Mach. Intell.* **4**, 367–377 (2022).
- Angelini, M. C. & Ricci-Tersenghi, F. Modern graph neural networks do worse than classical greedy algorithms in solving combinatorial optimization problems like maximum independent set. *Nat. Mach. Intell.* **5**, 29–31 (2022).
- Boettcher, S. Inability of a graph neural network heuristic to outperform greedy algorithms in solving combinatorial optimization problems. *Nat. Mach. Intell.* **5**, 24–25 (2023).
- Kadowaki, T. & Nishimori, H. Quantum annealing in the transverse Ising model. *Phys. Rev. E* **58**, 5355–5363 (1998).
- Farhi, E. et al. A quantum adiabatic evolution algorithm applied to random instances of an NP-complete problem. *Science* **292**, 472–475 (2001).
- Johnson, M. W. et al. Quantum annealing with manufactured spins. *Nature* **473**, 194–198 (2011).
- Hauke, P., Katzgraber, H. G., Lechner, W., Nishimori, H. & Oliver, W. D. Perspectives of quantum annealing: methods and implementations. *Rep. Prog. Phys.* **83**, 054,401 (2020).
- Hibat-Allah, M., Inack, E. M., Wiersema, R., Melko, R. G. & Carrasquilla, J. Variational neural annealing. *Nat. Mach. Intell.* **3**, 952–961 (2021).
- Farhi, E., Goldstone, J. & Gutmann, S. A quantum approximate optimization algorithm. Preprint at <https://arxiv.org/abs/1411.4028> (2014).
- Zhou, L., Wang, S. T., Choi, S., Pichler, H. & Lukin, M. D. Quantum approximate optimization algorithm: performance, mechanism, and implementation on near-term devices. *Phys. Rev. X* **10**, 021,067 (2020).
- Harrigan, M. P. et al. Quantum approximate optimization of non-planar graph problems on a planar superconducting processor. *Nat. Phys.* **17**, 332–336 (2021).
- Larkin, J., Jonsson, M., Justice, D. & Guerreschi, G. G. Evaluation of QAOA based on the approximation ratio of individual samples. *Quantum Sci. Technol.* **7**, 045,014 (2022).
- Pelofske, E., Bärtschi, A. & Eidenbenz, S. In *International Conference on High Performance Computing* 240–258 (Springer, 2023).
- Bittel, L. & Kliesch, M. Training variational quantum algorithms is NP-hard. *Phys. Rev. Lett.* **127**, 120502 (2021).
- Anschuetz, E. R. Critical points in quantum generative models. In *proceedings International Conference on Learning Representations*, <https://openreview.net/forum?id=2f1z55GVQN> (2022).
- McClellan, J. R., Boixo, S., Smelyanskiy, V. N., Babbush, R. & Neven, H. Barren plateaus in quantum neural network training landscapes. *Nat. Commun.* **9**, 4812 (2018).
- Ortiz Marrero, C., Kieferová, M. & Wiebe, N. Entanglement-induced barren plateaus. *PRX Quantum* **2**, 040316 (2021).
- Wang, S. et al. Noise-induced barren plateaus in variational quantum algorithms. *Nat. Commun.* **12**, 6961 (2021).
- Arrasmith, A., Cerezo, M., Czarnik, P., Cincio, L. & Coles, P. J. Effect of barren plateaus on gradient-free optimization. *Quantum* **5**, 558 (2021).
- Verdon, G. et al. Learning to learn with quantum neural networks via classical neural networks. Preprint at <https://arxiv.org/abs/1907.05415> (2019).
- Alam, M., Ash-Saki, A. & Ghosh, S. in *2020 Design, Automation & Test in Europe Conference & Exhibition (DATE)* 686–689 (IEEE, 2020).
- Khairy, S., Shaydulin, R., Cincio, L., Alexeev, Y. & Balaprakash, P. Learning to optimize variational quantum circuits to solve combinatorial problems. In *Proceedings of the AAAI Conference on Artificial Intelligence*, **34**, 2367–2375 (2020).
- Jain, N., Coyle, B., Kashefi, E. & Kumar, N. Graph neural network initialisation of quantum approximate optimisation. *Quantum* **6**, 861 (2022).
- Shaydulin, R., Marwaha, K., Wurtz, J. & Lotshaw, P. C. In *2021 IEEE/ACM Second International Workshop on Quantum Computing Software (QCS)*, Vol. 50, 64–71 (IEEE, 2021).
- Moussa, C., Wang, H., Bäck, T. & Dunjko, V. Unsupervised strategies for identifying optimal parameters in quantum approximate optimization algorithm. *EPJ Quantum Technol.* **9**, 11 (2022).
- Amosy, O., Danzig, T., Porat, E., Chechik, G. & Makmal, A. Iterative-free quantum approximate optimization algorithm using neural networks. Preprint at <https://arxiv.org/abs/2208.09888> (2022).
- Yao, J., Li, H., Bukov, M., Lin, L. & Ying, L. In *Mathematical and Scientific Machine Learning* 49–64 (PMLR, 2022).
- Xie, N., Lee, X., Cai, D., Saito, Y. & Asai, N. In *Journal of Physics: Conference Series*, Vol. 2595, 012001 (IOP Publishing, 2023).
- Tate, R., Farhadi, M., Herold, C., Mohler, G. & Gupta, S. Bridging classical and quantum with SDP initialized warm-starts for QAOA. *ACM Trans. Intell. Syst. Technol.* **4**, 1–39 (2023).
- Campos, E., Rabinovich, D., Akshay, V. & Biamonte, J. Training saturation in layerwise quantum approximate optimization. *Phys. Rev. A* **104**, L030401 (2021).
- Shaydulin, R., Lotshaw, P. C., Larson, J., Ostrowski, J. & Humble, T. S. Parameter transfer for quantum approximate optimization of weighted maxcut. *ACM Trans. Quantum Comput.* **4**, 1–15 (2023).
- Sack, S. H., Medina, R. A., Kueng, R. & Serbyn, M. Recursive greedy initialization of the quantum approximate optimization algorithm with guaranteed improvement. *Phys. Rev. A* **107**, 062404 (2023).
- Mele, A. A., Mbeng, G. B., Santoro, G. E., Collura, M. & Torta, P. Avoiding barren plateaus via transferability of smooth solutions in a hamiltonian variational ansatz. *Phys. Rev. A* **106**, L060401 (2022).
- Norouzi, M., Ranjbar, M. & Mori, G. Stacks of convolutional restricted boltzmann machines for shift-invariant feature learning. In *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2009*, 2735–2742 (2009).

35. Larocca, M., Ju, N., García-Martín, D., Coles, P. J. & Cerezo, M. Theory of overparametrization in quantum neural networks. *Nat. Comput. Sci.* **3**, 542–551 (2023).
36. Kim, J., Kim, J. & Rosa, D. Universal effectiveness of high-depth circuits in variational eigenproblems. *Phys. Rev. Res.* **3**, 023203 (2021).
37. Frazier, P. I. A tutorial on Bayesian optimization. Preprint at <https://arXiv.org/abs/1807.02811> (2018).
38. Eriksson, D., Pearce, M., Gardner, J., Turner, R. D. & Poloczek, M. In *Advances in Neural Information Processing Systems*, Vol. 32 (eds Wallach, H.) (Curran Associates, Inc., 2019).
39. Letham, B., Karrer, B., Ottoni, G. & Bakshy, E. Constrained Bayesian optimization with noisy experiments. *Bayesian Anal.* **14**, 495–519 (2019).
40. Letham, B., Calandra, R., Rai, A. & Bakshy, E. In *Advances in Neural Information Processing Systems*, Vol. 33 (eds Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. & Lin, H.) 1546–1558 (Curran Associates, Inc., 2020).
41. Yuan, Y. X. Recent advances in trust region algorithms. *Math. Program.* **151**, 249–281 (2015).
42. Powell, M. J. D. *A Direct Search Optimization Method That Models the Objective and Constraint Functions by Linear Interpolation* 51–67 (Springer, 1994).
43. Kingma, D. P. & Ba, J. Adam: A method for stochastic optimization. In *Proceedings International Conference on Learning Representations (ICLR)* (eds Bengio, Y. & LeCun, Y.) (2015).
44. Spall, J. Implementation of the simultaneous perturbation algorithm for stochastic optimization. *IEEE Trans. Aerosp. Electron. Syst.* **34**, 817–823 (1998).
45. Tibaldi, S., Vodola, D., Tignone, E. & Ercolessi, E. Bayesian optimization for QAOA. *IEEE Trans. Quantum Eng.* **4**, 1–11 (2023).
46. Self, C. N. et al. Variational quantum algorithm with information sharing. *Npj Quantum Inf.* **7**, 116 (2021).
47. Tamiya, S. & Yamasaki, H. Stochastic gradient line Bayesian optimization for efficient noise-robust optimization of parameterized quantum circuits. *Npj Quantum Inf.* **8**, 90 (2022).
48. Shaffer, R., Kocia, L. & Sarovar, M. Surrogate-based optimization for variational quantum algorithms. *Phys. Rev. A* **107**, 032415 (2023).
49. Gelbart, M. A., Snoek, J. & Adams, R. P. Bayesian optimization with unknown constraints. In *Proceedings of the Thirtieth Conference on Uncertainty in Artificial Intelligence* 250–259 (Arlington, Virginia, USA Quebec City, Quebec, Canada, AUAI Press, 2014).
50. Bravyi, S., Sheldon, S., Kandala, A., Mckay, D. C. & Gambetta, J. M. Mitigating measurement errors in multiqubit experiments. *Phys. Rev. A* **103**, 042605 (2021).
51. Nation, P. D., Kang, H., Sundaresan, N. & Gambetta, J. M. Scalable mitigation of measurement errors on quantum computers. *PRX Quantum* **2**, 040326 (2021).
52. Temme, K., Bravyi, S. & Gambetta, J. M. Error mitigation for short-depth quantum circuits. *Phys. Rev. Lett.* **119**, 180509 (2017).
53. Li, Y. & Benjamin, S. C. Efficient variational quantum simulator incorporating active error minimization. *Phys. Rev. X* **7**, 021050 (2017).
54. Eriksson, D. & Jankowiak, M. In *Uncertainty in Artificial Intelligence* 493–503 (PMLR, 2021).
55. Nayebi, A., Munteanu, A. & Poloczek, M. In *Proceedings of the 36th International Conference on Machine Learning, Proceedings of Machine Learning Research*, Vol. 97 (eds Chaudhuri, K. & Salakhutdinov, R.) 4752–4761 (PMLR, 2019).
56. Martinez-Cantin, R., Tee, K. & McCourt, M. In *International Conference on Artificial Intelligence and Statistics* 1722–1731 (PMLR, 2018).
57. Fröhlich, L., Klenske, E., Vinogradskaya, J., Daniel, C. & Zeilinger, M. In *International Conference on Artificial Intelligence and Statistics* 2262–2272 (PMLR, 2020).
58. Daulton, S. et al. In *International Conference on Machine Learning* 4831–4866 (PMLR, 2022).
59. Dave, A. et al. Autonomous optimization of non-aqueous li-ion battery electrolytes via robotic experimentation and machine learning coupling. *Nat. Commun.* **13**, 5454 (2022).
60. Zhang, Y., Apley, D. W. & Chen, W. Bayesian optimization for materials design with mixed quantitative and qualitative variables. *Sci. Rep.* **10**, 1–13 (2020).
61. Cheng, L., Yang, Z., Liao, B., Hsieh, C. & Zhang, S. Odbo: Bayesian optimization with search space prescreening for directed protein evolution. Preprint at <https://arXiv.org/abs/2205.09548> (2022).
62. Zhang, S. X., Hsieh, C. Y., Zhang, S. & Yao, H. Differentiable quantum architecture search. *Quantum Sci. Technol.* **7**, 045023 (2022).
63. Weidinger, A., Mbeng, G. B. & Lechner, W. Error mitigation for quantum approximate optimization. *Phys. Rev. A* **108**, 032408 (American Physical Society, 2023) <https://doi.org/10.1103/PhysRevA.108.032408>.
64. Mockus, J. *Bayesian Approach to Global Optimization: Theory and Applications*, Vol. 37 (Springer Science & Business Media, 2012).
65. Rasmussen, C. E. & Williams, C. K. I. *Gaussian Processes for Machine Learning* (MIT Press, 2006).
66. Gardner, J. R., Pleiss, G., Bindel, D., Weinberger, K. Q. & Wilson, A. G. In *Advances in Neural Information Processing Systems* (eds Bengio, S. et al.) (Curran Associates, Inc., 2018).
67. Srinivas, N., Krause, A., Kakade, S. & Seeger, M. Gaussian process optimization in the bandit setting: no regret and experimental design. In *Proceedings of the 27th International Conference on International Conference on Machine Learning* 1015–1022 (Omnipress, Madison, WI, USA, Haifa, Israel, 2010).
68. Zhang, S. X. et al. TensorCircuit: a quantum software framework for the NISQ era. *Quantum* **7**, 912 (2023).

Acknowledgements

We appreciate Jonathan Allcock’s helpful discussions and Zhaofeng Ye’s suggestions on the schematic workflow’s graphics design. We also thank the technical support from Maochun Dai, Zhenxing Zhang, and Dengfeng Li for the usage of the quantum device.

Author contributions

L.C., Y.C. and S.-X.Z. developed the idea, implemented the simulation and experiments, analyzed the results, and wrote the manuscript. L.C. performed the numerical experiments. L.C., Y.C., S.-X.Z. and S.Z. participated in the discussion for the research and revising the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s42005-024-01577-x>.

Correspondence and requests for materials should be addressed to Lixue Cheng, Shi-Xin Zhang or Shengyu Zhang.

Peer review information *Communications Physics* thanks the anonymous reviewers for their contribution to the peer review of this work.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher’s note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2024