

Learning parameter curves in feedback-based quantum algorithms

Vicente Peña Pérez ^{1,2,*} Matthew D. Grace ² Christian Arenz ^{1,†} and Alicia B. Magann ^{2,‡}

¹*School of Electrical, Computer, and Energy Engineering, Arizona State University, Tempe, Arizona 85287, USA*

²*Quantum Algorithms and Applications Collaboratory, Sandia National Laboratories, Albuquerque, New Mexico 87185, USA*



(Received 16 January 2026; accepted 4 September 2026; published 29 September 2026)

Feedback-based quantum algorithms (FQAs) operate by iteratively growing a quantum circuit to optimize a given task. At each step, feedback from qubit measurements is used to inform the next quantum circuit update. In practice, the sampling cost associated with these measurements can be significant. Here, we ask whether FQA parameter sequences can be predicted using classical machine learning, obviating the need for qubit measurements altogether. To this end, we train a teacher-student model to map problem instances to associated FQA parameter curves in a single classical inference step. We first study weighted MaxCut problem instances, where numerical experiments show that this machine learning model can accurately predict FQA parameter curves across a range of problem sizes, including sizes not seen during model training. We then extend the analysis to transverse-field Ising models on graphs and analyze how this learning strategy performs and generalizes beyond the original MaxCut setting. To evaluate performance, we compare the predicted parameter curves in simulation against FQA reference curves and linear quantum annealing schedules. We observe similar results to the former and performance improvements over the latter. These results suggest that machine learning can offer a heuristic, practical path to reducing sampling costs and resource overheads in quantum algorithms.

DOI: [10.1103/pgbz-zm8h](https://doi.org/10.1103/pgbz-zm8h)

I. INTRODUCTION

There is substantial interest in using quantum computers to approximate ground states of Hamiltonians. This is a challenging computational task [1,2] with diverse applications, including in combinatorial optimization [3]. Early work in the latter domain centered on quantum annealing (QA) [4–7], a conventionally continuous-time approach that aims to prepare the ground state of an Ising Hamiltonian using an adiabatic evolution, parametrized according to an annealing schedule. Later, variational quantum algorithms (VQAs) emerged as another approach for ground-state preparation geared toward implementations on digital, gate-model quantum computers [8,9]. These algorithms operate by classically optimizing over a set of free parameters in a quantum circuit, whose purpose is to prepare an approximation of the target ground state. One notable example of a VQA is the quantum approximate optimization algorithm (QAOA) [10], developed in analogy to QA for finding approximate solutions to combinatorial optimization problems.

Feedback-based quantum algorithms, such as the feedback-based algorithm for quantum optimization

(FALQON) [11,12], were subsequently developed in analogy to QAOA. These algorithms operate by assigning values to free parameters in quantum circuits in a sequential manner. At each step in the sequence, a feedback law is utilized to set the value of a given parameter according to the outcome of qubit measurements performed at the previous step. Then, additional qubit measurements are performed to inform the next step. The feedback law is derived from quantum Lyapunov control principles [13,14] to ensure that a cost function decreases monotonically as a function of step. In practice, however, the sampling cost associated with evaluating the feedback law can be significant, as each step requires many circuit repetitions to perform the needed qubit measurements. This sampling overhead can become a significant bottleneck toward quantum device implementations at scale. Here, we address this challenge by leveraging classical machine learning (ML) to predict FALQON parameters without the need for measurement data.

The sequence of FALQON parameters can be collected into a *parameter curve*, which is analogous to parameter sequences in QAOA and continuous annealing schedules in QA. FALQON parameter curves tend to display a typical shape, characterized by an initial peak region defined by large, rapid variations in the FQA parameters, followed by a subsequent tail region. Figure 1 illustrates this typical shape for FALQON parameter curves associated with the MaxCut problem. Across different MaxCut problem instances, the curves cluster tightly around the common profile. This phenomenon was observed but not analyzed in Ref. [11], and it suggests that FALQON parameter curves may be predictable from relatively simple problem features, without having to explicitly run the full feedback procedure for each instance.

*Contact author: vpenaper@asu.edu

†Contact author: carenz1@asu.edu

‡Contact author: abmagan@sandia.gov




FIG. 1. FALQON parameter curves found for solving MaxCut on 19 nonisomorphic, unweighted 3-regular graphs with ten nodes, where each curve corresponds to the FALQON solution for one of the 19 graphs. The consistent clustering of these curves around a common profile motivates our work exploring whether machine learning could be employed to predict them.

Motivated by this observation, in this work we develop and test an ML model that replaces the per-layer qubit measurements with a one-step classical prediction of the full parameter curve. We develop this ML model in the context of predicting FQA parameter curves associated with solving an Ising model formulation of the weighted MaxCut problem and analyze its performance. We then explore applications to quantum models, here considering transverse-field Ising models (TFIMs) defined on the same class of graphs considered for MaxCut. While our results appear encouraging, we emphasize that our ML approach is heuristic: What problem classes admit FQA parameter curves compatible with prediction using classical ML, and when ML models trained on one family of Hamiltonians can be productively transferred to others, remain open questions.

This direction is aligned with prior efforts that use ML to predict QAOA parameters, primarily to reduce the burden of classical optimization [15–19]. Ultimately, the ML-predicted curves we obtain could be used as stand-alone solutions that directly replace FALQON, QA, or QAOA solutions, or as warm starts for subsequent refinement through classical optimization or hybrid FALQON or QAOA strategies [12].

The remainder of this article is organized as follows. Section II reviews QA, QAOA, and FALQON, contrasting feedback-generated FALQON parameter curves with those of QAOA and QA. Section III introduces our teacher-student approach for predicting the full FALQON parameter curves, detailing the high-level architecture and training procedure. Section IV reports a variety of numerical results that explore the performance of the ML predictions. Section V summarizes and outlines future directions.

II. BACKGROUND

In this section, we provide relevant technical background on QA and QAOA, followed by an overview of FALQON.

A. Quantum annealing

Quantum annealing is a strategy for preparing or approximating the ground state of a target problem Hamiltonian,

H_p [4–7]. It operates by first initializing a set of qubits in the ground state of another, simpler Hamiltonian, denoted here by H_d . Then, the qubits are evolved over time under a Hamiltonian,

$$H(t) = A(t)H_d + B(t)H_p, \quad t \in [0, T], \quad (1)$$

that smoothly interpolates between the simpler Hamiltonian and the problem Hamiltonian of interest, according to a pre-defined annealing schedule captured by the time-dependent scalar functions $A(t)$ and $B(t)$, typically selected so that $A(0) = B(T) = 1$ and $A(T) = B(0) = 0$. The performance of QA depends strongly on the choice of QA schedule, and a commonly used heuristic choice is a linear schedule, given by

$$A(t) = 1 - \frac{t}{T}, \quad B(t) = \frac{t}{T}, \quad (2)$$

which interpolates from H_d at $t = 0$ to the problem Hamiltonian H_p at $t = T$. While Eq. (2) is a simple and commonly used schedule, it is not generally optimal. If some coarse-grained information about the instantaneous spectral gap is available, then the schedule can be modified to slow down near small-gap regions and speed up where the gap is larger. Such local-adiabatic strategies can improve over a purely linear interpolation, although they generally require additional spectral information about the problem Hamiltonian [20]. If the evolution proceeds slowly enough, the system remains close to the instantaneous ground state and transitions toward the target ground state according to the quantum adiabatic theorem [21]. Quantum annealing is conventionally an analog procedure, carried out continuously in time. In practice, the slowly varying schedules required to maintain adiabaticity can demand evolution times that are prohibitively long [22,23].

It is also possible to digitize QA for implementation on gate-based quantum computers [24,25]. The latter has been explored in contexts such as adiabatic quantum state preparation, which has been considered as an approach for preparing initial states that are then fed into quantum algorithms such as quantum phase estimation, or algorithms for Hamiltonian simulation. Digitization can be accomplished through Hamiltonian simulation [26], e.g., by discretizing the interval $[0, T]$ into ℓ steps of size $\Delta t = T/\ell$ and applying a first-order Trotter formula [27,28] to yield the product

$$U_{QA}(T) \approx \prod_{j=1}^{\ell} e^{-i\Delta t B(t_j)H_p} e^{-i\Delta t A(t_j)H_d}, \quad (3)$$

where $t_j = j \Delta t$ and the error of the approximation can be controlled via the choice of Δt according to Trotter error bounds [27,28]. Equation (3) describes a unitary sequence that can be used for a digitized implementation. In practice, this would be further compiled down to a gate sequence. Due to the long times that are often required for QA, such digital formulations can involve circuit depths that are substantial.

B. Quantum approximate optimization algorithm

More recently, VQAs have emerged as another paradigm for ground-state preparation on digital, gate-model quantum computers [8,9]. The idea is to execute short parametrized quantum circuits on a quantum computer, with the aim that these circuits prepare the ground state of a target Hamiltonian,

or close to it. This goal is quantified using a cost function, which is typically chosen to be the target Hamiltonian expectation value, as this is minimized by the ground state. Then, a classical computer is used to iteratively search for the values of the quantum circuit parameters that minimize the cost function. Depending on the problem at hand, the resulting classical optimization landscape can be highly nonconvex and can exhibit many local minima, making convergence challenging when not enough variational parameters are used [29]. Moreover, variational training can suffer from barren plateaus, where gradients vanish exponentially with system size, rendering gradient estimation itself resource intensive [30,31].

One relevant example of a VQA is the QAOA [10,32], developed for finding approximate solutions to combinatorial optimization problems. The roots of QAOA can be traced back to QA. In particular, QAOA utilizes a parametrized quantum circuit ansatz whose structure matches that of a digitized simulation of QA in Eq. (3). In particular, a depth- p QAOA circuit prepares the quantum state

$$|\psi_p(\boldsymbol{\gamma}, \boldsymbol{\beta})\rangle = \left(\prod_{j=1}^p e^{-i\beta_j H_d} e^{-i\gamma_j H_p} \right) |\psi_0\rangle, \quad (4)$$

where $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_p)$ and $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)$ are free parameters and $|\psi_0\rangle$ is the initial state typically taken to be an eigenstate of H_d . That is, rather than using a fixed QA schedule, QAOA allows the digitized evolutions to be freely parametrized. The parameter values are iteratively adjusted by a classical optimizer to minimize a cost function, $\langle \psi_p(\boldsymbol{\gamma}, \boldsymbol{\beta}) | H_p | \psi_p(\boldsymbol{\gamma}, \boldsymbol{\beta}) \rangle$, given by the expectation value of the problem Hamiltonian. This cost function is minimized by the ground state. The properties of optimal parameter curves in QAOA and QA have been analyzed previously in Refs. [33–35].

The classical optimization loop in QAOA allows for going beyond slow, predefined QA schedules; instead, in QAOA, the free parameters are iteratively optimized in an effort to achieve faster convergence than adiabatic QA or straightforward digitizations of QA, using quantum circuits with a relatively shallow depth. Nonetheless, as the size and complexity of the combinatorial problems of interest increase, the classical optimization cost in QAOA associated with searching through the associated parameter space can become a bottleneck. A variety of proposals have aimed to address this situation, e.g., through the use of bootstrapping solutions [15,17,36], warm starting [37,38] the classical optimization problem, or utilizing adaptively constructed circuits [39–44].

C. Feedback-based algorithm for quantum optimization

Feedback-based quantum algorithms constitute another approach for ground-state preparation [11,12,45–52]. Here, we consider the FALQON, which was developed in close analogy to QAOA. In the following, we provide details on FALQON, beginning by motivating the relevant background in continuous-time quantum Lyapunov control theory [13,53–60].

Consider a quantum system whose dynamics (in units of $\hbar = 1$) are described by the Schrödinger

equation

$$i \frac{d}{dt} |\psi(t)\rangle = [H_p + \beta(t) H_d] |\psi(t)\rangle, \quad (5)$$

where $|\psi(t)\rangle$ denotes the state of the system at time t , H_p denotes here the portion of the Hamiltonian that is time independent, and H_d denotes the so-called driver Hamiltonian that couples a time-dependent control function $\beta(t)$ to the system.

Quantum Lyapunov control is a technique that can be applied to a variety of quantum control problems. Here, we consider the problem of designing a control function, $\beta(t)$, for controlling, and specifically minimizing, the expectation value of an observable, taken here to be H_p . For this application, quantum Lyapunov control operates by seeking $\beta(t)$ to satisfy the following derivative condition:

$$\frac{d}{dt} \langle H_p \rangle_t \leq 0, \quad (6)$$

and in so doing, ensure that $\langle H_p \rangle_t$ decreases monotonically over time. Here, we have used the abbreviated notation $\langle \cdot \rangle_t \equiv \langle \psi(t) | \cdot | \psi(t) \rangle$ for the expectation value. Evaluating the left side of Eq. (6) leads to

$$\frac{d}{dt} \langle H_p \rangle_t = i \langle [H_d, H_p] \rangle_t \beta(t) = A(t) \beta(t), \quad (7)$$

with

$$A(t) \equiv i \langle [H_d, H_p] \rangle_t. \quad (8)$$

There are many choices of $\beta(t)$ that satisfy Eq. (6) [14]. A common choice is to utilize a quantum Lyapunov control law of the form

$$\beta(t) = -A(t) \quad (9)$$

to ensure that the derivative condition in Eq. (6) is satisfied at all times t . This monotonic descent alone does not guarantee eventual convergence to the ground state of H_p , since the equality in Eq. (6) is obtained whenever $A(t) = 0$ (see Ref. [12]). This means that in practice, quantum Lyapunov control is often utilized as a heuristic control approach. Numerical evidence supports that it can often achieve good convergence in practice [11,12], with performance depending on the choice of initial state $|\psi(0)\rangle$, specification of Hamiltonians H_p and H_d , and the functional form and parameters in the control law in Eq. (9).

We now describe how digitizing this continuous-time quantum Lyapunov control approach leads to FALQON as a quantum circuit-level algorithm. This mirrors the discussion of digitizing QA in Sec. II A. First, we discretize time into steps of length Δt and take $\beta(t)$ to be constant over each step, and second, we Trotterize the time evolution over a single time step, giving

$$|\psi_j\rangle = e^{-i \Delta t H_d \beta_j} e^{-i \Delta t H_p} |\psi_{j-1}\rangle, \quad (10)$$

where β_j is the sampled control at $t_j = j \Delta t$, $|\psi_j\rangle$ denotes the associated state, and $j = 1, 2, \dots$ labels the step or layer. The choice of Δt has a large impact on the algorithm execution. Selecting this parameter to be too large can lead to violations of the derivative condition in Eq. (6), while selecting it to be very small can lead to slow convergence as a function of layer. In practice, Δt can be treated as a hyperparameter; it is desirable to tune it to be as large as possible while still satisfying Eq. (6).




FIG. 2. Schematic of the steps for implementing FALQON. A weighted graph G (left) defines the problem Hamiltonian H_p . Each quantum circuit layer j applies $U_p = e^{-i\Delta t H_p}$ followed by $U_d(\beta_j) = e^{-i\Delta t \beta_j H_d}$. After layer j , the expectation value $A_j \equiv \langle i[H_d, H_p] \rangle_j$ is estimated from qubit measurements and used to set the next parameter via the feedback law $\beta_{j+1} = -A_j$. Repeating for ℓ layers generates the parameter curve $\{\beta_j\}_{j=1}^\ell$, presented at right. At the end of the protocol, one may either estimate the final expectation value $\langle H_p \rangle_\ell$ or sample the terminal state to obtain a candidate MaxCut bit string solution. In the present work, we examine the prospect of replacing the steps of FALQON, and associated sampling costs, with an ML model trained to predict the full FALQON parameter curve. This figure is adapted from Ref. [11].

In many cases, the digitized form in Eq. (10) is readily compatible with a quantum circuit-level decomposition, and it represents one step of FALQON. In other cases, e.g., where H_p and H_d contain multiple noncommuting terms, we can perform a subsequent Trotterization of each exponential in Eq. (10) to obtain a form that is then decomposable into a quantum circuit [46].

At each step of FALQON, β_j is determined using the quantum Lyapunov control law from Eq. (9). In FALQON, this control law is evaluated using feedback obtained from the previous step, with the typical choice of feedback law being

$$\beta_j = -A_{j-1}, \quad (11)$$

with $A_{j-1} = i\langle \psi_{j-1} | [H_d, H_p] | \psi_{j-1} \rangle$. Based on these principles, the steps for executing FALQON are depicted in Fig. 2 and described in the associated figure caption. Crucially, at each step the quantity A_{j-1} is obtained by estimating the expectation value of $i[H_d, H_p]$. This observable can, in general, be expanded into a linear combination of Pauli operators. These Pauli operators can be grouped into mutually commuting subsets, and the expectation values of each subset can be estimated separately. Then, the output results can be combined classically to evaluate A_{j-1} [46]. Thus, the feedback law is experimentally accessible, but it can require many repeated circuit executions at each layer to estimate the required expectation value. We refer to the FALQON sequence $\{\beta_j\}_{j=1}^\ell$ produced by the feedback law in Eq. (11) as a parameter curve. This is analogous to the (γ, β) parameter sequences in QAOA and the continuous annealing schedules defined by $A(t), B(t)$ in QA.

To generate a FALQON parameter curve, each expectation value A_{j-1} is estimated from qubit measurements, performed on repeated executions of the FALQON circuit up to layer

$j - 1$. In practice, FALQON can require deep circuits with many layers, increasing the measurement cost. In the following, we explore the prospect of training a machine learning model to predict these curves classically, thereby removing the need for qubit measurements.

III. MACHINE LEARNING PREDICTION OF FALQON PARAMETER CURVES

This section presents the ML model we use to predict FALQON parameter curves for two classes of problems: MaxCut on weighted graphs and transverse-field Ising models on graphs. The former is a widely studied combinatorial optimization problem that has been extensively studied in both classical [61] and quantum settings, and has been utilized as a standard benchmark for QAOA and FALQON [10,11]. The latter is also a widely utilized benchmark problem and expands the scope of our work to “quantum” Hamiltonians, i.e., whose ground states can possess entanglement. We first overview the application of FALQON toward each of these problem classes, and then go on to motivate and describe the high-level architecture of the ML model, with further details captured in Fig. 3. We then discuss the training procedure used in this work. Additional architectural details, a list of hyperparameters, and alternative pooling and readout choices are provided in the Appendix.

A. FALQON for MaxCut

Given an undirected graph $G = (V, E)$ with vertex set V of size $|V| = n$ and edge set E , a cut is a partition of V into two disjoint subsets. The cut size is the number of edges that cross the partition, and the (unweighted) MaxCut problem asks for




FIG. 3. Teacher-student ML framework for predicting FALQON parameter curves for MaxCut. (a) Teacher: The input graph G and five scalar features s are processed by a graph neural network (GNN); pooled graph features feed a small conditioning module that outputs a full parameter curve $\hat{\beta}_T$. (b) Student: A lighter graph network produces an embedding g and an auxiliary estimate $\hat{s}$ (supervised only during training); a small readout maps $[g, \hat{s}]$ to the predicted curve $\hat{\beta}_S$. At test time, only $\hat{\beta}_S$ is used. Further details are given in the Appendix.

the partition that maximizes this number. In the weighted variant we consider here, each edge is assigned a positive weight w_{ij} , and the objective becomes to maximize the total weight of edges crossing the cut. Figure 2 illustrates this premise in the case of a 3-node graph, where the dashed line partitioning the graph on the right-hand side of the figure represents the maximum cut.

The weighted MaxCut problem can be cast as a ground-state preparation task by defining the problem Hamiltonian

$$H_p^{\text{MC}} = - \sum_{(i,j) \in E} w_{ij} \frac{1 - Z_i Z_j}{2}, \quad (12)$$

whose ground state encodes the maximum cut, where Z_i denotes the Pauli-Z operator on qubit i and w_{ij} are the edge weights. The driver Hamiltonian considered in the present work is given by

$$H_d = \sum_{i=1}^n X_i, \quad (13)$$

where X_i denotes the Pauli-X operator on qubit i . Qubits are initialized in the ground state of H_d , given by

$$|\psi_0\rangle = |-\rangle^{\otimes n}, \quad (14)$$

where $|-\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$.

We evaluate performance using two figures of merit. The first is the approximation ratio, defined at layer j as

$$r_{A,j} = \frac{E_j}{E_{\min}}, \quad (15)$$

where $E_j = \langle \psi_j | H_p^{\text{MC}} | \psi_j \rangle$ and $E_{\min}$ denotes the minimum eigenvalue of the problem Hamiltonian. In our numerical simulations, $E_{\min}$ is computed exactly and used only to evaluate performance. It is not assumed to be known during execution of FALQON or during inference with the trained ML model.

In a setting where exact normalization is unavailable, one could instead report the energy directly, compare against the best sampled cut value, or estimate approximation quality relative to a classical benchmark when available.

The second figure of merit we consider is the squared overlap with the ground-state subspace, defined at the j th layer as

$$\phi_j = \sum_i |\langle g_i | \psi_j \rangle|^2, \quad (16)$$

where $\{|g_i\rangle\}$ denotes the set of degenerate ground states of H_p^{MC} . Equivalently, ϕ_j is the fidelity with the ground-state subspace of H_p^{MC} . For MaxCut, this quantity can also be interpreted as the probability that a computational-basis measurement samples an optimal cut. Thus, $r_{A,j}$ captures the average cut quality, while ϕ_j captures the probability weight assigned to globally optimal solutions.

B. FALQON for the transverse-field Ising model

The TFIM describes interacting spins acted on by a transverse field. Here, we consider spins with pairwise interactions, where the coupling graphs are sampled from the same distributions as the graphs considered for MaxCut. The TFIM Hamiltonian is given by

$$H_p^{\text{TFIM}} = H_{ZZ} + hH_X, \quad (17)$$

with

$$H_{ZZ} = \sum_{(i,j) \in E} w_{ij} Z_i Z_j, \quad H_X = \sum_{i=1}^n X_i. \quad (18)$$

Here, w_{ij} denotes the edge weight and h is the transverse-field strength. In the numerical experiments below, we set $h = 1$. We sample the edge weights independently from the uniform distribution on $[0, 2]$. Also, the unweighted baseline

is obtained by setting all edge weights equal to 1. For studies of the TFIM, we consider the same driver Hamiltonian and initial state given in Eqs. (13) and (14).

Since the transverse-field term in H_p^{TFIM} is proportional to the driver Hamiltonian, one layer of FALQON is given by

$$|\psi_j\rangle = e^{-i\Delta t (\beta_j + h) H_X} e^{-i\Delta t H_{ZZ}} |\psi_{j-1}\rangle, \quad (19)$$

with the feedback law generated from the commutator

$$\beta_j = -A_{j-1}, \quad A_{j-1} = i\langle\psi_{j-1}|[H_X, H_p^{\text{TFIM}}]|\psi_{j-1}\rangle. \quad (20)$$

Because $[H_X, hH_X] = 0$, the feedback is determined by the noncommuting Ising component H_{ZZ} , while the performance is evaluated with respect to the full TFIM Hamiltonian in Eq. (17).

We evaluate TFIM performance using the same figures of merit $r_{A,j}$ and ϕ_j introduced in Eqs. (15) and (16), but using H_p^{TFIM} . Unlike H_p^{MC} , whose spectrum is nonpositive by construction, the TFIM Hamiltonian H_p^{TFIM} has an indefinite spectrum. Thus, for an arbitrary state, the normalized quantity $r_{A,j}^{\text{TFIM}}$ could in principle be negative. The feedback law is designed to decrease $\langle H_p^{\text{TFIM}} \rangle_j$ with layer, and our simulations remain in the negative-energy regime. Therefore, along the exact FALQON trajectories and sufficiently accurate approximations to them, $r_{A,j}$ remains between 0 and 1. For other choices of initial state, this interpretation need not hold.

C. Machine learning model

Our ML model is based on a teacher-student framework: A higher-capacity teacher produces high-quality reference curves, and a lighter student is subsequently trained to approximate those curves [62]. The teacher is trained to fit FALQON curves directly, and the student is trained to mimic the teacher (a distillation loss on the whole curve), with an auxiliary loss on a set of predicted scalars $\hat{s}$. Details of the loss functions are given in the Appendix. At a high level, the student and teacher models share the following flow. A weighted graph as input is first fed into a GNN, where each node aggregates information from its neighbors and connected edges in a message-passing fashion [63]. We use a GNN here because it is a natural fit for our graph Hamiltonian problems; i.e., the aggregation it performs depends only on adjacency and is permutation invariant to vertex relabelings. This is relevant because it prevents spurious label memorization given that the MaxCut and TFIM objectives are invariant to vertex relabeling. This is followed by multiple convolutional layers, i.e., given by graph isomorphism network with edge features (GINEConv) and the mask label prediction (TransformerConv) layers in the teacher and the student, respectively [64,65], and the outputs are pooled to capture both local statistics and global structure. This is then sent through a small readout network that outputs the FALQON parameter curve $\hat{\beta}$, where the hat is used to denote that it is a model prediction. The teacher and student share this overall pattern; the teacher uses a more expressive, conditioned readout, whereas the student is streamlined. In our numerical analyses, we have found that simpler models, e.g., using dense networks only, ignoring graph structure, or only matching values, produce worse predictions.

We consider weighted, 3-regular graphs with $n \in \{4, 6, 8, 10, 12\}$ and with edge weights sampled uniformly

from $[0, 2]$. In total, we generate 2240 instances and split them into 60% training (of which 80% are used for parameter updates and 20% for validation) and 40% testing. For each instance, we compute the FALQON parameter curve exactly in numerical simulation with time step $\Delta t = 0.01$. We train with batch size 1 so that each optimization step processes a full parameter curve, and we use Adam [66]. Training runs for 100 epochs with learning rate 10^{-3} , with hyperparameters tuned using Optuna, an open source optimization toolkit that automates search over model settings [67]. The model outputs parameter curves of length $\ell = 1001$ to match the FALQON depth used in our numerical experiments. The forward path follows Fig. 3: three message-passing blocks (1–3), multi-pooling heads (A–E), and a two-stage readout (I and II). The teacher uses a conditioned (dynamic) readout whose weights are produced by a small hypernetwork fed by the scalars s described in the Appendix; the student uses a lightweight readout and additionally predicts $\hat{s}$ as an auxiliary target. After training, we use only the student to predict parameter curves; the teacher is not needed at test time.

The resulting ML workflow used in this work is as follows. For the MaxCut study, we generate a training set of weighted graph instances and compute reference FALQON parameter curves by classical simulation. After training the teacher-student model, a new weighted graph is passed through the trained student model, which outputs a predicted parameter curve in a single classical inference step. This curve can then be replayed as a fixed schedule on the corresponding weighted Hamiltonian for validation, if desired. The model architecture and workflow were designed for the MaxCut setting. The TFIM study is included as an initial exploration of whether the same learning strategy can be useful beyond MaxCut. For TFIM, the same teacher-student architecture and training procedure are used, but we train a new model on TFIM-generated FALQON curves. Thus, the TFIM results test whether the same approach and architecture can capture parameter curves for a different graph-defined Hamiltonian.

IV. RESULTS

In this section, we present numerical results exploring how the ML model described in Sec. III performs at predicting FALQON parameter curves. We organize the section around four questions: (1) How faithfully does the student model reproduce FALQON results layer by layer for weighted MaxCut? (2) How well does the model extrapolate to MaxCut problem sizes not seen during training? (3) For MaxCut, how do the learned curves perform when compared with a standard digitized linear-annealing schedule? (4) Does the same learning strategy extend to a graph-defined transverse-field Ising model? Questions (1)–(4) are analyzed in Figs. 4–6, and 7 and 8, respectively. Figure 4 compares, layer by layer, the reference FALQON parameters β_j with the predictions produced by our student ML model on weighted 3-regular graphs with $n \in \{8, 10, 12\}$. These results are plotted in blue. For a given weighted graph instance, the reference FALQON parameters are computed from an exact simulation of running FALQON for solving MaxCut on that weighted graph instance. The ML model predictions are obtained using the same weighted graph as input to the student model to generate




FIG. 4. MaxCut performance is plotted vs layer for the ML model predictions (blue) as well as the unweighted baseline (red) for different figures of merit and different problem sizes. Solid curves show the mean across problem instances, shaded bands indicate the associated standard deviation, and dashed curves indicate the maximum error across instances. The three columns of panels correspond to graph sizes $n \in \{8, 10, 12\}$. The top row shows the parameter error $|\Delta\beta_j|$ on a logarithmic scale, the middle row shows the approximation-ratio error $|\Delta r_{A,j}|$, and the bottom row shows the ground-state-subspace-overlap error $|\Delta\phi_j|$. For a given weighted graph instance, these errors are computed with respect to the true FALQON curve and its performance for solving MaxCut on that instance. The ML model predictions are obtained using the same weighted graph as input, and the unweighted baseline corresponds to using the FALQON parameter curve computed for the same topology with all weights set equal to one. The logarithmic scale resolves the tail-region behavior of the parameter error, where the ML predictions remain close to the weighted FALQON reference and generally improve over the unweighted baseline.

an ML-predicted parameter curve. This latter ML curve is then used to parametrize an exact simulation of a FALQON circuit for solving MaxCut on this weighted graph instance, and its performance is evaluated in terms of the absolute deviations between the predicted and reference values, according to

$$|\Delta\beta_j| \equiv |\hat{\beta}_j - \beta_j|, \quad (21)$$

$$|\Delta r_{A,j}| \equiv |\hat{r}_{A,j} - r_{A,j}|, \quad (22)$$

and

$$|\Delta\phi_j| \equiv |\hat{\phi}_j - \phi_j|, \quad (23)$$

at each layer j , where hats denote ML student model predictions. We first observe that across $n \in \{8, 10, 12\}$, the ML predictions closely track FALQON for β_j . This matters because any layerwise deviations $\Delta\beta_j$ can introduce errors that then accumulate as a function of layer. Here, we do not observe this issue, and the good agreement between the ML predictions and the FALQON reference curves extends to r_A and ϕ .

The results in red in Fig. 4, meanwhile, show the performance of an “unweighted-baseline” model, included here to benchmark the ML model results. For a MaxCut problem instance defined by a given weighted graph, this unweighted baseline is computed as follows. First, we run FALQON for this graph, but with all edge weights set equal to 1. This outputs a parameter curve associated with solving the unweighted version of the MaxCut problem. The performance of this unweighted-baseline curve is then evaluated in simulation against the original weighted graph. For edge weights near 1, this baseline is expected to be strong because the weighted Hamiltonian is close to its unweighted counterpart. As the weight distribution broadens, however, the baseline becomes a more stringent test of whether the ML model learns weight-dependent effects. By utilizing the unweighted solution as a baseline to compare against, we can therefore check our ML model’s ability to learn these edge-weight effects, in addition to the effects of the underlying graph topology. We see in Fig. 4 that the ML model results are consistently better than the results of the unweighted baseline, confirming that the ML




FIG. 5. MaxCut performance is plotted vs layer for the ML model predictions (blue) as well as the unweighted baseline (red) for different figures of merit and for problem sizes $n \in \{14, 16, 18, 20\}$ not seen during training. The plot conventions match those in Fig. 4. The logarithmic scale in the top row makes the long-layer behavior of $|\Delta\beta_j|$ visible after the initial peak region. The ML model continues to reproduce key features of the FALQON parameter curves outside the training range, although the performance gradually deteriorates as the problem size moves farther from the training distribution.

model learns weight-dependent parameter curves. We observe that performance is strongest at $n = 12$. We attribute this to greater training diversity: There are 85 nonisomorphic graphs for $n = 12$ versus 5 and 19 for $n = 8$ and $n = 10$, respectively; this richer variety appears to regularize the predictor and improve generalization within the training range.

The ability of the ML model to generalize beyond the training sizes is explored next, with results presented in Fig. 5. For unseen graphs with $n \in \{14, 16, 18, 20\}$, the absolute deviations $|\Delta\beta_j|$, $|\Delta r_{A,j}|$, and $|\Delta\phi_j|$ remain small on average, with modest variability across layers. Notably, $|\Delta r_{A,j}|$ and $|\Delta\phi_j|$ stay controlled even as n increases, suggesting that an ML model can serve as a reliable warm-start surrogate for FALQON's iterative measurement-feedback loop on larger problem sizes.

Our final MaxCut analysis compares the learned curves with a linear digitized quantum annealing heuristic. For this comparison, we use the same number of layers ℓ and the same time step Δt as in the corresponding FALQON simulation, so that the total digitized evolution time is $T = \ell\Delta t$. The linear schedule is evaluated at the discrete times $t_j = j\Delta t$, with $A(t_j) = 1 - t_j/T$ and $B(t_j) = t_j/T$. We do not separately optimize the total annealing time for each problem size or instance. Thus, this comparison should be interpreted as an equal-layer, equal-total-time comparison against a simple linear schedule, rather than as a comparison against an optimized annealing protocol.

In Fig. 6, we compare the MaxCut performance, as quantified by r_A and ϕ , of the ML-predicted parameter curves against standard linear QA schedules, digitized for implementation on circuit-model quantum computers. Given that the identification of optimal QA schedules is known to be difficult, the aim of this comparison is to explore whether the ML model predictions could have value for improving performance over other more commonly used QA schedule heuristics. In line with previous work [12], this comparison is carried out by fixing the total number of layers and comparing the methods' performance with this parameter held fixed. Dashed curves show the mean performance of the parameter curves predicted by the student model for different problem sizes, while dotted curves show the performance when a linear QA schedule is used. As a reference, exact FALQON results are shown as solid curves. We find that the ML model results track the FALQON mean curves closely for both r_A and ϕ and consistently outperform the linear schedule across the problem sizes we consider, $n \in \{8, 10, 12\}$.

Having established this behavior for weighted MaxCut, we repeat the analysis for the TFIM introduced in Sec. III B. For each weighted TFIM instance, we compare the weighted FALQON reference curve against two alternatives in line with our analysis for the MaxCut results, i.e., the first being the ML-predicted curve obtained from the weighted graph input and then replayed on the same weighted TFIM Hamiltonian, and the second being an unweighted graph baseline, obtained




FIG. 6. Mean approximation ratio and mean ground-state-subspace overlap are plotted vs layer for the exact FALQON reference (solid curves), a digitized linear-annealing baseline (dotted curves), and the ML model predictions (dashed curves) for solving weighted MaxCut. Panel (a) shows the mean approximation ratio r_A and panel (b) shows the mean ground-state-subspace overlap ϕ ; for $n \in \{8, 10, 12\}$, the means are computed over 40, 152, and 680 weighted graph instances, respectively. The magnitude of the associated standard deviations (not shown) of the FALQON reference curves and ML predictions are similar, and smaller on average compared with the linear-annealing baseline.

by running FALQON on the corresponding unweighted graph and replaying that parameter curve on the weighted TFIM Hamiltonian.

Figure 7 shows the TFIM results for $n \in \{8, 10, 12\}$, organized in the same format as Fig. 4. Across these system sizes, the ML-predicted curves remain close to the weighted FALQON reference and generally yield smaller errors than the unweighted baseline, indicating that the ML model is able to learn how the edge weights modify the FALQON trajectory for the TFIM.

The energy and overlap errors show the effect of the predicted curve on the actual TFIM dynamics. Since $r_{A,j}$ measures closeness to the ground-state energy and ϕ_j measures probability weight in the ground-state subspace, smaller errors indicate that the model captures weight-dependent information relevant to the ground-state preparation behavior of the FALQON trajectory.

We next test TFIM extrapolation beyond the training sizes, with results for $n \in \{14, 16\}$ shown in Fig. 8. For $n = 14$, the ML predictions remain reasonably close to the weighted FALQON reference and generally improve over the unweighted baseline. This behavior is consistent with the MaxCut extrapolation results and suggests that the model has learned structure that is useful beyond the exact training sizes. For $n = 16$, however, the quality of the extrapolation degrades, suggesting a finite generalization window. That is, the trained ML model captures useful low-dimensional structure in the FALQON parameter curves, but its reliability decreases as the system size moves farther from the training distribution.

Thus, the TFIM analysis supports the broader applicability of the approach beyond MaxCut, while also showing that robust extrapolation to larger systems will likely require larger training sets or transfer-learning strategies.

V. CONCLUSION AND OUTLOOK

In this work, we have explored the prospect of predicting FALQON parameter curves with classical machine learning. We have found that the teacher-student ML model developed here can effectively predict FALQON parameter curves associated with MaxCut and TFIM problems when fed a weighted graph as input. For both settings, the ML model reproduces key features of the reference FALQON curves and generally improves over an unweighted baseline within the regimes studied. These results indicate that once trained, an ML-driven approach can be a viable, measurement-free, and classical-optimization-free alternative to the iterative, measurement-based feedback procedure utilized in FALQON. We have also illustrated that the ML-predicted parameter curves can offer benefits in a digitized QA setting when used in place of conventional, linear QA schedules. The latter finding opens prospects for utilizing ML to design schedules for use in adiabatic quantum state preparation more broadly.

There are also important practical considerations for hardware implementations. First, the circuit depths considered in our simulations may be challenging for near-term quantum devices, where noise can accumulate over many layers. Second, the measurement savings discussed here concern the measurements needed for the generation of the FALQON parameter curve, i.e., as the trained ML model can remove the need for these measurements by outputting the full curve classically. This does not, however, remove the need to measure the final circuit output. In practice, one would still need to sample the terminal state to read out the desired results (e.g., problem Hamiltonian expectation value or sampled bit string).

It is also useful to relate these learned curves to ideas from quantum control and counterdiabatic driving [68–70]. Counterdiabatic protocols aim to suppress diabatic transitions by adding control terms constructed from spectral information about the instantaneous Hamiltonian. FALQON curves considered here also differ qualitatively from Floquet-engineered implementations of counterdiabatic driving, in which periodic modulations of the available controls are used to synthesize approximate counterdiabatic terms [71]. In contrast, FALQON parameter curves studied here are dominated by a smooth peak-and-tail structure and do not exhibit a pronounced high-frequency component. The learned FALQON curves studied here are not counterdiabatic ansätze in this formal sense, since they do not introduce additional Hamiltonian terms and are not derived from the adiabatic gauge potential. Nevertheless, they can play a related heuristic role since the learned schedule can deviate from a simple linear-annealing path in a way that improves performance.

There are a variety of ways in which this work could be extended. One important aspect is to further analyze the scalability of the approach. We carried out an initial investigation of this point in Fig. 5, where we found that the ML predictions begin to lag for $n = 18, 20$, indicating limited generalization and scalability to problem sizes as they move farther beyond




FIG. 7. Performance is plotted vs layer for the TFIM test instances with $n \in \{8, 10, 12\}$. The ML model predictions are shown in blue and the unweighted baseline is shown in red. Solid curves show the mean across problem instances, shaded bands indicate the associated standard deviation, and dashed curves indicate the maximum error across instances. The top row shows the parameter error $|\Delta\beta_j|$ on a logarithmic scale, the middle row shows the normalized-energy error $|\Delta r_{A,j}|$, and the bottom row shows the ground-state-subspace-overlap error $|\Delta\phi_j|$. Both the ML-predicted curves and the unweighted-baseline curves are replayed on the same weighted TFIM Hamiltonians and compared against the weighted FALQON reference. The ML model generally produces smaller errors than the unweighted baseline, indicating that it captures weight-dependent features of the TFIM dynamics.

the largest training size of $n = 12$. The TFIM results in Fig. 8 exhibit a similar trend where the ML predictions remain reasonably close to the FALQON reference for $n = 14$, while the quality of the extrapolation degrades for $n = 16$. Looking ahead, it would be interesting to train and validate the ML model on larger problem sizes to probe its ability to handle the growing complexity that arises as n is increased further. It would also be valuable to study the number of problem instances needed for training ML models, and how this scales as the models are extended. In the present work, we utilized 1344 instances for training our model, where the generation of each instance involved an exact numerical simulation of FALQON to output the reference parameter curve. Given that exact numerical simulations have costs that scale exponentially, in general, with the size of the problem, training data could alternatively be obtained from runs of FALQON on quantum hardware. Either way, the cost of generating sufficient training data must be balanced against the utility and scalability of the resulting ML model. Another direction that could support scalability would be to explore training an ML model to learn a small set of parameters entering into a parametrized function defining a FALQON parameter curve, rather than the

full curve directly. The use of a parametrized function could constitute a more parsimonious approach with the potential to scale more easily, if a suitably expressive functional form could be identified.

Another future direction could involve extending our approach beyond the problem classes considered here, by exploring how ML models trained on one ensemble of graphs (e.g., defined by fixed degree, given edge-weight distribution, or topology) perform when applied to different graph families or weight priors. Techniques such as domain adaptation [72–74] or metalearning [75,76] could be explored for helping ML models adjust to unseen instances, preserving performance without retraining from scratch. Broadening out further, ML models could also be explored for transfer between different problem classes, such as MaxCut and TFIM, as well as to additional problems to test the flexibility of the method and reveal the specific features of the problem that a generalized architecture must capture. In the present work, separate models are trained for MaxCut and TFIM, so whether a model trained on one Hamiltonian family can transfer directly to another remains an open question. Yet another aspect that could be considered in future work is to




FIG. 8. Extrapolation performance for TFIM instances with $n \in \{14, 16\}$, which were not included in training. The plot conventions match those of Fig. 7. For $n = 14$, the ML-predicted curves remain close to the weighted FALQON reference and generally improve over the unweighted baseline. For $n = 16$, the prediction becomes less stable, as shown by the larger parameter fluctuations and the increased energy and overlap errors at later layers. These results suggest that the model learns useful TFIM parameter-curve structure beyond the training sizes, while also revealing the onset of extrapolation limitations.

explore the impact of noise on the predictions, given that today's quantum computers exhibit numerous types of noise including pulse distortions, calibration errors, and decoherence. Given the freedom in choosing the driver Hamiltonian, another possibility could be to train an ML model to select it or to output a targeted, layerwise sequence of driver Hamiltonians to use. Across these different future directions, it will be important to identify suitable points of comparison to benchmark ML models' performance. Finally, the results reported here also open the prospect of using ML models, like the one we have designed, for predicting quantum controls, which could have broad applications across quantum science and technology. This prospect is in line with recent work using ML to learn control fields directly from classical optimal-control solvers [77].

ACKNOWLEDGMENTS

We gratefully acknowledge discussions with Jennie Si, James Larsen, Dominic Messina, Prakriti Biswas, and Lili Ye. This research was supported in part by an appointment to the National Nuclear Security Administration Minority Serving Institutions Internship Program (NNSA-MSIIP), sponsored by the National Nuclear Security Administration and administered by the Oak Ridge Institute for Science and Education. This work was supported by the Laboratory Directed Research and Development program (Projects No. 225952 and No. 233972) at Sandia National Laboratories, a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc. for the U.S. Department of Energy's National Nuclear Security

Administration under Contract No. DE-NA0003525. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

DATA AVAILABILITY

The data that support the findings of this article are not publicly available. The data are available from the authors upon reasonable request.

APPENDIX: MODEL ARCHITECTURE AND TRAINING DETAILS

In this Appendix, we provide further details on the ML model architecture and training procedure used in this work.

1. Data, features, scalars, and normalization

Each MaxCut instance is a weighted 3-regular graph $G = (V, E, w)$ with positive edge weights w_{ij} . We encode each graph into the input tensors of a GNN encoder, i.e., a matrix that contains information about node features and a list with the edge attributes. Because the graphs do not have intrinsic node attributes, we use a constant node feature $x_i = 1$ for every vertex i ; all information enters the GNN through connectivity and edge weights. In addition to message passing on the graph (see Sec. 2 of the Appendix), during training both teacher and student receive a small set of scalar features s that contains the following: the node count $|V|$; the ground-state eigenvalue of the problem Hamiltonian; the spectral gap of the problem Hamiltonian; the spectral gap of the midpoint Hamiltonian, defined as $(H_d + H_p)/2$; and an optional index feature that is determined by the index (after canonical labeling) of the minimum diagonal entry of H_p , where canonical labeling means a deterministic relabeling that assigns the same vertex ordering to isomorphic graphs so that indices are comparable across instances. We found the index feature to have negligible effect in our tests, so we omit it from the nominal model. To provide explicit size cues, we append $\sqrt{|V|}$ and $\ln |V|$ to the set of scalar features. Thus, we have five instance scalars (if the optional index feature is included) plus two size cues. All scalar components are standardized using training-split statistics (zero mean, unit variance), and the same procedure is applied at validation and test time. We train with batch size 1.

2. Encoder (message passing)

Given a weighted graph, a GNN encoder performs message passing for ℓ_{mp} GNN layers. Here, message passing refers to each node repeatedly gathering information from its neighboring nodes, which are weighted by the edges connecting them, and updating its own internal vector to reflect that neighborhood. That is, in each layer, every node updates its embedding by aggregating a learned function of its neighbors' embeddings and the incident edge weights. After ℓ_{mp} layers, the encoder produces node embeddings $\{h_v^{(\ell_{\text{mp}})} \in \mathbb{R}^d : v \in V\}$, i.e., learned d -dimensional vectors that summarize each node's local weighted neighborhood. Because the

neighbor-aggregation step is permutation invariant, the encoder is invariant to vertex relabelings.

3. Pooling and the graph-level vector

After the neighbor-update steps, each node v carries a vector $h_v^{(\ell_{\text{mp}})}$ that encodes its local environment. To describe the whole graph with a single object, we combine these node vectors into one fixed-length graph-level vector z . We do this by taking a few simple summaries over all nodes,

$$\mu = \frac{1}{|V|} \sum_v h_v, \quad \sigma = \sum_v h_v, \quad \pi = \max_v h_v,$$

and by using two standard learned summaries that adaptively weight or scan over the set of nodes: *GlobalAttention* (a learned, weighted average) and *Set2Set* (a small recurrent module applied to the unordered node set) [63,78,79]. We then concatenate these pieces to form

$$z = [\mu; \sigma; \pi; g_{\text{GA}}; g_{\text{S2S}}],$$

such that z collects both simple whole-graph statistics (mean/sum/max) and learned global summaries that can capture more structured correlations across the graph.

4. Readouts and teacher-student roles

Teacher (conditioned readout). The teacher stacks three message-passing convolutional blocks (GINEConv with edge attributes) to produce node embeddings, applies the multipooling heads A–E (A: mean, B: sum, C: max, D: GlobalAttention, and E: Set2Set), and concatenates them to z . A small *hypernetwork* H then takes five graph-level scalars s and outputs the weights and bias ($W(s)$, $b(s)$) of a two-stage readout (I and II), yielding the reference curve

$$\hat{\beta}_T = W(s)z + b(s).$$

This lets the final map adapt smoothly across instance sizes and coarse difficulty without retraining upstream layers.

Student (lightweight backbone with auxiliary scalar head). The student uses three *TransformerConv* layers with edge attributes, the same multipooling to obtain z , and two heads: (1) an auxiliary scalar head f_{scal} that predicts $\hat{s}$ (supervised only during training) and (2) a compact two-layer multilayer perceptron (MLP) (*curve head*) that takes $[z, \hat{s}]$ and outputs $\hat{\beta}^S$. Having the student also learn the graph statistics s encourages it to capture the same information that guides the teacher. At inference, the student builds a graph summary z , forms an internal estimate $\hat{s}$ from this summary using a small readout trained to match s , and then uses $[z, \hat{s}]$ to output the curve $\hat{\beta}^S$; we report only $\hat{\beta}^S$ and do not use the teacher.

5. Distillation objective (values, differences, and light smoothing)

By “distill” we mean that the student matches the teacher as a function of layer index, minimizing discrepancy on the full curve and its first/second finite differences, while lightly suppressing high-frequency oscillations. Let ∇ and ∇^2 denote first and second finite differences along the layer axis, $\mathcal{F}$ the real-input FFT, and TV total variation. The training objective

is

$$L_{\text{distill}} = c_1 \|\hat{\beta}^S - \hat{\beta}^T\|_2^2 + c_2 \|\nabla \hat{\beta}^S - \nabla \hat{\beta}^T\|_2^2 + c_3 \|\nabla^2 \hat{\beta}^S - \nabla^2 \hat{\beta}^T\|_2^2 + c_4 \sum_{\omega} \omega^4 |\mathcal{F}\{\hat{\beta}^S\}(\omega)|^2 + c_5 \text{TV}(\hat{\beta}^S) + c_6 \|\hat{s} - s\|_2^2. \quad (\text{A1})$$

Definitions (finite differences and transforms).

$$(\nabla \beta)_j = \beta_{j+1} - \beta_j, \quad 0 \leq j \leq \ell - 2, \quad (\text{A2})$$

$$(\nabla^2 \beta)_j = \beta_{j+1} - 2\beta_j + \beta_{j-1}, \quad 1 \leq j \leq \ell - 2, \quad (\text{A3})$$

$$\mathcal{F}\{\beta\}(\omega_m) = \sum_{j=0}^{\ell-1} \beta_j e^{-i\omega_m j}, \quad \omega_m := \frac{2\pi m}{\ell}, \quad m = 0, \dots, \ell - 1, \quad (\text{A4})$$

$$\text{TV}(\beta) := \sum_{j=0}^{\ell-2} |\beta_{j+1} - \beta_j|. \quad (\text{A5})$$

Intuition. c_1 matches level; c_2 matches slope (peak placement and tail decay rate); c_3 matches curvature (peak sharpness and tail bend); c_4 reduces high-frequency chatter (via ω^4); c_5 discourages step-to-step jumps; and c_6 aligns the scalar conditioning channel. Values are given in Table I.

6. Data, training protocol, and hyperparameters

We consider 3-regular graphs with $n \in \{4, 6, 8, 10, 12\}$. We generate 2240 instances and split them into 60% training

TABLE I. Model hyperparameters (Optuna tuned unless marked fixed).

Hyperparameter	Value
Teacher (fixed)	
Hidden dimension	96
Hypernetwork hidden	224
Output length ℓ	1001
GlobalAttention/Set2Set	True/True
Nominal Student [TransformerConv; loss weights in Eq. (A1)]	
c_1 (level match)	1.970 063 3
c_2 (slope match)	0.355 368 1
c_3 (curvature match)	1.761 085 2
c_4 (spectral penalty)	8.7863×10^{-2}
c_5 (total variation)	$5.554 861 8 \times 10^{-1}$
c_6 (auxiliary scalar head)	tuned; best c_6^*

Tunable values selected via Optuna (TPE) on validation loss; teacher settings are fixed.

(of which 80% are used for parameter updates and 20% for validation) and 40% testing. We train for 100 epochs with Adam (learning rate 10^{-3}), batch size 1 (per-graph curve supervision). The model outputs curves of length $\ell = 1001$ to match the FALQON depth used in our numerical experiments. The coefficients in Eq. (A1) are those in Table I. The forward path follows Fig. 3: three message-passing blocks (1–3), multipooling heads (A–E), and a two-stage readout (I and II). The teacher uses a conditioned (hypernetwork) readout fed by s ; the student adds the auxiliary scalar head $\hat{s}$ that is estimated during training by a small MLP and included in the loss function. Training the nominal student on a single GPU takes on the order of hours, while a trained student predicts a full curve in milliseconds per graph.

- [1] J. Kempe, A. Kitaev, and O. Regev, The complexity of the local Hamiltonian problem, *SIAM J. Comput.* **35**, 1070 (2006).
- [2] D. Poulin and P. Wocjan, Preparing ground states of quantum many-body systems on a quantum computer, *Phys. Rev. Lett.* **102**, 130503 (2009).
- [3] A. Lucas, Ising formulations of many NP problems, *Front. Phys.* **2**, 5 (2014).
- [4] T. Kadowaki and H. Nishimori, Quantum annealing in the transverse Ising model, *Phys. Rev. E* **58**, 5355 (1998).
- [5] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, Quantum computation by adiabatic evolution, *arXiv:quant-ph/0001106*.
- [6] A. Das and B. K. Chakrabarti, *Colloquium*: Quantum annealing and analog quantum computation, *Rev. Mod. Phys.* **80**, 1061 (2008).
- [7] T. Albash and D. A. Lidar, Adiabatic quantum computation, *Rev. Mod. Phys.* **90**, 015002 (2018).
- [8] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, Variational quantum algorithms, *Nat. Rev. Phys.* **3**, 625 (2021).
- [9] K. Bharti, A. Cervera-Lierta, T. H. Kyaw, T. Haug, S. Alperin-Lea, A. Anand, M. Degroote, H. Heimonen, J. S. Kottmann, T. Menke, W.-K. Mok, S. Sim, L.-C. Kwek, and A. Aspuru-Guzik, Noisy intermediate-scale quantum algorithms, *Rev. Mod. Phys.* **94**, 015004 (2022).
- [10] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, *arXiv:1411.4028*.
- [11] A. B. Magann, K. M. Rudinger, M. D. Grace, and M. Sarovar, Feedback-based quantum optimization, *Phys. Rev. Lett.* **129**, 250502 (2022).
- [12] A. B. Magann, K. M. Rudinger, M. D. Grace, and M. Sarovar, Lyapunov-control-inspired strategies for quantum combinatorial optimization, *Phys. Rev. A* **106**, 062414 (2022).
- [13] S. Grivopoulos and B. Bamieh, Lyapunov-based control of quantum systems, in *42nd IEEE International Conference on Decision and Control* (IEEE, Piscataway, NJ, 2003), Vol. 1, pp. 434–438.
- [14] S. Cong and F. Meng, A survey of quantum Lyapunov control methods, *Sci. World J.* **2013**, 967529 (2013).
- [15] M. Alam, A. Ash-Saki, and S. Ghosh, Accelerating quantum approximate optimization algorithm using machine learning, in *2020 Design, Automation & Test in Europe Conference & Exhibition* (IEEE, Piscataway, NJ, 2020), pp. 686–689.
- [16] N. Jain, B. Coyle, E. Kashefi, and N. Kumar, Graph neural network initialisation of quantum approximate optimisation, *Quantum* **6**, 861 (2022).

- [17] N. Xie, X. Lee, D. Cai, Y. Saito, and N. Asai, Quantum approximate optimization algorithm parameter prediction using a convolutional neural network, *J. Phys.: Conf. Ser.* **2595**, 012001 (2023).
- [18] P. Zentilini, S. Corli, and E. Prati, Emulating QAOA via graph neural networks, in *2024 IEEE International Conference on Quantum Computing and Engineering (QCE)* (IEEE, Piscataway, NJ, 2024), pp. 472–473.
- [19] I. Tyagin, M. H. Farag, K. Sherbert, K. Shirali, Y. Alexeev, and I. Safro, QAOA-GPT: Efficient generation of adaptive and regular quantum approximate optimization algorithm circuits, in *2025 IEEE International Conference on Quantum Computing and Engineering (QCE)* (2025), Vol. 1, pp. 1505–1515.
- [20] J. Roland and N. J. Cerf, Quantum search by local adiabatic evolution, *Phys. Rev. A* **65**, 042308 (2002).
- [21] M. Born and V. Fock, Beweis des adiabatensatzes, *Z. Phys.* **51**, 165 (1928).
- [22] L. P. García-Pintos, L. T. Brady, J. Bringewatt, and Y.-K. Liu, Lower bounds on quantum annealing times, *Phys. Rev. Lett.* **130**, 140601 (2023).
- [23] L. P. García-Pintos, M. Sahasrabudhe, and C. Arenz, Tighter lower bounds on quantum annealing times, *SciPost Phys.* **18**, 159 (2025).
- [24] G. B. Mbeng, R. Fazio, and G. Santoro, Quantum annealing: A journey through digitalization, control, and hybrid quantum variational schemes, [arXiv:1906.08948](https://arxiv.org/abs/1906.08948).
- [25] L. K. Kovalsky, F. A. Calderon-Vargas, M. D. Grace, A. B. Magann, J. B. Larsen, A. D. Baczewski, and M. Sarovar, Self-healing of Trotter error in digital adiabatic state preparation, *Phys. Rev. Lett.* **131**, 060602 (2023).
- [26] S. Lloyd, Universal quantum simulators, *Science* **273**, 1073 (1996).
- [27] H. F. Trotter, On the product of semi-groups of operators, *Proc. Am. Math. Soc.* **10**, 545 (1959).
- [28] M. Suzuki, Generalized Trotter's formula and systematic approximants of exponential operators and inner derivations with applications to many body problems, *Commun. Math. Phys.* **51**, 183 (1976).
- [29] M. Wiedmann, D. Burgarth, G. Dirr, T. Schulte-Herbrüggen, E. Malvetti, and C. Arenz, On the convergence of the variational quantum eigensolver and quantum optimal control, [arXiv:2509.05295](https://arxiv.org/abs/2509.05295).
- [30] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, *Nat. Commun.* **9**, 4812 (2018).
- [31] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, Cost function dependent barren plateaus in shallow parametrized quantum circuits, *Nat. Commun.* **12**, 1791 (2021).
- [32] S. Hadfield, Z. Wang, B. O'Gorman, E. G. Rieffel, D. Venturelli, and R. Biswas, From the quantum approximate optimization algorithm to a quantum alternating operator ansatz, *Algorithms* **12**, 34 (2019).
- [33] 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).
- [34] L. T. Brady, C. L. Baldwin, A. Bapat, Y. Kharkov, and A. V. Gorshkov, Optimal protocols in quantum annealing and quantum approximate optimization algorithm problems, *Phys. Rev. Lett.* **126**, 070505 (2021).
- [35] L. T. Brady, L. Kocia, P. Bienias, A. Bapat, Y. Kharkov, and A. V. Gorshkov, Behavior of analog quantum algorithms, [arXiv:2107.01218](https://arxiv.org/abs/2107.01218).
- [36] 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:1812.04170](https://arxiv.org/abs/1812.04170).
- [37] D. J. Egger, J. Marecek, and S. Woerner, Warm-starting quantum optimization, *Quantum* **5**, 479 (2021).
- [38] S. H. Sack and M. Serbyn, Quantum annealing initialization of the quantum approximate optimization algorithm, *Quantum* **5**, 491 (2021).
- [39] L. Zhu, H. L. Tang, G. S. Barron, F. A. Calderon-Vargas, N. J. Mayhall, E. Barnes, and S. E. Economou, Adaptive quantum approximate optimization algorithm for solving combinatorial problems on a quantum computer, *Phys. Rev. Res.* **4**, 033029 (2022).
- [40] A. B. Magann, S. E. Economou, and C. Arenz, Randomized adaptive quantum state preparation, *Phys. Rev. Res.* **5**, 033227 (2023).
- [41] R. Wiersema and N. Killoran, Optimizing quantum circuits with Riemannian gradient flow, *Phys. Rev. A* **107**, 062421 (2023).
- [42] H. L. Tang, Y. Chen, P. Biswas, A. B. Magann, C. Arenz, and S. E. Economou, Nonvariational ADAPT algorithm for quantum simulations, *Phys. Rev. Res.* **7**, 023275 (2025).
- [43] E. Malvetti, C. Arenz, G. Dirr, and T. Schulte-Herbrüggen, Randomized gradient descents on Riemannian manifolds: Almost sure convergence to global minima in and beyond quantum optimization, [arXiv:2405.12039](https://arxiv.org/abs/2405.12039).
- [44] N. A. McMahon, M. Pervez, and C. Arenz, Equating quantum imaginary time evolution, Riemannian gradient flows, and stochastic implementations, *Phys. Rev. Res.* **8**, 023024 (2026).
- [45] H. G. Clausen, S. A. Rahman, Ö. Karabacak, and R. Wisniewski, Measurement-based control for minimizing energy functions in quantum systems, *IFAC-Pap.* **56**, 5171 (2023).
- [46] J. B. Larsen, M. D. Grace, A. D. Baczewski, and A. B. Magann, Feedback-based quantum algorithms for ground state preparation, *Phys. Rev. Res.* **6**, 033336 (2024).
- [47] S. Abdul Rahman, O. Karabacak, and R. Wisniewski, Weighted feedback-based quantum algorithm for excited states calculation, in *2024 IEEE International Conference on Quantum Computing and Engineering (QCE)* (IEEE, Piscataway, NJ, 2024), pp. 169–175.
- [48] R. K. Malla, H. Sukeno, H. Yu, T.-C. Wei, A. Weichselbaum, and R. M. Konik, Feedback-based quantum algorithm inspired by counterdiabatic driving, *Phys. Rev. Res.* **6**, 043068 (2024).
- [49] L. Bertuzzi, J. P. Engster, E. C. R. da Rosa, and E. I. Duzzioni, Shadow measurements for feedback-based quantum optimization, *Phys. Rev. A* **112**, 022419 (2025).
- [50] D. Arai, K. N. Okada, Y. Nakano, K. Mitarai, and K. Fujii, Scalable circuit depth reduction in feedback-based quantum optimization with a quadratic approximation, *Phys. Rev. Res.* **7**, 013035 (2025).
- [51] L. A. M. Rattighieri, G. E. L. Peixe, B. L. Bernardo, and F. F. Fanchini, Accelerating feedback-based quantum algorithms through time rescaling, *Phys. Rev. A* **112**, 042607 (2025).

- [52] S. Abdul Rahman, Ö. Karabacak, and R. Wisniewski, Feedback-based quantum strategies for constrained combinatorial optimization problems, *Future Gener. Comput. Syst.* **174**, 107979 (2026).
- [53] R. Kosloff, A. D. Hammerich, and D. Tannor, Excitation without demolition: Radiative excitation of ground-surface vibration by impulsive stimulated Raman scattering with damage control, *Phys. Rev. Lett.* **69**, 2172 (1992).
- [54] M. Sugawara and Y. Fujimura, Control of quantum dynamics by a locally optimized laser field. Application to ring puckering isomerization, *J. Chem. Phys.* **100**, 5646 (1994).
- [55] M. Sugawara and Y. Fujimura, Control of quantum dynamics by a locally optimized laser field. Multi-photon dissociation of hydrogen fluoride, *Chem. Phys.* **196**, 113 (1995).
- [56] Y. Ohtsuki, Y. Yahata, H. Kono, and Y. Fujimura, Application of a locally optimized control theory to pump-dump laser-driven chemical reactions, *Chem. Phys. Lett.* **287**, 627 (1998).
- [57] M. Sugawara, General formulation of locally designed coherent control theory for quantum system, *J. Chem. Phys.* **118**, 6784 (2003).
- [58] M. Mirrahimi, G. Turinici, and P. Rouchon, Reference trajectory tracking for locally designed coherent quantum controls, *J. Phys. Chem. A* **109**, 2631 (2005).
- [59] D. J. Tannor, R. Kosloff, and A. Bartana, Laser cooling of internal degrees of freedom of molecules by dynamically trapped states, *Faraday Discuss.* **113**, 365 (1999).
- [60] V. Engel, C. Meier, and D. J. Tannor, Local control theory: Recent applications to energy and particle transfer processes in molecules, in *Advances in Chemical Physics* (John Wiley & Sons, Hoboken, NJ, 2009), pp. 29–101.
- [61] M. X. Goemans and D. P. Williamson, Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming, *J. ACM* **42**, 1115 (1995).
- [62] G. Hinton, O. Vinyals, and J. Dean, Distilling the knowledge in a neural network, this work was presented at the NIPS 2014 Deep Learning Workshop, [arXiv:1503.02531](https://arxiv.org/abs/1503.02531).
- [63] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl, Neural message passing for quantum chemistry, in *Proceedings of the 34th International Conference on Machine Learning (ICML 2017)* (PMLR, Cambridge, MA, 2017), Vol. 70, pp. 1263–1272.
- [64] W. Hu, B. Liu, J. Gomes, M. Zitnik, P. Liang, V. S. Pande, and J. Leskovec, Strategies for pre-training graph neural networks, in *8th International Conference on Learning Representations (ICLR 2020)*, Addis Ababa, Ethiopia, April 26–30, 2020 (OpenReview.net, Amherst, MA, 2020).
- [65] Y. Shi, Z. Huang, S. Feng, H. Zhong, W. Wang, and Y. Sun, Masked label prediction: Unified message passing model for semi-supervised classification, in *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence (IJCAI-21)*, edited by Z.-H. Zhou (International Joint Conferences on Artificial Intelligence Organization, California, 2021), pp. 1548–1554.
- [66] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, *International Conference on Learning Representations (ICLR)* (2015).
- [67] T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama, Optuna: A next-generation hyperparameter optimization framework, in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Association for Computing Machinery, New York, NY, USA, 2019), pp. 2623–2631.
- [68] 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).
- [69] D. Sels and A. Polkovnikov, Minimizing irreversible losses in quantum systems by local counterdiabatic driving, *Proc. Natl. Acad. Sci. USA* **114**, E3909 (2017).
- [70] M. Kolodrubetz, D. Sels, P. Mehta, and A. Polkovnikov, Geometry and non-adiabatic response in quantum and classical systems, *Phys. Rep.* **697**, 1 (2017).
- [71] 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).
- [72] S. J. Pan and Q. Yang, A survey on transfer learning, *IEEE Trans. Knowl. Data Eng.* **22**, 1345 (2010).
- [73] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky, Domain-adversarial training of neural networks, *J. Mach. Learn. Res.* **17**, 1 (2016).
- [74] E. Tzeng, J. Hoffman, K. Saenko, and T. Darrell, Adversarial discriminative domain adaptation, in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (IEEE, Piscataway, NJ, 2017), pp. 2962–2971.
- [75] C. Finn, P. Abbeel, and S. Levine, Model-agnostic meta-learning for fast adaptation of deep networks, in *Proceedings of the 34th International Conference on Machine Learning (ICML)* (PMLR, Cambridge, MA, 2017), Vol. 70, pp. 1126–1135.
- [76] T. Hospedales, A. Antoniou, P. Micaelli, and A. Storkey, Meta-learning in neural networks: A survey, *IEEE Trans. Pattern Anal. Mach. Intell.* **44**, 5149 (2022).
- [77] E. Khaleghian, A. Fath Lipaei, A. Bahrampour, M. Nikaeen, and A. Bahrampour, Development of neural network-based optimal control pulse generator for quantum logic gates using the GRAPE algorithm in NMR quantum computer *Quantum Inf. Process.* **25**, 131 (2026).
- [78] Y. Li, D. Tarlow, M. Brockschmidt, and R. Zemel, Gated graph sequence neural networks, *International Conference on Learning Representations (ICLR)* (2016).
- [79] O. Vinyals, S. Bengio, and M. Kudlur, Order matters: Sequence to sequence for sets, *International Conference on Learning Representations (ICLR)* (2016).