

Quantum Annealing Initialization of the Quantum Approximate Optimization Algorithm

Stefan H. Sack and Maksym Serbyn

IST Austria, Am Campus 1, 3400 Klosterneuburg, Austria

June 28, 2021

The quantum approximate optimization algorithm (QAOA) is a prospective near-term quantum algorithm due to its modest circuit depth and promising benchmarks. However, an external parameter optimization required in the QAOA could become a performance bottleneck. This motivates studies of the optimization landscape and search for heuristic ways of parameter initialization. In this work we visualize the optimization landscape of the QAOA applied to the MaxCut problem on random graphs, demonstrating that random initialization of the QAOA is prone to converging to local minima with sub-optimal performance. We introduce the initialization of QAOA parameters based on the Trotterized quantum annealing (TQA) protocol, parameterized by the Trotter time step. We find that the TQA initialization allows to circumvent the issue of false minima for a broad range of time steps, yielding the same performance as the best result out of an exponentially scaling number of random initializations. Moreover, we demonstrate that the optimal value of the time step coincides with the point of proliferation of Trotter errors in quantum annealing. Our results suggest practical ways of initializing QAOA protocols on near-term quantum devices and reveal new connections between QAOA and quantum annealing.

1 Introduction

Recent technological advances have led to a large number of implementations [1–4] of so-called Noisy Intermediate-Scale Quantum (NISQ) devices [5]. These machines, which allow to manipulate a small number of imperfect qubits with limited coherence time, inspired the search for practical quantum algorithms. The quantum approximate optimization algorithm (QAOA) [6]

Stefan H. Sack: stefan.sack@ist.ac.at

Maksym Serbyn : maksym.serbyn@ist.ac.at

has emerged as a promising candidate for such NISQ devices [7–9].

The QAOA is a variational hybrid quantum algorithm where the classical computer operates a NISQ device. The computer is responsible for the optimization of the cost function over a set of variational parameters. The cost function is calculated using a NISQ device that prepares a quantum state corresponding to chosen parameters and performs quantum measurements. In QAOA of depth p the wave function is prepared by a unitary circuit parametrized by $2p$ parameters, see Fig. 1(a). Each of the p layers consist of two unitaries: the first is generated by a classical Hamiltonian H_C that encodes the cost function of a combinatorial optimization problem, and the second is generated by the mixing quantum Hamiltonian, H_B .

While the $p = 1$ limit of QAOA allows for analytic considerations and derivation of performance guarantees [6], subsequent work suggested that higher depth p may be required in order to achieve a quantum advantage [8, 10]. However, increasing p leads to a progressively more complex optimization landscape, that is characterized by a large number of local suboptimal minima [7, 9, 11, 12], see Fig. 1(c). The convergence of classical optimization algorithms into such sub-optimal solutions was demonstrated to be a potential bottleneck of QAOA performance as finding a nearly optimal minimum usually requires exponential in p number of initializations of the classical optimization algorithm [6, 7]. Note, that the problem of sub-optimal local minima is different from that of barren plateaus [13, 14], i.e. large regions in parameter space with vanishing gradients, since barren plateaus are associated with circuit depths p polynomial in system size N [15], beyond what is typically considered in the QAOA.

The complexity of the energy landscape of large- p QAOA has motivated the search for heuristic ways of improving the convergence to a (nearly) optimal minimum values of the variational parameters. Recent work has demonstrated a concentration of the QAOA landscape for typical problem instances [16], which implies the existence of a typical landscape and hints at the fact that the same variational parameter choice may

work between different problem instances or sizes. A particular example of such a heuristic was proposed in Ref. [7] which constructs a good initialization for the QAOA at level $p + 1$ using the solution at level p , thus requiring a polynomial in p number of optimization runs. Other approaches, such as reusing parameters from similar graphs [12], using an initial state that encodes the solution of a relaxed problem [17], or utilizing machine learning techniques to predict QAOA parameters [18, 19] were also proposed.

In this work we propose a different approach to the QAOA initialization, based on the relation between QAOA and the quantum annealing algorithm. Quantum annealing uses adiabatic time evolution to find the lowest energy state of H_C , but often requires unfeasible evolution time T [20]. We explore the observation that Trotterization of unitary evolution in quantum annealing provides a particular choice of parameters for the QAOA [6]. This leads us to introduce a one-parameter family of Trotterized quantum annealing (TQA) initializations for QAOA, controlled by the time step or, equivalently, total time used in adiabatic evolution.

The central result of our work is the demonstration that TQA initialization for QAOA gives comparable performance to the search over an exponentially scaling number of random initializations. To this end, we establish that TQA initialization leads to convergence of the QAOA to a nearly optimal minimum for a certain range of time steps, see Fig. 1(c) for visualization. Furthermore, we identify the optimal time step of the TQA initialization and suggest a purely experimental way of fixing this parameter.

Our work reveals a connection between intermediate- p QAOA and short-time quantum annealing. Previous studies [6–8] established a correspondence between quantum annealing with long annealing times and the QAOA protocol with large p (potentially increasing exponentially with the problem size). More recent work proposed quantum annealing inspired initialization strategies for the so-called ‘bang-bang’ modification of the QAOA [21] that however also corresponds to large circuit depths. Our work is different from this context, since we establish that the best performance is achieved for a very *coarse discretization* of quantum annealing, resulting in a realistic circuit depth. We show the existence of an optimal step for TQA discretization that does not depend on problem size and QAOA depth. This suggests an intimate relation between QAOA and TQA, since the optimal value of the time step is in close correspondence to the point where proliferation of Trotter error occurs in TQA [22].

The remainder of the paper is organized as follows: in Section 2 we introduce the QAOA, visualize its optimization landscape and show that most random initial-

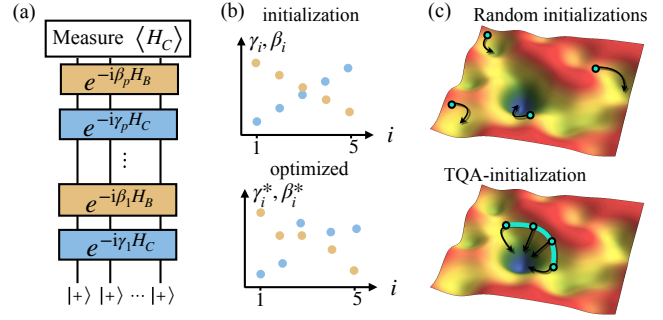


Figure 1: (a) The circuit that prepares a quantum state in the QAOA is parametrized by a set of $2p$ angles γ_i, β_i . For the MaxCut problem, that is considered in the main text, the unitaries can be expressed using single and two qubit gates that are readily available on current NISQ devices. (b) The optimization of $\langle H_C \rangle$ is launched from a certain guess of parameters, state preparation and measurements are iterated until the algorithm converges to a set of optimized angles γ_i^*, β_i^* . (c) The cartoon of the cost function $\langle H_C \rangle$ landscape as a function of variational parameters shows that random initializations are prone to converge to sub-optimal local minima. In contrast, the family of TQA initializations proposed in this work converges to the (nearly) optimal minimum.

izations concentrate around sub-optimal local minima. Next, in Section 3 we discuss TQA and the corresponding initialization and show that it avoids converging at sub-optimal local optima. Finally, in Section 4 we summarize the results, discuss its implications and potential future work.

2 Optimization landscape of QAOA

2.1 QAOA for MaxCut problems

As we discussed in the Introduction, the QAOA is often applied to hard combinatorial optimization problems. In what follows we concentrate on the problem of finding a maximal cut (MaxCut) in a given graph which has become one of standard tasks used to benchmark the QAOA [7, 9]. Finding the maximum cut is an *NP*-hard combinatorial optimization problem, though efficient classical algorithms exist that yield good approximate solutions. Notably, the Goemans-Williamson algorithm yields a cut that is at least 88% of the size of the maximum cut in polynomial time [23].

Given a graph $G = (V, E)$ with vertices $V = 1, 2, \dots, N$ and edges $E = \{\langle i, j \rangle\}$, the *maximal cut* is defined as the partition that splits the vertices into two groups, maximizing the number of edges that connect vertices from different groups. Mathematically, such partition amounts to finding the global minimum of a cost function, $C(\vec{z}) = \sum_{\langle i, j \rangle \in E} z_i z_j$, where the binary variables z_i correspond to the vertices of the graph, and

their value $z_i = \pm 1$ encodes which partition the given vertex i belongs to. The cost function $C(\vec{z})$ can be mapped into a classical spin Hamiltonian by promoting the binary variable z_i to the quantum spin-1/2 operator σ_i^z . The resulting Hamiltonian,

$$H_C = \sum_{\langle i,j \rangle \in E} \sigma_i^z \sigma_j^z, \quad (1)$$

operates on N spins that reside on the vertices V of the corresponding graph and interact with each other when connected by an edge.

The QAOA uses a NISQ device to prepare the following quantum state [see Fig. 1(a)]:

$$|\vec{\gamma}, \vec{\beta}\rangle = \prod_{i=1}^p e^{-i\beta_i H_B} e^{-i\gamma_i H_C} |0\rangle_B, \quad (2)$$

where H_C is classical Hamiltonian introduced above, and $H_B = -\sum_i^N \sigma_i^x$ the mixing Hamiltonian, as proposed by Farhi *et al.* [6]. Both operators operate on the Hilbert space corresponding to N spins or, equivalently, qubits, and the initial state $|0\rangle_B = |+\rangle^{\otimes N}$ corresponds to all qubits pointing along x -direction, thus yielding the ground state of H_B . The variational parameters are obtained by minimizing the expectation value $\langle H_C \rangle_{\vec{\gamma}, \vec{\beta}}$ as:

$$(\vec{\gamma}^*, \vec{\beta}^*) = \arg \min_{(\vec{\gamma}, \vec{\beta})} \langle \vec{\gamma}, \vec{\beta} | H_C | \vec{\gamma}, \vec{\beta} \rangle, \quad (3)$$

which is typically carried out with numerical optimization routines. To benchmark the QAOA it is useful to define the approximation ratio,

$$r_{\vec{\gamma}, \vec{\beta}} = \frac{\langle \vec{\gamma}, \vec{\beta} | H_C | \vec{\gamma}, \vec{\beta} \rangle}{C_{\min}}, \quad (4)$$

which quantifies how close the expectation value of the classical Hamiltonian over the QAOA wave function is to the ground state energy of H_C , denoted as C_{\min} . For QAOA at depth $p = 1$ the algorithm is guaranteed to find a cut that is at least 69% the size of the optimal cut [6], while for $p > 1$ analytic results are limited [24].

The performance of the QAOA is typically investigated over an ensemble of graphs rather than an individual realization. Below we focus on the ensemble of random 3-regular graphs, where each vertex is connected to three other vertices chosen at random. However, in the Appendix we also consider weighted 3-regular graphs and Erdős-Rényi graph ensembles in order to illustrate the general applicability of our results.

2.2 Visualizing optimization landscape

The performance of the classical optimization in Eq. (3) strongly depends on the properties of the optimization

landscape. While this landscape can be readily visualized for $p = 1$, the dependence of approximation ratio $r_{\vec{\gamma}, \vec{\beta}}$ on $2p$ angles parametrizing QAOA was suggested to become progressively more complex for larger values of p . In order to visualize the properties of this high-dimensional landscape, we focus below on points where $1 - r_{\vec{\gamma}, \vec{\beta}}$ achieves (local) minima.

We quantify properties of minima using two different characteristics. First, we measure the difference between the approximation ratio of the given minimum characterized by angles $\vec{\gamma}, \vec{\beta}$ and the global minimum characterized by angles $\vec{\gamma}^*, \vec{\beta}^*$, $\Delta r_{\vec{\gamma}, \vec{\beta}} = r_{\vec{\gamma}^*, \vec{\beta}^*} - r_{\vec{\gamma}, \vec{\beta}}$. This definition implies that the smallest possible value of $\Delta r_{\vec{\gamma}, \vec{\beta}}$ is 0, and larger values of $\Delta r_{\vec{\gamma}, \vec{\beta}}$ corresponds to local minima with poor performance (i.e. much larger value of cost function) compared to the global minimum. The second characteristic measures the distance between minima in parameter space,

$$d_{\vec{\gamma}, \vec{\beta}} = \sum_{i=1}^p (|\beta_i - \beta_i^*|_{\frac{\pi}{2}} + |\gamma_i - \gamma_i^*|_{\pi}), \quad (5)$$

where $|\dots|_{\alpha}$ denotes the absolute value modulo α which takes into account symmetries, see Appendix A.

We calculate values of $\Delta r_{\vec{\gamma}, \vec{\beta}}$ and $d_{\vec{\gamma}, \vec{\beta}}$ numerically. For a given graph realization we use 2^p different random initializations of variational parameters $\vec{\gamma}, \vec{\beta}$ and optimize them using the iterative BFGS algorithm [25–28]. The algorithm is accessed via the `scipy.optimize` Python module with default parameters [29]. Convergence is achieved when the norm of the gradient is less than 10^{-5} , maximum number of iterations is set to $400p$, where p is the QAOA depth. In our simulations the routine typically converged before using up the maximum number of allowed iterations. We use the converged angles with the lowest value of $1 - r_{\vec{\gamma}, \vec{\beta}}$ as an estimate for the global minimum γ_i^*, β_i^* .

Figure 2 visualizes the structure of local minima via the joint probability distribution of $\Delta r_{\vec{\gamma}, \vec{\beta}}$ and $d_{\vec{\gamma}, \vec{\beta}}$ for 50 different graphs using Kernel Density Estimation [30, 31]. We observe that for QAOA with $p = 5$ the most typical local minima reached from random initialization are far away from the best minimum (corresponding to $\Delta r_{\vec{\gamma}^*, \vec{\beta}^*} = 0$ and $d_{\vec{\gamma}^*, \vec{\beta}^*} = 0$) both in terms of quality of approximation ratio and parameter values. While this figure illustrates a particular choice of system size and QAOA depth, a similar trend is observed for different N , p , and other graph ensembles, see Appendix A.

The tendency of random initialization to converge to suboptimal solutions highlights the importance of better initialization methods. In the next section we investigate a family of initializations inspired by quantum annealing and demonstrate that it achieves a good approximation ratio with a suitable choice of parameters.

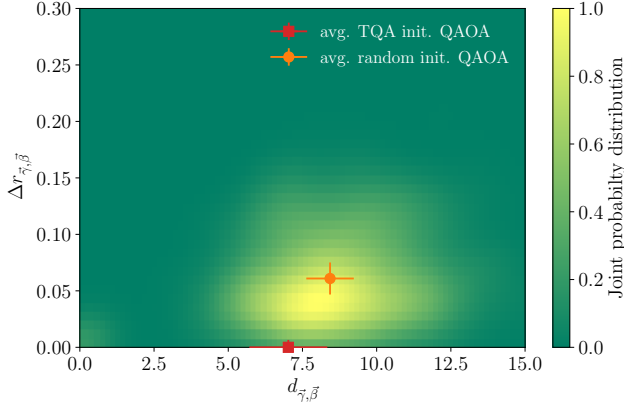


Figure 2: Joint probability distribution of distance to the global minimum in parameter space $d_{\gamma,\beta}$ and in terms of approximation ratio $\Delta r_{\gamma,\beta}$ reveals that the most probable outcome of random initialization is a convergence to sub-optimal local minima (yellow region). The orange dot corresponds to average values of $d_{\gamma,\beta}$, $\Delta r_{\gamma,\beta}$ for random initialization. In contrast, TQA initialization leads to a local minimum with a better approximation ratio that occasionally outperforms the best random initialization (red square, shifted from slightly negative values to $\Delta r_{\gamma,\beta} = 0$ for improved visibility). The data is averaged over 50 random unweighted 3-regular graphs with $N = 12$ vertices and QAOA at level $p = 5$.

3 Trotterized quantum annealing as initialization

3.1 Optimal time for TQA

Quantum annealing [32, 33] was among the first algorithms proposed for quantum computing [34, 35], and was demonstrated to be universal for $T \rightarrow \infty$ and equivalent to digital quantum computing [36]. The general idea of quantum annealing is to prepare the ground state $|0\rangle_C$ of a classical Hamiltonian H_C starting from the ground state $|0\rangle_B$ of the mixing Hamiltonian H_B using adiabatic time evolution under the Hamiltonian $H(t) = (1 - t/T)H_B + (t/T)H_C$. Practical execution of quantum annealing on NISQ devices requires discretization to represent such unitary evolution via a sequence of gates, resulting in the TQA algorithm. The first order Suzuki-Trotter decomposition allows to approximate the time evolution with $H(t)$ over time interval Δt as $e^{-i\Delta t H(t)} \approx e^{-i\beta H_B} e^{-i\gamma H_C} + \mathcal{O}(\Delta t^2)$ with $\beta = (1 - t/T)\Delta t$ and $\gamma = (t/T)\Delta t$.

Applying such decomposition to the quantum annealing protocol that is uniformly discretized on a grid of evolution times $t_i = i\Delta t$ with $i = 1, \dots, p$ and time step $\Delta t = T/p$, we obtain the unitary circuit equivalent to

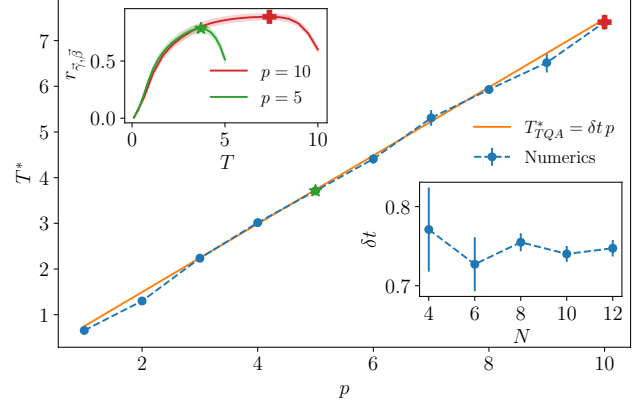


Figure 3: Optimal time of TQA evolution T^* increases linearly with number of discretization steps p . Top inset illustrates that optimal performance of TQA at time T^* is followed by the rapid decrease in approximation ratio at longer times T^* . Data is shown for $N = 12$. Bottom inset shows finite size scaling of the time step δt , determined by the slope of the T^* vs p dependence, that assumes approximately constant value with the graph size. All averaging is performed over 50 random instances of unweighted 3-regular graphs.

the depth- p QAOA ansatz (2) with angles being

$$\gamma_i = \frac{i}{p} \Delta t, \quad \beta_i = \left(1 - \frac{i}{p}\right) \Delta t. \quad (6)$$

In what follows we refer to such choice of angles as TQA initialization, controlled by the time step Δt at a fixed depth p .

The mapping between TQA and QAOA along with the universality of quantum annealing for $T \rightarrow \infty$ was previously used as an argument for the existence of good QAOA protocols at depths $p \rightarrow \infty$ [6]. Typically the required evolution time of quantum annealing is inversely proportional to the square of the minimal energy gap $T \propto \Delta^{-2}$ encountered in the Hamiltonian $H(t)$ over the time evolution. Numerous studies established that the time required for a good performance often blows up exponentially due to the encounter of exponentially small gaps in N [20].

In contrast to previous studies, we investigate TQA performance in a different setting that is motivated by its subsequent usage as a QAOA initialization. The QAOA is characterized by a fixed circuit depth, p . Therefore, we fix p and study the performance of TQA as a function of total time T or, equivalently, time step Δt , related as $T = p \Delta t$. Generally the performance of quantum annealing tends to increase with the total annealing time. However in case of fixed p , longer annealing time corresponds to a coarser discretization, which leads to larger Trotter errors that scale proportionally to $\mathcal{O}(\Delta t^2)$ at small values of Δt . It is the interplay between increased efficiency and Trotter errors that leads

to the existence of an optimal annealing time in the present setting. This is illustrated in Fig. 3 (top inset), where the approximation ratio for the TQA protocol increases with T for small times, reaching a maximum at time T^* followed by a sharp downturn. The sharp decrease of QA performance after T^* was reported by Heyl *et al.* [22], who attributed it to a phase transition caused by a proliferation of Trotter errors.

Main panel of Fig. 3 reveals a linear scaling of the optimal time T^* with the number of time steps p . This is equivalent to the existence of an *optimal time step* δt , that determines T^* as

$$T_{\text{TQA}}^* = \delta t p. \quad (7)$$

The bottom inset in Fig. 3 shows that the time step δt defined as a slope of a linear fit of T^* with p converges with the problem size N . This gives a strong evidence that δt is a well-defined quantity in the thermodynamic limit $N \rightarrow \infty$. For the family of the 3-regular graphs considered here we observe that the optimal time step tends to value $\delta t \approx 0.75$. The existence of an optimal time step that is of order one holds for three other graph ensembles, considered in Appendix B, although the numerical value of this time step depends on the specific graph ensemble.

We use the TQA initialization in Eq. (6) with time step $\Delta t = 0.75$ for the QAOA and observe in Fig. 2 that it allows to avoid the local minima and helps the QAOA to converge to a minimum that is very close to the global minimum in terms of approximation ratio. This result motivates the systematic analysis of the performance of the TQA initialization.

3.2 TQA initialization of QAOA

We continue with a detailed study of the TQA initialization defined in Eq. (6) as a function of time T at fixed p . The green line in Fig. 4(a) reveals that the approximation ratio remains constant for a range of times, denoted as $[T_{\min}^*, T_{\max}^*]$. This figure shows results for $p = 5$ QAOA applied to graphs with $N = 12$ vertices, but a similar trend holds for other values of depth, problem sizes, and graph ensembles. The constant approximation ratio in a range of T is naturally explained by the convergence of parameter optimization routine to the same minimum for $T \in [T_{\min}^*, T_{\max}^*]$, see cartoon in Fig. 1(c). In order to discriminate between different times in the above range, we study the distance between initialization parameters and optimized values of $\vec{\gamma}, \vec{\beta}$. The red line Fig. 4(a) shows that this distance has a well-pronounced minimum at a time denoted as T_d^* that is contained within the same interval $[T_{\min}^*, T_{\max}^*]$. The TQA initialization with time T_d^* is

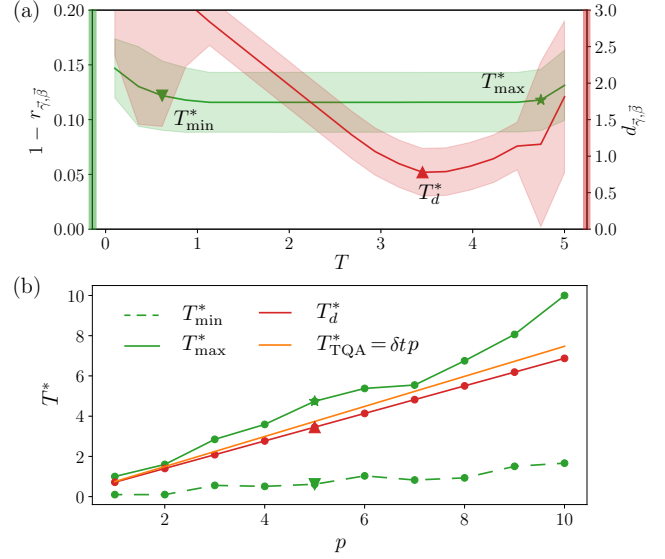


Figure 4: (a) Approximation ratio of the $p = 5$ QAOA as a function of TQA initialization time T reveals that a range of initialization times $[T_{\min}^*, T_{\max}^*]$ (green triangle and star) yield the performance within 1% of the minimal $1 - r_{\vec{\gamma}, \vec{\beta}}$. On the other hand, the study of the distance between the TQA initialization and the converged value of the angles reveals the existence of a time T_d^* where the QAOA performs the smallest parameter updates. (b) All three times T_{\min}^* , T_{\max}^* , and T_d^* defined in panel (a) increase linearly with QAOA circuit depth p . Moreover, T_d^* is very close to the time where the TQA protocol itself achieves optimal performance, T_{TQA}^* , see Fig. 2. Data was obtained for $N = 12$ and averaged over 50 random graphs.

closest to the local minimum achieved from it in a sense of distance defined in Eq. (5).

All three times T_{\min}^* , T_{\max}^* , and T_d^* were defined above using the QAOA with fixed depth p . Figure 4(b) reveals that all three times scale approximately linearly with p . This allows to define a range of time steps for the TQA initialization that yield the same performance of optimized QAOA, $\Delta t \in [0.16, 0.92]$ for the present graph ensemble. Moreover, the time T_d^* nearly coincides with the optimal TQA time $T_{\text{TQA}}^* = \delta t p$ obtained in the previous section, implying that $\Delta t = \delta t = 0.75$ is the optimal value of time step. This result also holds for the MaxCut problem on other graph families, see Appendix.

The similarity between the optimal time of the TQA protocol to the time where the angular distance $d_{\vec{\gamma}, \vec{\beta}}$ between the initial and final protocol is minimized, suggests that the performance of the QAOA is bounded by the same phase transition that occurs in TQA [22]. However, the QAOA is able to provide a significant improvement over TQA by doing additional optimizations of variational parameters. Recent work [7] suggested that such performance improvement may be due to uti-

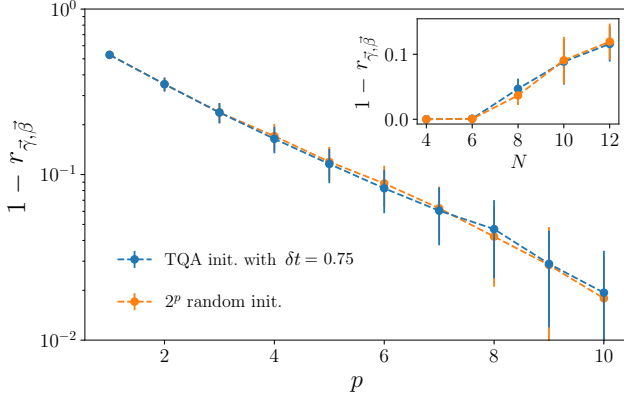


Figure 5: A single optimization run of the QAOA with TQA initialization with time $T = \delta t p$ yields equivalent performance to the best out of 2^p random initializations. System size is $N = 12$. Inset reveals that the comparable performance persists over the entire range of considered system sizes, circuit depth is $p = 10$. Averaging was performed over 50 random graphs.

lization of “diabatic pumps” that allow to return the population from excited states back to the ground state. This could potentially explain the systematic deviation of the QA protocol from TQA initialization as seen in Fig. 8 in Appendix C.

Finally, we compare the performance of QAOA that used 2^p random initializations to the QAOA launched from TQA initialization with optimal time step δt . Surprisingly, Fig. 5 shows that TQA initialization yields the *same* performance as the best result for random initialization even for QAOA protocols with depth comparable to the problem size, N . Moreover, the inset of Fig. 5 illustrates that the excellent performance of TQA initialization holds true for a broad range of system sizes N , while Appendix D presents equally encouraging results for other graph ensembles. Note that the QAOA performance for fixed p decreases with system size N , which was attributed to the fact that the QAOA with fixed p cannot “probe” the whole graph. In order for the QAOA to achieve constant performance for increasing problem size N , the depth of QAOA should increase at least as $\log N$ [7].

4 Summary and discussion

Our central result is the establishment of a family of TQA initializations for the QAOA parametrized by a time step Δt . We find that TQA initialization allows the QAOA to find a solution close to the global optima for a broad range of parameter Δt . In this range our initialization scheme achieves results similar to the best outcome of 2^p random initializations, with a single optimization run. Moreover we establish a heuristic way to

identify the optimal Δt for the TQA initialization from the performance of the TQA protocol.

Our results open the door to more time-efficient practical implementations of the QAOA on NISQ devices. To this end, we propose a two-step practical NISQ algorithm that capitalizes on the success of TQA initialization and uses the heuristic results to establish an optimal value of the time step. The first two steps of Algorithm 1 implement the TQA protocol on a NISQ device, thus obtaining an estimate for the optimal time in the TQA initialization. This can be readily carried out on today’s NISQ devices [37]. The second part of the algorithm consists of running the QAOA optimization loop using values of variational parameters according to Eq. (6).

Algorithm 1 QAOA with TQA initialization

- 1: Implement QAOA ansatz with circuit depth p .
 - 2: Estimate time step δt using TQA:
find optimal time $T^* \leftarrow \arg \min_T \langle H_C \rangle_p$
and set $\delta t \leftarrow \frac{T^*}{p}$, see Fig. 3.
 - 3: Use TQA initialization $\gamma_i \leftarrow \frac{i}{p} \delta t$ and $\beta_i \leftarrow (1 - \frac{i}{p}) \delta t$.
 - 4: Run the QAOA parameter optimization, see Fig. 1.
-

Numerical simulations presented above suggest good performance of the above algorithm in the idealized case when presence of noise, gate errors, and other imperfections are neglected. Moreover, the fact that TQA initialization converges to a good minimum for the range of times (equivalently, time steps) $T \in [T_{\min}^*, T_{\max}^*]$, see Fig. 4, suggests that this algorithm has a high tolerance towards imperfections in determining the value of δt . Determining the performance of this algorithm on a real NISQ device or incorporating some of the imperfections into the numerical simulation remains an interesting open problem.

In our studies we restricted our attention to the Max-Cut problem and demonstrated success of our approach for three different random graph ensembles. We expect that these results also hold for other graph ensembles, provided that the concentration of the QAOA landscape is true [16]. It is also interesting to check if our findings hold true beyond the MaxCut problem. Furthermore, it will be interesting to study the finite size scaling for problem sizes $N > 12$ considered here using matrix product states [38] or neural-network quantum states [39, 40].

In addition to practical NISQ algorithms, our finding suggest a previously unknown connection between the QAOA at relatively small circuit depth and quantum annealing. The fact that quantum annealing inspired initializations belong to a basin of attraction of a high-quality minimum in the QAOA landscape, see Fig. 1(c),

invites a more comprehensive study of the QAOA landscape from this perspective. How many good quality minima typically exist in such landscape? How different are they from each other and what are their basins of attraction? Can one use other information measures such as entanglement or Fisher information [41] to characterize the QAOA landscape? Finding answers to such questions may lead to other prospective families of QAOA initializations.

While TQA provides a good initialization, the subsequent QAOA optimization is able to significantly improve the performance. Understanding the underlying mechanisms of such performance improvement is an outstanding challenge. In particular, there remains an intriguing possibility that the QAOA optimization routine implements some of the techniques, developed to improve the annealing fidelity, such as diabatic pumps [7], shortcuts to adiabaticity [42], and counterdiabatic driving [43, 44]. The fact that the optimal time step coincides with the point of proliferation of Trotter errors [22], thus effectively taking maximal possible value suggests interesting parallels to the Pontryagin’s minimum principle considered in context of variational quantum algorithms [45].

To conclude, we hope that TQA initialization of the QAOA established in this work will help to achieve practical quantum advantage by executing the QAOA on available devices and inspire future research that could lead to better understanding of what happens under the hood of QAOA optimization.

Data and code availability

Data is available upon reasonable request, a brief tutorial for the TQA initialization can be found in Ref. [46]

Acknowledgments

We would like to thank D. Abanin and R. Medina for fruitful discussions and A. Smith and I. Kim for valuable feedback on the manuscript. We acknowledge support by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (Grant Agreement No. 850899).

A Optimization landscape for different graph ensembles

We start by reviewing all graph ensembles used in the main text and Appendices. In particular, we focus on symmetries that allow to reduce the space of QAOA parameters.

3-regular unweighted graphs represent the graph ensemble considered in the main text. Each vertex is connected exactly to three other vertices chosen at random. In order to sample graphs from this ensemble we use the `networkx` Python package [47]. For 3-regular unweighted graphs the space of variational parameters can be restricted using the fact that the classical Hamiltonian has integer eigenvalues (thus γ_i are defined modulo π) and that shifting any of angles β_i by $\pi/2$ is equivalent to a spin flip of H_C that has no effect [7]. This allows to restrict $\beta_i \in [-\frac{\pi}{4}, \frac{\pi}{4})$ and $\gamma_i \in [-\frac{\pi}{2}, \frac{\pi}{2})$, and is reflected in the definition of distance in Eq. (5) in the main text.

3-regular weighted graphs are characterized by presence of random weights w_{ij} assigned to each edge $\langle i, j \rangle$. These weights are chosen to be $w_{ij} \in [0, 1)$. Presence of random weights does not allow to restrict the domain of γ_i angles as before, though restriction $\beta_i \in [-\frac{\pi}{4}, \frac{\pi}{4})$ still works. Therefore the analogue of Eq. (5) for this and other weighted ensembles reads $d_{\vec{\gamma}, \vec{\beta}}^{(w)} = \sum_{i=1}^p (|\beta_i - \beta_i^*|_{\frac{\pi}{2}} + |\gamma_i - \gamma_i^*|)$.

Erdős-Rényi graphs represent a random graph ensemble where two edges are connected on random with a fixed probability, chosen to be $q = 0.5$. In contrast to above examples, the fixed value of q implies that edge connectivity increases with number of vertices as qN . Erdős-Rényi graphs exhibit the same symmetries as 3-regular unweighted graphs.

The presence of an unbounded region of parameters γ_i in the weighted graph ensemble represents an additional challenge in visualizing the QAOA optimization landscape and choice of initialization parameter. In order to explore the importance of large values of $|\gamma_i|$, we consider the sequence of enlarged intervals $\gamma_i \in [-k\frac{\pi}{2}, k\frac{\pi}{2})$ with $k = 1, 2$. Figure 6 shows the joint probability distributions similar to Fig. 2. We see that for 3-regular weighted graphs the enlarged initialization interval $k = 2$ leads to a concentration of local optima further away from the global solution compared to the $k = 1$ interval. When we repeat the same analysis for Erdős-Rényi graphs, we observe that $\Delta r_{\vec{\gamma}, \vec{\beta}}$ is unaffected by the enlarged $k = 2$ interval. This numerically confirms the symmetry considerations from above and allows us to restrict $\vec{\gamma}$ to the $k = 1$ interval in all further analysis. For unweighted graphs such restriction relies on symmetry, and for weighted graphs this is motivated by the fact that an extended region of γ_i worsens the performance of random initialization in the QAOA.

B Optimal time for TQA

Below we discuss the dependence of the optimal time step δt of the TQA algorithm on the graph ensemble.

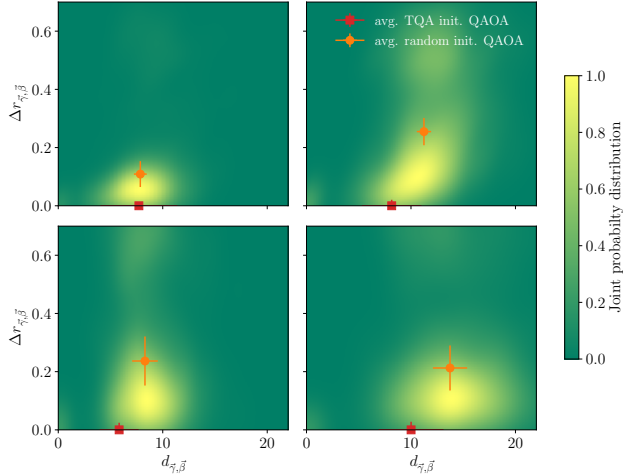


Figure 6: Comparing the joint probability distribution of the distance to the global minimum in parameter space $d_{\gamma,\beta}$ and in terms of approximation ratio $\Delta r_{\gamma,\beta}$ for weighted 3-regular (top) and Erdős-Rényi graphs with edge probability 0.5 (bottom) reveals that the distribution is dependent on the initialization interval for weighted 3-regular graphs. We initialize the parameters for $k = 1$ (left) and $k = 2$ (right) and observe that for weighted 3-regular graphs the enlarged interval leads to an increased spread of the local optimas in $\Delta r_{\gamma,\beta}$ (yellow region). The spread in $\Delta r_{\gamma,\beta}$ for Erdős-Rényi graphs remains largely unaffected, as expected from the symmetry considerations. Similarly to Fig. 2, red squares correspond to the QAOA minimum achieved from TQA initialization (shifted from small negative values of $\Delta r_{\gamma,\beta}$ to zero for improved visibility), orange dots correspond to the average performance of random initialization. Data is for 50 random graphs with $N = 10$ and $p = 5$.

An analytical *upper bound* on the number of Trotter steps p needed to approximate the time evolution with precision ϵ in terms of operator trace distance was obtained in Ref. [48]. Translating this bound into the scaling of δt we obtain $\delta t \propto 1/(\|H_C\|_F N)$, where $\|H_C\|_F$ is the Frobenius norm of the classical Hamiltonian. This norm exponentially diverges with N , suggesting very small values of δt at large system sizes. This is not surprising, since the bound of Ref. [48] operates on the distance between two many-body unitary operators. In contrast, the performance of the TQA algorithm is studied using the approximation ratio that quantifies how close the expectation value of the local observable H_C , is to the ground state energy.

The effect of Trotterization on local observables was considered in Ref. [22]. This work conjectured the existence of a finite value of the time step of order one, at which the discretization of time evolution fails to approximate the local observables. This value of the time step may be related to the convergence radius of the Baker-Campbell-Hausdorff series expansion, which

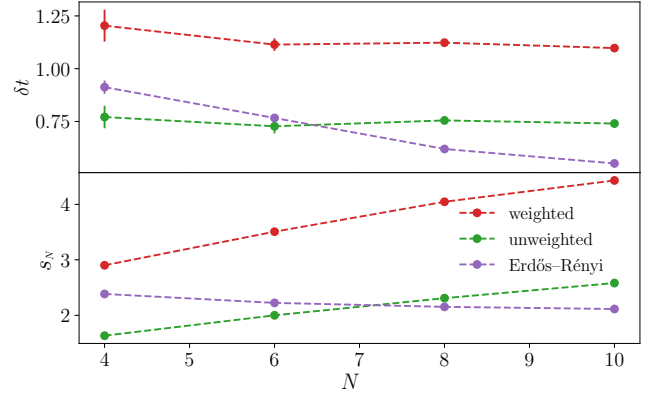


Figure 7: (Top) Optimal time step of TQA evolution δt is largely independent of system size and scales qualitatively similar to Eq. (8) shown in the bottom panel.

is governed by the norm of the classical Hamiltonian and its commutator with H_B . Phenomenologically, the Frobenius norm divided by the square root of Hilbert space dimension and problem size N ,

$$s_N = \frac{N2^{N/2}}{\|H_C\|_F}, \quad (8)$$

is expected to be N -independent in the thermodynamic limit.

Figure 7 compares the dependence of δt on the system size with the phenomenological scaling s_N defined in Eq. (8). We observe that the expression s_N qualitatively matches the numerical scaling that we observe for δt between different graph ensembles. In particular, the value of the time step is largest for weighted 3-regular graphs that are expected to have the smallest norm of the classical Hamiltonian. However, s_N fails to capture δt quantitatively, highlighting the need to develop a better analytical understanding of the point that governs the phase transition from localization to quantum chaos for local observables according to Ref. [22].

C Patterns in optimized parameters

The QAOA is inspired by TQA and is thus universal for $p \rightarrow \infty$. However, for finite p the converged QAOA parameters also display stark similarity to a QA protocol which was noticed in some earlier works [7, 8]. In Fig. 8 we compare the TQA initialization and final QAOA parameters. The QAOA parameters show only slight alterations at the beginning of the protocol and remain close to their original values throughout the rest of the protocol. This holds true for the three graph types that we considered in our analysis. In addition, the small variation between optimal parameters for different graph instances is in line with the concentration

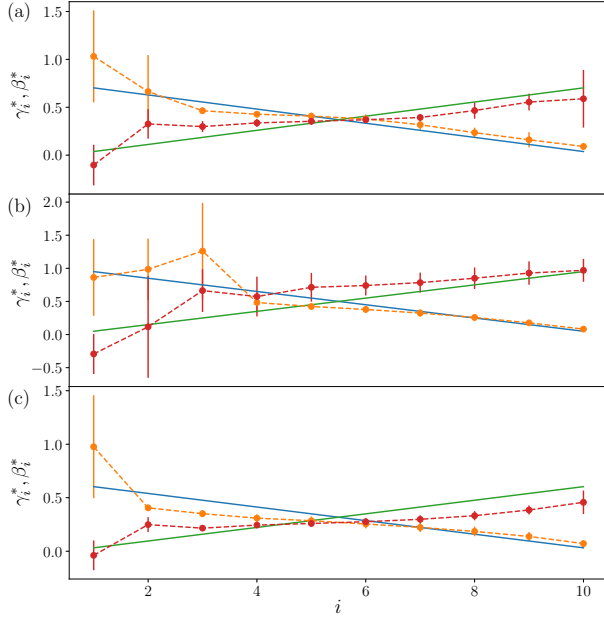


Figure 8: Converged parameters $\vec{\gamma}^*$ (red) and $\vec{\beta}^*$ (orange) show only slight alterations from the TQA initialization indicated by the green and blue lines respectively. The QAOA optimization modifies parameters at small i , while they remain TQA-like in the rest of the protocol. The results were averaged over 50 random unweighted 3-regular graphs (a), weighted 3-regular graphs (b) and Erdős-Rényi graphs (c), all data is for $p = 10$ and $N = 10$.

of the QAOA landscape demonstrated analytically at low p in Ref. [16].

D Random vs TQA initialization for other graph ensembles

In addition to the unweighted 3-regular graphs, discussed in the main text, we also test TQA initialization on weighted 3-regular graphs and Erdős-Rényi graphs. We find that TQA initialization yields the same performance as the best of random initializations for weighted 3-regular graphs, see Fig. 9. For Erdős-Rényi, TQA initialization even outperforms the best of 2^p random initializations.

References

[1] F. Arute *et al.*, Quantum Approximate Optimization of Non-Planar Graph Problems on a Planar Superconducting Processor, arXiv e-prints , arXiv:2004.04197 (2020), arXiv:2004.04197 [quant-ph] .

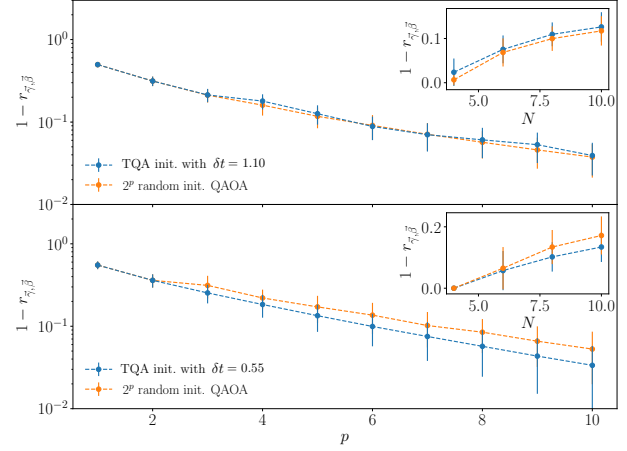


Figure 9: TQA initialization leads to the same QAOA performance as the best of 2^p random initializations for both weighted 3-regular graphs (top) and Erdős-Rényi graphs (bottom). We average the results over 50 graph realizations, the main plot was obtained for system size $N = 10$, inset is for circuit depth $p = 10$.

[2] F. Arute *et al.*, Hartree-Fock on a superconducting qubit quantum computer, *Science* **369**, 1084 (2020), arXiv:2004.04174 [quant-ph] .

[3] F. Arute *et al.*, Observation of separated dynamics of charge and spin in the Fermi-Hubbard model, arXiv e-prints , arXiv:2010.07965 (2020), arXiv:2010.07965 [quant-ph] .

[4] K. Wright *et al.*, Benchmarking an 11-qubit quantum computer, *Nature Communications* **10**, 5464 (2019), arXiv:1903.08181 [quant-ph] .

[5] J. Preskill, Quantum Computing in the NISQ era and beyond, arXiv e-prints , arXiv:1801.00862 (2018), arXiv:1801.00862 [quant-ph] .

[6] E. Farhi, J. Goldstone, and S. Gutmann, A Quantum Approximate Optimization Algorithm, arXiv e-prints , arXiv:1411.4028 (2014), arXiv:1411.4028 [quant-ph] .

[7] 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).

[8] G. E. Crooks, Performance of the Quantum Approximate Optimization Algorithm on the Maximum Cut Problem, arXiv e-prints , arXiv:1811.08419 (2018), arXiv:1811.08419 [quant-ph] .

[9] M. Willsch, D. Willsch, F. Jin, H. De Raedt, and K. Michielsen, Benchmarking the quantum approximate optimization algorithm, *Quantum Information Processing* **19**, 197 (2020), arXiv:1907.02359 [quant-ph] .

- [10] S. Bravyi, A. Kliesch, R. Koenig, and E. Tang, Obstacles to State Preparation and Variational Optimization from Symmetry Protection, arXiv e-prints , arXiv:1910.08980 (2019), [arXiv:1910.08980 \[quant-ph\]](#) .
- [11] G. G. Guerreschi and A. Y. Matsuura, QAOA for Max-Cut requires hundreds of qubits for quantum speed-up, *Scientific Reports* **9**, 6903 (2019).
- [12] R. Shaydulin, I. Safro, and J. Larson, Multistart methods for quantum approximate optimization, in *2019 IEEE High Performance Extreme Computing Conference (HPEC)* (2019) pp. 1–8.
- [13] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, *Nature Communications* **9**, 4812 (2018), [arXiv:1803.11173 \[quant-ph\]](#) .
- [14] Z. Holmes, K. Sharma, M. Cerezo, and P. J. Coles, Connecting ansatz expressibility to gradient magnitudes and barren plateaus, arXiv e-prints , arXiv:2101.02138 (2021), [arXiv:2101.02138 \[quant-ph\]](#) .
- [15] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, Cost function dependent barren plateaus in shallow parametrized quantum circuits, *Nature Communications* **12**, 1791 (2021), [arXiv:2001.00550 \[quant-ph\]](#) .
- [16] F. G. S. L. Brandao, M. Broughton, E. Farhi, S. Gutmann, and H. Neven, For Fixed Control Parameters the Quantum Approximate Optimization Algorithm’s Objective Function Value Concentrates for Typical Instances, arXiv e-prints , arXiv:1812.04170 (2018), [arXiv:1812.04170 \[quant-ph\]](#) .
- [17] D. J. Egger, J. Marecek, and S. Woerner, Warm-starting quantum optimization, arXiv e-prints , arXiv:2009.10095 (2020), [arXiv:2009.10095 \[quant-ph\]](#) .
- [18] M. Alam, A. Ash-Saki, and S. Ghosh, Accelerating Quantum Approximate Optimization Algorithm using Machine Learning, arXiv e-prints , arXiv:2002.01089 (2020), [arXiv:2002.01089 \[cs.ET\]](#) .
- [19] S. Khairy, R. Shaydulin, L. Cincio, Y. Alexeev, and P. Balaprakash, Learning to Optimize Variational Quantum Circuits to Solve Combinatorial Problems, arXiv e-prints , arXiv:1911.11071 (2019), [arXiv:1911.11071 \[cs.LG\]](#) .
- [20] T. Albash and D. A. Lidar, Adiabatic quantum computation, *Rev. Mod. Phys.* **90**, 015002 (2018).
- [21] D. Liang, L. Li, and S. Leichenauer, Investigating quantum approximate optimization algorithms under bang-bang protocols, *Phys. Rev. Research* **2**, 033402 (2020).
- [22] M. Heyl, P. Hauke, and P. Zoller, Quantum localization bounds Trotter errors in digital quantum simulation, *Science Advances* **5**, [10.1126/sciadv.aau8342](#) (2019).
- [23] M. X. Goemans and D. P. Williamson, Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming, *J. ACM* **42**, 1115 (1995).
- [24] J. Wurtz and P. J. Love, Bounds on MAXCUT QAOA performance for $p > 1$, arXiv e-prints , arXiv:2010.11209 (2020), [arXiv:2010.11209 \[quant-ph\]](#) .
- [25] C. G. BROYDEN, The Convergence of a Class of Double-rank Minimization Algorithms 1. General Considerations, *IMA Journal of Applied Mathematics* **6**, 76 (1970).
- [26] R. Fletcher, A new approach to variable metric algorithms, *The Computer Journal* **13**, 317 (1970).
- [27] D. Goldfarb, A family of variable-metric methods derived by variational means, *Mathematics of Computation* **24**, 23 (1970).
- [28] D. F. Shanno, Conditioning of quasi-Newton methods for function minimization, *Mathematics of Computation* **24**, 647 (1970).
- [29] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors, SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python, *Nature Methods* **17**, 261 (2020).
- [30] M. Rosenblatt, Remarks on some nonparametric estimates of a density function, *Ann. Math. Statist.* **27**, 832 (1956).
- [31] E. Parzen, On estimation of a probability density function and mode, *Ann. Math. Statist.* **33**, 1065 (1962).
- [32] T. Kadowaki and H. Nishimori, Quantum annealing in the transverse Ising model, *Phys. Rev. E* **58**, 5355 (1998).
- [33] J. Brooke, D. Bitko, T. F. Rosenbaum, and G. Aeppli, Quantum Annealing of a Disordered Magnet, *Science* **284**, 779 (1999).
- [34] E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda, A Quantum Adiabatic Evolution Algorithm Applied to Random Instances of an NP-Complete Problem, *Science* **292**, 472 (2001), [arXiv:quant-ph/0104129 \[quant-ph\]](#) .

- [35] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, Quantum Computation by Adiabatic Evolution, arXiv e-prints , quant-ph/0001106 (2000), [arXiv:quant-ph/0001106 \[quant-ph\]](#) .
- [36] D. Aharonov, W. van Dam, J. Kempe, Z. Landau, S. Lloyd, and O. Regev, Adiabatic Quantum Computation Is Equivalent to Standard Quantum Computation, *SIAM Review* **50**, 755 (2008).
- [37] A. Smith, M. S. Kim, F. Pollmann, and J. Knolle, Simulating quantum many-body dynamics on a current digital quantum computer, *npj Quantum Information* **5**, 106 (2019), [arXiv:1906.06343 \[quant-ph\]](#) .
- [38] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, *Annals of Physics* **326**, 96 (2011), [arXiv:1008.3477 \[cond-mat.str-el\]](#) .
- [39] G. Carleo and M. Troyer, Solving the quantum many-body problem with artificial neural networks, *Science* **355**, 602 (2017), [arXiv:1606.02318 \[cond-mat.dis-nn\]](#) .
- [40] M. Medvidovic and G. Carleo, Classical variational simulation of the Quantum Approximate Optimization Algorithm, arXiv e-prints , [arXiv:2009.01760 \(2020\)](#), [arXiv:2009.01760 \[quant-ph\]](#) .
- [41] A. Abbas, D. Sutter, C. Zoufal, A. Lucchi, A. Figalli, and S. Woerner, The power of quantum neural networks, arXiv e-prints , [arXiv:2011.00027 \(2020\)](#), [arXiv:2011.00027 \[quant-ph\]](#) .
- [42] D. Guéry-Odelin, A. Ruschhaupt, A. Kiely, E. Torrontegui, S. Martínez-Garaot, and J. G. Muga, Shortcuts to adiabaticity: Concepts, methods, and applications, *Rev. Mod. Phys.* **91**, 045001 (2019).
- [43] D. Sels and A. Polkovnikov, Minimizing irreversible losses in quantum systems by local counterdiabatic driving, *Proceedings of the National Academy of Sciences* **114**, E3909 (2017).
- [44] P. W. Claeys, M. Pandey, D. Sels, and A. Polkovnikov, Floquet-engineering counterdiabatic protocols in quantum many-body systems, *Phys. Rev. Lett.* **123**, 090602 (2019).
- [45] Z.-C. Yang, A. Rahmani, A. Shabani, H. Neven, and C. Chamon, Optimizing variational quantum algorithms using Pontryagin’s minimum principle, *Phys. Rev. X* **7**, 021027 (2017).
- [46] S. H. Sack, Trotterized quantum annealing initialization of the QAOA, <https://github.com/shsack/TQA-init.-for-QAOA> (2021).
- [47] A. Hagberg, P. Swart, and D. S Chult, *Exploring network structure, dynamics, and function using NetworkX*, Tech. Rep. (Los Alamos National Lab.(LANL), Los Alamos, NM (United States), 2008).
- [48] D. W. Berry, G. Ahokas, R. Cleve, and B. C. Sanders, Efficient Quantum Algorithms for Simulating Sparse Hamiltonians, *Communications in Mathematical Physics* **270**, 359 (2007), [arXiv:quant-ph/0508139 \[quant-ph\]](#) .