

End-To-End Quantum Machine Learning Implemented with Controlled Quantum Dynamics


Re-Bing Wu,^{1,2,*} Xi Cao,¹ Pinchen Xie³ and Yu-xi Liu⁴

¹*Department of Automation, Tsinghua University, Beijing 100084, China*

²*Beijing National Research Center for Information Science and Technology, Beijing 100084, China*

³*Program in Applied and Computational Mathematics, Princeton University, Princeton New Jersey 08544, USA*

⁴*Institute of Microelectronics, Tsinghua University, Beijing 100084, China*

 (Received 1 April 2020; revised 13 July 2020; accepted 23 November 2020; published 7 December 2020)

To achieve quantum machine learning on imperfect noisy intermediate-scale quantum (NISQ) processors, the entire physical implementation should include as few as possible hand-designed modules with only a few ad hoc parameters to be determined. This work presents such a hardware-friendly end-to-end quantum-machine-learning scheme that can be implemented with imperfect NISQ processors. The proposal transforms the machine-learning task to the optimization of controlled quantum dynamics, in which the learning model is parameterized by experimentally tunable control variables. Our design also enables automated feature selection by encoding the raw input to quantum states through agent control variables. Comparing with the gate-based parameterized quantum circuits, the proposed end-to-end quantum-learning model is easy to implement as there are only a few ad hoc parameters to be determined. Numerical simulations on the benchmarking MNIST (Mixed National Institute of Standards and Technology) dataset of handwritten digits demonstrate that the model can achieve high performance using only 3–5 qubits without downsizing the dataset, which shows great potential for accomplishing large-scale real-world learning tasks on NISQ processors.

DOI: [10.1103/PhysRevApplied.14.064020](https://doi.org/10.1103/PhysRevApplied.14.064020)

I. INTRODUCTION

Quantum computing has entered the noisy intermediate-scale quantum (NISQ) era [1] in which it may surpass classical computing with even imperfect quantum hardware [2]. As one of the most promising applications, quantum machine learning is drawing intense attention [3,4] for its potential supremacy on solving large-scale real-world learning tasks with quantum computers. With ideal programmable and error-tolerant quantum computers, many quantum subroutines such as quantum Fourier transform or Grover search can be applied to speed up the training or inference process, e.g., quantum supporting vector machine for classification problems [5], quantum principal component analysis [6], and quantum generative adversarial learning [7,8].

To enable quantum-machine-learning algorithms on NISQ processors, a realistic approach is to construct quantum neural-network (NN) models with parameterized quantum circuits (PQCs) [9,10] that is trained by classical optimization algorithms. Such hybrid quantum-classical models have universal approximation capabilities and are able to achieve classically intractable feature learning

tasks [11]. Various applications have been put forward for quantum simulation [12], combinatorial optimization [13], and machine-learning problems [14]. Moreover, feature-selection schemes were also proposed, such as the quantum kitchen sink [15] and quantum metric learning [16], to efficiently encode the data into quantum states by learning the representation.

The PQC-based machine learning has also been experimentally demonstrated by shallow circuits on NISQ processors (with no greater than 20 qubits) [10] on classification [17], clustering [18], and generative [19] learning tasks. To our knowledge, few of them were tested with real-world datasets, and the achieved performance was still poor unless the dataset is downsized. For example, the precision of recognizing handwritten digits in the benchmarking Mixed National Institute of Standards and Technology (MNIST) dataset [20] is no higher than what can be achieved by a simple classical logistic regression model, and in most cases the original images have to be downsampled to make compromises with scarce quantum resources (e.g., limited number of qubits and decoherence time) [14,15,21–25].

On top of limited quantum resources, the architecture of the current PQC ansatz also challenges the development of practical NISQ processors, because the designed

*rbwu@tsinghua.edu.cn

gate sequences cannot be trivially implemented due to the incompatibility of the PQC topology with the available quantum processors with sparse qubit-qubit connectivity. For example, nonlocal gates may need to be realized through a series of intermediate operations (e.g., SWAP) due to the lack of interactions between target qubits, while local gates may not be directly implementable when the qubit-qubit coupling cannot be freely turned off. Hence, the designed circuit must be properly mapped to the actual interconnect topology of a quantum chip, and consequently the compiled quantum circuit is usually deeper and more complicated.

All these demands call for a hardware-friendly quantum-machine-learning scheme that can be efficiently deployed on NISQ processors with full consideration of their on-chip interconnect topologies. The entire physical implementation should include as few as possible hand-designed modules with only a few ad hoc parameters to be determined, so that as few as possible errors are introduced to the quantum-machine-learning process. In other words, the scheme should provide an end-to-end data pipeline that automatically extracts the features for inference from the input data, so that the overall performance can be better improved. This has become an influential trend in classical deep learning, especially in big-data applications [26].

In this paper, we propose that such an end-to-end quantum-learning model can be naturally realized by dynamical quantum evolution manipulated by the laboratory hardware control devices. In the following, we show in Sec. II how this can be done by reparameterizing the PQC with control variables and by introducing a data-to-control interface for automatic feature selection. Then, in Sec. III we provide the hybrid quantum-classical training algorithm, following which simulation examples with 3–5 qubits are given to demonstrate the effectiveness of the proposed end-to-end learning scheme in Sec. IV. Finally, concluding remarks are made in Sec. V.

II. QUANTUM END-TO-END LEARNING MODEL

To illustrate how the end-to-end quantum learning can be implemented by controlled quantum dynamics, let us start from a classification learning task with a set of Z training samples $[x^{(k)}, y^{(k)}]$, where $x^{(k)} \in \mathbb{R}^d$ is the input data represented by d -dimensional vectors and $y^{(k)} \in \{1, \dots, r\}$ is the corresponding label. For most quantum-machine-learning models, the input x is first transferred to the quantum state $|\Psi(x)\rangle$ of the register through an encoder circuit. After being processed by a succeeding quantum circuit represented by a parameterized unitary transformation $U(\mathbf{w})$, where \mathbf{w} is the hyperparameters (e.g., rotating angles of X , Y , or Z gates) of the circuit. The output state is measured under a positive operator-valued measurement (POVM) $\{M_1, \dots, M_r\}$, in which the operator M_k is associated with

the k th class to be discriminated. The conditional probability of obtaining y for a given input x and circuit $U(\mathbf{w})$ is then $P(y|x, \mathbf{w}) = \langle \Psi(x) | U^\dagger(\mathbf{w}) M_y U(\mathbf{w}) | \Psi(x) \rangle$, based on which we can define the empirical loss:

$$L[\mathbf{w}] = 1 - Z^{-1} \sum_{k=1}^Z P[y^{(k)} | x^{(k)}, \mathbf{w}]. \quad (1)$$

The commonly applied PQC model for machine learning usually consists of layered parameterized one-qubit or two-qubit quantum gates. As discussed above, the assigned gates in such a black model may not be directly implementable due to the limited qubit-qubit connectivity, and hence a more hardware-friendly scheme is desired.

A. From circuit model to control model

In practice, the PQCs performed on quantum chips are always realized through a set of hardware control and measurement devices. For example, in the experimental superconducting quantum-computing system [27–29] shown in Fig. 1, the entire PQC is dictated by the control pulses produced from an arbitrary waveform generator for implementing the individual designed quantum gates, and the inference is made by readout the qubit states. Therefore, it is natural to replace by these control amplitudes the gate parameters in the unitary transformation $U(\mathbf{w})$ realized by

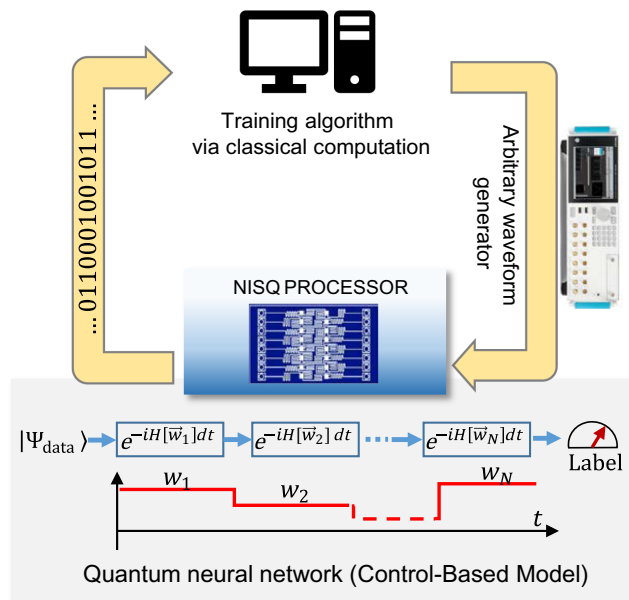


FIG. 1. The physical control associated with a hybrid quantum-classical algorithm deployed on a NISQ processor. The classical computer runs the training algorithm to instruct the arbitrary waveform generator (AWG) to iteratively update the control pulses according to the qubit measurement outcomes. The controlled quantum evolution forms an equivalent quantum NN parameterized by time-dependent AWG parameters.

the PQC. An obvious and significant advantage of such a control-based model is the hardware friendliness because the control parameters are directly tunable in experiments.

The control-based model can also be treated as a layered quantum feedforward neural network represented by its time-evolving quantum dynamics steered by the Schrödinger equation:

$$i\dot{|\Psi(t)\rangle} = -i \left[H_0 + \sum_{\ell=1}^m w_{\ell}(t) H_{\ell} \right] |\Psi(t)\rangle, \quad (2)$$

where $|\Psi(t)\rangle$ is the quantum state [initially prepared at $|\Psi(t_0)\rangle = |0\rangle$] of the entire system. When the AWG pulses $w_1(t), \dots, w_m(t)$ consists of M piecewise-constant sub-pulses over M sampling periods, the states $|\Psi(t_k)\rangle$ ($k = 1, 2, \dots, M$) at the end of each subinterval form the layers of the quantum NN. The interconnection between these layers are realized by the unitary evolution operators over these subintervals, which are parameterized by control variables $\vec{w}_k = [w_1(t_k), \dots, w_M(t_k)]$. The equivalent depth of the quantum NN is thus the number of AWG sampling periods during the entire quantum evolution, while the width is determined by the number of qubits.

The control-based model is a generalization of the gate-based PQC model because any gate operation must be eventually realized through physical control pulses. In special cases when the control Hamiltonians are mutually commutable, it is equivalent with a PQC model, because each control parameter governs a parameterized gate. However, under more general circumstances with limited tuning ability of qubit-qubit couplings, the compilation of gate-based PQCs becomes much more complicated, but the control-based scheme can easily adapt to the on-chip interconnect topology without having to artificially split the model into separate gates.

B. From hand-designed to autoselected features

In most PQC-based learning models, the data vector is mapped to the quantum state using a preselected encoder to represent the set of hand-designed features. As schematically shown in Fig. 2(a), the encoder first “translates” the data vector to a quantum state, and then applies a control pulse to physically prepare the system in this state. The complexity of the control design depends on the encoded state, which could be very expensive when the state is highly entangled (e.g., in the amplitude encoding scheme for exploiting the superposition of quantum states). The scheme also becomes impractical when dealing with large-size datasets because every single sample needs an individually designed control pulse.

We propose that the “translation” from the data vector to the quantum state can be designed in an *implicit* and *automatic* manner. As shown in Fig. 2(b), a data-to-control interface (e.g., a classical NN) is introduced to transform

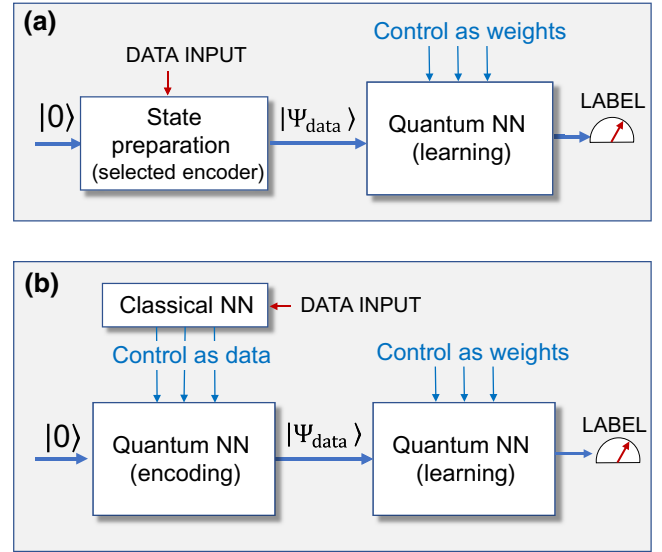


FIG. 2. Quantum control-based end-to-end learning models. (a) The data is encoded to the quantum state through a preselected encoder; (b) the data is encoded to the quantum state through a data-to-control interface (a classical NN) via a selected set of agent control variables. Following the encoding control processes, inference control pulses are applied as weights to infer the class that the input belongs to.

the data vector into a selected set of agent control variables, which then conveys the information about the received data to the quantum state they steer to. The encoded state is not explicitly (and nor necessarily) known unless being reconstructed through quantum-state tomography.

The introduced data-to-control interface can be taken as a hidden NN layer that feeds the classical data into the quantum NN. It can be trained together with the rest part of the quantum NN. Once the interface is determined, the encoding control pulse will be automatically generated in response to the input sample.

It should be noted that the encoding scheme also brings favored nonlinearity through the nonlinear control-to-state mapping, which is crucial for improving the model expressivity in complex learning tasks. Later we show in the simulation examples that the nonlinearity in control-to-state mapping plays a crucial role in achieving high learning performance.

III. THE TRAINING PROCESS

Now we have built up a machine-learning model based on the controlled quantum dynamics. The training of the learning model can thus be naturally transformed to a quantum optimal control problem [30]. Suppose that the involved control pulses contains M sampling periods. We assign the control variables $w^{\text{code}} = \{\vec{w}_1, \dots, \vec{w}_{M_0}\}$ in the first M_0 sampling periods for data encoding, and

$w^{\text{infer}} = \{\vec{w}_{M_0+1}, \dots, \vec{w}_M\}$ in the rest sampling periods for the inference.

The data-to-control interface can be chosen as arbitrary linear or nonlinear function. For illustration, we select a perceptron layer, i.e., each element of w^{code} is

$$w_i^{\text{code}} = B \frac{e^{\sum_j W_{ij} x_j + b_i} - 1}{e^{\sum_j W_{ij} x_j + b_i} + 1}, \quad (3)$$

where x_j is the j th element of the input $x \in \mathbb{R}^d$, $W = \{W_{ij}\}$ and $b = \{b_i\}$ are the weight matrix and bias vector of the perceptrons, and $B > 0$ is the bound of the control amplitudes. Because the bias term b can be merged into W by extending x to $(x^T, 1)^T$ and W to (W, b) , we ignore b for simplicity. Thus, the hyperparameters to be trained are $\mathbf{w} = (W, w^{\text{infer}})$.

The model is trained by minimizing the empirical loss $L[\mathbf{w}]$ defined by Eq. (1). Similar to most hybrid quantum-classical algorithms, these hyperparameters are to be tuned along gradient-descent directions of $L[\mathbf{w}]$. Since the gradient vector is not directly computable on the NISQ processors, we need to sequentially perturb each hyperparameter, evaluate the change of empirical loss via ensemble measurements on $L(\mathbf{w})$, and estimate its gradient with respect to $\mathbf{w} = (W, w^{\text{infer}})$ via the finite difference:

$$\frac{\partial L}{\partial \mathbf{w}_j} \approx \frac{L(\mathbf{w} + \Delta \cdot \mathbf{e}_j) - L(\mathbf{w})}{\Delta}, \quad (4)$$

where \mathbf{e}_j is the unit vector along which the j th element of \mathbf{w} is perturbed by Δ .

Let n be the number of qubits, and each qubit is manipulated by c independent control fields (e.g., the bias field for frequency tuning or Rabi driving fields for flipping qubits). Then there are $N_{\text{code}} = cnM_0$ encoding control variables to be generated by $(d+1)N_{\text{code}}$ weight variables in W and $N_{\text{infer}} = cn(M - M_0)$ inference control variables to be directly tuned. This implies that, to evaluate the gradient with a given input sample, about $(d+1)N_{\text{code}} + N_{\text{infer}}$ ensemble measurements will be required on the conditional probability $P(y|x, \mathbf{w})$. The experimental overhead can easily exceed the ability of current NISQ processors for large-size and high-dimensional datasets.

Nonetheless, observing that the gradient of $P(y|x, \mathbf{w})$ with respect to the entries of W can be decomposed (via the chain rule) as

$$\frac{\partial P(y|x, \mathbf{w})}{\partial W_{ij}} = \frac{\partial P(y|x, \mathbf{w})}{\partial w_i^{\text{code}}} \frac{B^2 - (w_i^{\text{code}})^2}{2B} x_j, \quad (5)$$

where $1 \leq i \leq N_{\text{code}}$ and $1 \leq j \leq d+1$, we only need to experimentally measure $\frac{\partial P(y|x, \mathbf{w})}{\partial w_i^{\text{code}}}$, with the rest parts handled by a classical computer. In this way, the experimental burden can be greatly relieved because the number

Input: training dataset $\{x^{(k)}, y^{(k)}\}$, the mini-batch size m , and the learning rate α .

Output: the hyper-parameter $\mathbf{w} = (W, w^{\text{infer}})$.

repeat

Sample a mini-batch from the training dataset.

for $i = 1$ **to** m **do**

Feed the i th sample in the mini-batch, say (\bar{x}, \bar{y}) , to generate the encoding control variables w^{code} via W .

Synthesize the full control pulse with current w^{code} and w^{infer} .

Perturb the control variables and measure their gradients $\frac{\partial P(\bar{y}|\bar{x}, \mathbf{w})}{\partial w^{\text{code}}}$ and $\frac{\partial P(\bar{y}|\bar{x}, \mathbf{w})}{\partial w^{\text{infer}}}$, and use Eq. (5) to evaluate $\frac{\partial P(\bar{y}|\bar{x}, \mathbf{w})}{\partial W}$ from $\frac{\partial P(\bar{y}|\bar{x}, \mathbf{w})}{\partial w^{\text{code}}}$.

end for

Compute the average gradient $\frac{\partial L[\mathbf{w}]}{\partial W}$ and $\frac{\partial L[\mathbf{w}]}{\partial w^{\text{infer}}}$ over the selected mini-batch, and update \mathbf{w} with a selected optimizer (e.g., Adam).

until Empirical loss is sufficiently small.

Algorithm 1. Quantum end-to-end learning.

of required ensemble experiments is reduced from $(d+1)N_{\text{code}} + N_{\text{infer}}$ to $N_{\text{code}} + N_{\text{infer}}$ (the total number of control variables), which is not explicitly dependent on the dimensionality of the data space.

Based on the measured gradient, we can apply the widely used stochastic gradient algorithms for machine learning, which had been demonstrated to be powerful in robust quantum control [31] and quantum approximate optimization algorithms [32]. Roughly speaking, in each iteration we randomly select a small batch of samples, apply the encoding and inference control fields, and measure the conditional probability $P(y|x, \mathbf{w})$ and its gradient for each sample. The averaged gradient over these samples is then used to update the model hyperparameters $\mathbf{w} = (W, w^{\text{infer}})$. The detailed pseudocode of the training algorithm can be found in Algorithm 1.

IV. SIMULATION RESULTS

Now we apply the proposed end-to-end learning model to the MNIST dataset for recognition of handwritten digits. To demonstrate the effectiveness and efficiency, we use a simple chain system of $n = 3 \sim 5$ qubits, which is typical in solid-state quantum computing, as the physical realization of the NISQ processor. The Hamiltonian reads

$$H(t) = \sum_{1 \leq i < n} g_{ij} \sigma_z^i \sigma_z^{i+1} + \sum_{k=1}^n \left[w_{kx}(t) \sigma_x^k + w_{ky}(t) \sigma_y^k \right], \quad (6)$$

where σ_α^k ($\alpha = x, y, z$) are standard Pauli matrices for the qubits. The neighboring qubit-qubit coupling strengths are $g_{12} = 1.5$ MHz, $g_{23} = 2.0$ MHz, $g_{34} = 2.5$ MHz, and $g_{45} = 3.0$ MHz, respectively. These qubits are addressed

by control fields $w_{kx}(t)$ and $w_{ky}(t)$ along the x and y axis, respectively. In all simulations, we fix the AWG sampling periods as 5 ns and set the control bounds $B = 25$ MHz.

A. The training process

To train the learning model, we use 46 993 samples associated to eight digits $\{0, 2, 3, 4, 5, 6, 8, 9\}$ (because three-qubit models can discriminate at most eight digits). The POVM for inference is chosen to be $\{M_j = |j\rangle\langle j|, j = 000, 001, \dots, 111\}$ under the σ_z basis of the first three qubits. The original 28×28 -pixel sample images are converted to $d = 28 \times 28 + 1 = 785$ dimensional vectors after merging the bias vector b into W .

We first train learning models with fixed depth (all using ten encoding layers and ten inference layers) and vary the number of qubits from three to five, and an additional three-qubit model with 50 coding layers and 50 inference layers for comparison. From the learning curves shown in Fig. 3, the deeper model learns remarkably better, in which the empirical loss can be reduced to below 10% after a few epochs (an epoch means that all training samples are traversed for once). In contrast, the performance remains almost unimproved when using more qubits (i.e., wider quantum NNs), because it is sufficient to encode the principal features of handwritten digits with a few qubits.

Figure 4 displays the optimized control pulses applied in a trained three-qubit model with a randomly picked input sample. The first and second 50 sampling periods correspond to, respectively, the sample-dependent encoding layers (blue) and the sample-independent inference layers (red). Most encoding control variables reach the set bound $B = 25$ MHz. This pattern is observed in almost all simulations, implying that the encoding network may be further simplified (e.g., fix the control amplitudes and vary

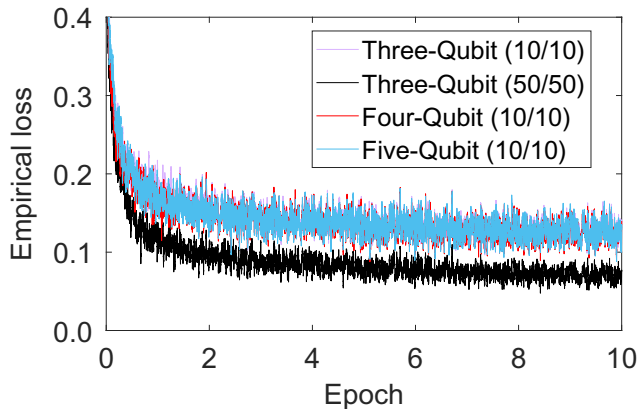


FIG. 3. The learning curves for the training of quantum end-to-end learning models. The empirical loss evaluated over each minibatch (smoothed over the past 100 iterations) during the training process for models with 3–5 qubits and various numbers of encoding and inference layers.




FIG. 4. The control pulses applied in the three-qubit end-to-end learning model with 50 encoding layers (blue color, 1–250 ns with 50 sampling periods) and 50 inference layers (red color, 251–500 ns with 50 sampling periods).

only the switching times) so as to reduce the model complexity. However, the saturated encoding control variables will lead to vanishing gradients along W variables, which may slow down the training process on a landscape plateau [33].

B. The testing results

The generalizability of the trained end-to-end learning models is tested with the validation dataset (containing 7837 independent samples belonging to the selected eight classes). Table I lists both the empirical loss evaluated on the validation dataset and the error rates evaluated by an equivalent softmax function:

$$R_{\text{error}} = P \left[\arg \max_y \langle \Psi_{x^{(j)}}(\mathbf{w}) | M_y | \Psi_{x^{(j)}}(\mathbf{w}) \rangle \neq y^{(j)} \right], \quad (7)$$

which infers the label as y for input x if the probability of producing y through the measurement is the largest. These indices are consistent with the learning curves tested on minibatches, and are close to those of classical NN models. To our knowledge, such high performance (error rate lower than 10%) is only reachable with PQC-based models either on a downsize dataset (e.g., binary classification or with downsampled images [14,21,22]) or with more (≥ 9) qubits [15,23–25].

The confusion matrix listed in Table II provides more details for the validation results on the trained three-qubit

TABLE I. The empirical losses and error rates evaluated on the validation MNIST dataset.

Model	Loss	Error rate
Three-qubit (10/10)	13.82%	8.15%
Three-qubit (50/50)	9.56%	3.30%
Four-qubit (10/10)	13.78%	8.18%
Five-qubit (10/10)	14.06%	8.80%

model with 50 encoding and 50 inference layers. It can be seen that the digit “0” has the highest precision, meaning that it is the least probable to be misclassified as other digits. The digit “4” has the highest and lowest recall rate, i.e., it is the least probable digit for other digits to be misclassified as. Among all the digits, “0” is the relatively best recognized digit by the quantum-learning model.

C. Roles of quantum and classical NN layers

We carry out additional numerical experiments to further understand the respective roles of the classical and quantum layers played in the end-to-end learning process.

We first randomly pick a W matrix and fix it, and then train the inference layers. It turns out that the error rate can hardly be under 30%. However, if we remove the inference layers and train the encoding W alone, the error rate [see Fig. 5(a)] can be easily reduced to about 10% with only ten encoding layers. The performances can be further reduced using deeper encoding networks, but gradually saturates when the number of encoding layers is over 30.

The above simulations imply that, with sufficiently deep encoding networks, the inference layers can be reduced to an identity mapping followed by measurements. To see this more clearly, we fix the trained W with 5, 10, and 50 encoding layers and train various numbers of inference layers. The results shown in Fig. 5(b) indicate that these additional inference layers do not significantly improve the performance except when there are only few encoding layers. This is consistent with the practice of classical deep learning in which feature selection is dominantly more relevant than the inference from the selected features.

Because the size of the data-to-control interface linearly increases with the number of quantum encoding layers, it is necessary to verify whether the quantum network actually takes effect in the high performance achieved in Fig. 5(a). We follow the linear baseline rule [15] to specify the contributions of the quantum and classical layers by excluding their nonlinear effects that power up machine-learning algorithms. In our model, the nonlinearity comes from the sigmoid function in the data-to-control interface and the control-to-state mapping. We first remove the sigmoid function from the data-to-control interface in the three-qubit model with 50 encoding layers and no inference layers. With the remaining control-to-state nonlinearity, the error rate is only increased from 3.87% to 6.04%. However, if we further remove the control-to-state nonlinearity by replacing the encoding layers by a linear classical mapping, which forms a $784 \times 300 \times 3$ two-layer linear network, the achievable error rate rises much higher to 25.24%. These results clearly show the role of the control-to-state nonlinearity associated with quantum encoding layers.

D. Robustness to noises

The proposed quantum end-to-end learning model can also be trained to be robust to noises that are common on NISQ devices. The online training algorithm does not need to be changed because the system also receives “samples” of the noises in addition to the samples fed from the dataset. In other words, the learning model is jointly trained by the dataset and the noises, from which the model gains simultaneously generalizability on the learning task and robustness to the noises [31].

For illustration, we add white flux noises $n_k(t)$ to the nominal three-qubit Hamiltonian $H(t)$ defined in Eq. (6) through the following Hamiltonian:

$$H_n(t) = \sum_{k=1}^n n_k(t) \sigma_z^k. \quad (8)$$

TABLE II. The confusion matrix, precision (%), and recall rates (%) for the eight digits classified by the three-qubit (50/50) quantum end-to-end learning model.

	0	2	3	4	5	6	8	9	Prec.
0	967	4	2	0	2	4	1	0	98.7
2	7	1000	9	3	0	1	10	2	96.9
3	0	11	982	0	4	0	11	2	97.2
4	0	4	0	953	0	9	1	15	97.1
5	5	0	14	0	851	8	9	5	95.4
6	4	2	0	2	14	931	5	0	97.2
8	5	8	8	6	5	4	935	3	96.0
9	5	2	15	11	3	0	14	959	95.0
Rec.	97.4	97.0	95.3	97.7	96.8	97.3	94.8	97.3	96.7

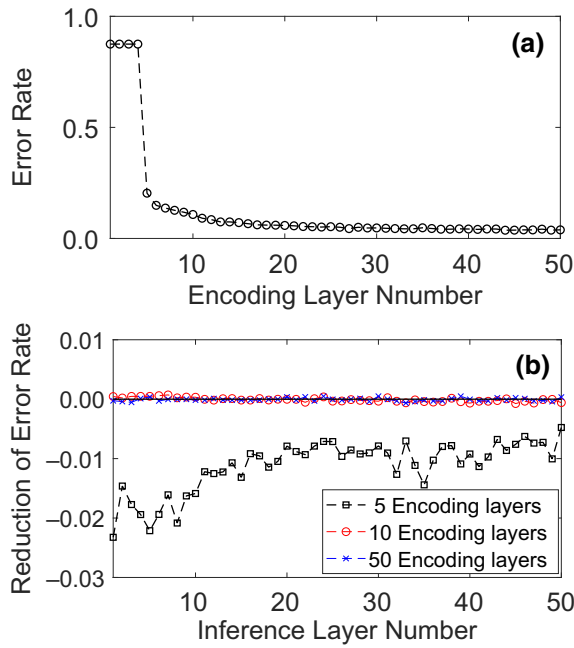


FIG. 5. The simulated performance of the quantum end-to-end learning. (a) Error rates of three-qubit models with 1–50 encoding layers and no inference layers; (b) the error-rate reduction of three-qubit models by inference layers based on pre-trained 5-layer, 10-layer, and 50-layer encoding networks in (b), whose baseline error rates are 20.38%, 10.78%, and 3.87%, respectively.

In the simulation, we train end-to-end learning models with 50 encoding and 50 inference layers at various noise levels (characterized by the variance δ). Then, we compare their accuracies with that of the model trained without noise, where the testing is done with the same MNIST

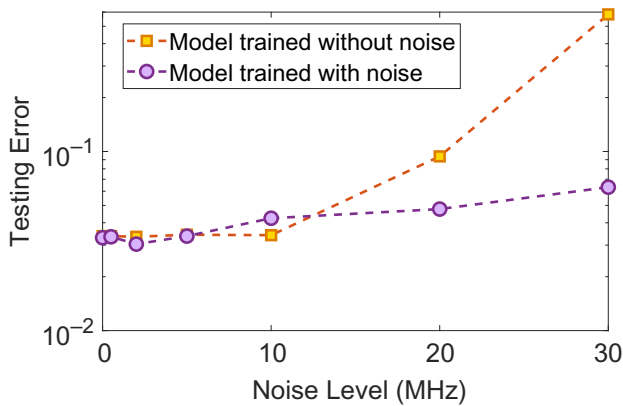


FIG. 6. The testing errors of the end-to-end learning models trained with and without flux noises, respectively, in which the testing is done at various noise levels from $\delta = 0$ MHz to $\delta = 30$ MHz. The jointing training with both MNIST samples and flux noises can keep the error rate below 10% under all simulated noise levels, while the training without noises leads to much worse performance when the noises are strong.

testing set and flux noises at the same level. As is shown in Fig. 6, the accuracies of the two models are not very different when the noises are relatively weak ($\delta < 10$ MHz). When the noises are stronger, the error of the model trained with noises only slightly increases, while the model trained without noises performs much worse and becomes totally unreliable. The comparison shows that the practically trained learning model is inherently robust to noises.

V. CONCLUDING REMARKS

To summarize, we propose a hardware-friendly quantum end-to-end learning model that can be conveniently deployed on NISQ processors. The model implements the controlled quantum dynamics as a quantum NN parameterized by experimental addressable control pulses, and the embedded data-to-control interface can automatically select appropriate features for inference. Numerical tests on the benchmarking MNIST dataset demonstrate that high performance can be attained with only a few qubits without downsizing the images even in the presence of noises in the physical systems. Taking into account the precision, the size of dataset, and the model complexity, the scheme achieves good overall performance, exhibiting great potential for real-world learning tasks.

Our proposal turns the machine-learning process into an optimal control problem, both of which can be resolved with stochastic gradient-descent algorithms. This interesting connection can be dated back to the invention of famous back-propagation (BP) algorithm for training NNs, whose was shown to be closely connected with Pontryagin’s maximum principle (PMP) in optimal control theory [34]. Recently, it was rediscovered to train deep NNs [35,36]. From the opposite side, the design of robust quantum controls [31] and quantum optimizers [32,37] can be taken as the design of a generalizable learning model. We expect to develop more efficient and noise-resilient training algorithms by unifying these two different but connected fields.

From the point of view of control theory, the capacity of the quantum end-to-end learning model can be partially understood through the controllability of the underlying control system (i.e., the ability of generating arbitrary unitary transformations), which is jointly determined by the physical qubit-qubit connectivity and the bandwidth and length of the applied control pulses. Full controllability is seemingly not required for many-qubit quantum processors, as quantum supremacy can be approached only with those transformations reachable in polynomial time [38]. However, we still suggest that the quantum hardware should be as controllable as possible, not only for larger model capacity but also for efficient search for high-performance learning models, because the underlying control landscape (for training process) tends to be free of traps [39,40].

Finally, we indicate that the proposed learning scheme can be easily extended to NISQ processors containing more qubits and other components (e.g., cavity modes or multilevel atoms), and the framework is also transplantable to any other learning tasks. In the simulation examples on the relatively simple MNIST dataset, the full power of the quantum end-to-end learning has not been fully released, and we expect to explore its potential power on more complicated learning tasks in future studies.

ACKNOWLEDGMENTS

This work is supported by the National Key R&D Program of China (Grants No. 2017YFA0304304 and No. 2018YFA0306703), NSFC (Grants No. 61833010 and No. 61773232), the Key-Area R&D Program of Guangdong Province (Grant No. 2018B030326001) and a Grant from the Institute for Guo Qiang, Tsinghua University. Invaluable discussions with Professor Changshui Zhang are greatly appreciated.

-
- [1] John Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
- [2] Frank Arute *et al.*, Quantum supremacy using a programmable superconducting processor, *Nature* **574**, 505 (2019).
- [3] Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd, Quantum machine learning, *Nature* **549**, 195 (2017).
- [4] Vedran Dunjko and Hans J. Briegel, Machine learning & artificial intelligence in the quantum domain: A review of recent progress, *Rep. Prog. Phys.* **81**, 074001 (2018).
- [5] Patrick Rebentrost, Masoud Mohseni, and Seth Lloyd, Quantum Support Vector Machine for Big Data Classification, *Phys. Rev. Lett.* **113**, 130503 (2014).
- [6] Seth Lloyd, Masoud Mohseni, and Patrick Rebentrost, Quantum principal component analysis, *Nat. Phys.* **10**, 631 (2014).
- [7] Seth Lloyd and Christian Weedbrook, Quantum Generative Adversarial Learning, *Phys. Rev. Lett.* **121**, 040502 (2018).
- [8] Ling Hu, Shu-Hao Wu, Weizhou Cai, Yuwei Ma, Xianghao Mu, Yuan Xu, Haiyan Wang, Yipu Song, Dong-Ling Deng, Chang-Ling Zou, and Luyan Sun, Quantum generative adversarial learning in a superconducting quantum circuit, *Sci. Adv.* **5**, eaav2761 (2019).
- [9] D. Zhu, N. M. Linke, M. Benedetti, K. A. Landsman, N. H. Nguyen, C. H. Alderete, A. Perdomo-Ortiz, N. Korda, A. Garfoot, C. Brecque, L. Egan, O. Perdomo, and C. Monroe, Training of quantum circuits on a hybrid quantum computer, *Sci. Adv.* **5**, eaaw9918 (2019).
- [10] Marcello Benedetti, Erika Lloyd, Stefan Sack, and Mattia Fiorentini, Parameterized quantum circuits as machine learning models, *arXiv:1906.07682* (2019).
- [11] Seth Lloyd, Quantum approximate optimization is computationally universal, *arXiv:1812.04877* (2018).
- [12] Abhinav Kandala, Antonio Mezzacapo, Kristan Temme, Maika Takita, Markus Brink, Jerry M. Chow, and Jay M. Gambetta, Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets, *Nature* **549**, 242 (2017).
- [13] E. Farhi, J. Goldstone, S. Gutmann, and H. Neven, Quantum algorithms for fixed qubit architectures, *arXiv:1703.06199* (2017).
- [14] Edward Farhi and Hartmut Neven, Classification with quantum neural networks on near term processors, *arXiv:1802.06002* (2018).
- [15] C. M. Wilson, J. S. Otterbach, N. Tezak, R. S. Smith, A. M. Pollreno, Peter J. Karalekas, S. Heidel, M. Sohaib Alam, G. E. Crooks, and M. P. da Silva, Quantum kitchen sinks: An algorithm for machine learning on near-term quantum computers, *arXiv:1806.08321* (2018).
- [16] Seth Lloyd, Maria Schuld, Aroosa Ijaz, Josh Izaac, and Nathan Killoran, Quantum embeddings for machine learning, *arXiv:1812.04877* (2020).
- [17] Vojtěch Havlíček, Antonio D. Córcoles, Kristan Temme, Aram W. Harrow, Abhinav Kandala, Jerry M. Chow, and Jay M. Gambetta, Supervised learning with quantum-enhanced feature spaces, *Nature* **567**, 209 (2019).
- [18] J. S. Otterbach *et al.*, Unsupervised machine learning on a hybrid quantum computer, *arXiv:1712.05771* (2017).
- [19] Kathleen E. Hamilton, Eugene F. Dumitrescu, and Raphael C. Pooser, Generative model benchmarks for superconducting qubits, *Phys. Rev. A* **99**, 062323 (2019).
- [20] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, Gradient-based learning applied to document recognition, *Proc. IEEE* **86**, 2278 (1998).
- [21] Edward Grant, Marcello Benedetti, Shuxiang Cao, Andrew Hallam, Joshua Lockhart, Vid Stojevic, Andrew G. Green, and Simone Severini, Hierarchical quantum classifiers, *npj Quantum Inf.* **4**, 65 (2018).
- [22] William Huggins, Piyush Patil, Bradley Mitchell, K. Birgitta Whaley, and E. Miles Stoudenmire, Towards quantum machine learning with tensor networks, *Quantum Sci. Technol.* **4**, 024001 (2019).
- [23] Iordanis Kerenidis and Alessandro Luongo, Quantum classification of the mnist dataset via slow feature analysis, *arXiv preprint: 1805.08837* (2018).
- [24] Ryan Sweke, Frederik Wilde, Johannes Meyer, Maria Schuld, Paul K. Fährmann, Barthélémy Meynard-Piganeau, and Jens Eisert, Stochastic gradient descent for hybrid quantum-classical optimization, *arXiv:1910.01155* (2019).
- [25] Maxwell Henderson, Samridhi Shakya, Shashindra Pradhan, and Tristan Cook, Quantvolutional neural networks: Powering image recognition with quantum circuits, *Quantum Machine Intelligence* **2**, 1 (2020).
- [26] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton, Deep learning, *Nature* **521**, 436 (2015).
- [27] M. W. Johnson *et al.*, Quantum annealing with manufactured spins, *Nature* **473**, 194 (2011).
- [28] R. Barends, J. Kelly, A. Megrant, D. Sank, E. Jeffrey, Y. Chen, Y. Yin, B. Chiaro, J. Mutus, C. Neill, P. O'Malley, P. Roushan, J. Wenner, T. C. White, A. N. Cleland, and John M. Martinis, Coherent Josephson Qubit Suitable for Scalable Quantum Integrated Circuits, *Phys. Rev. Lett.* **111**, 080502 (2013).
- [29] Xiu Gu, Anton Frisk Kockum, Adam Miranowicz, Yu-xi Liu, and Franco Nori, Microwave photonics with superconducting quantum circuits, *Phys. Rep.* **718**, 1 (2017).

- [30] S. J. Glaser, U. Boscain, T. Calarco, C. P. Koch, W. Kockenberger, R. Kosloff, I. Kuprov, B. Luy, S. Schirmer, T. Schulte-Herbruggen, D. Sugny, and F. K. Wilhelm, Training schrodinger's cat: Quantum optimal control strategic report on current status, visions and goals for research in europe, *Eur. Phys. J. D* **69**, 279 (2015).
- [31] Re-Bing Wu, Haijin Ding, Daoyi Dong, and Xiaoting Wang, Learning robust and high-precision quantum controls, *Phys. Rev. A* **99**, 042327 (2019).
- [32] Yulong Dong, Xiang Meng, Lin Lin, Robert Kosut, and K. Birgitta Whaley, Robust control optimization for quantum approximate optimization algorithm, [arXiv:1911.00789](https://arxiv.org/abs/1911.00789) (2019).
- [33] Léon Bottou, Frank E. Curtis, and Jorge Nocedal, Optimization methods for large-scale machine learning, *SIAM Rev.* **60**, 223 (2018).
- [34] Y. LeCun, in *Proceedings of the 1988 Connectionist Models Summer School*, edited by D. Touretzky, G. Hinton, and T. Sejnowski (Morgan Kaufmann, Pittsburg, PA, 1988), p. 21.
- [35] Qianxiao Li, Long Chen, Cheng Tai, and E. Weinan, Maximum principle based algorithms for deep learning, *J. Machine Learning Res.* **18**, 5998 (2017).
- [36] Qianxiao Li and Shuji Hao, in *Proceedings of the 35th ICML*, edited by Jennifer Dy and Andreas Krause, *Proceedings of Machine Learning Research* Vol. 80 (Stockholm Sweden, PMLR, 10–15 July 2018), p. 2985.
- [37] Aditya Dendukuri, Blake Keeling, Arash Fereidouni, Joshua Burbridge, Khoa Luu, and Hugh Churchill, Defining quantum neural networks via quantum time evolution, [arXiv:1905.10912](https://arxiv.org/abs/1905.10912) (2019).
- [38] Christian Arenz and Herschel Rabitz, Controlling Qubit Networks in Polynomial Time, *Phys. Rev. Lett.* **120**, 220503 (2018).
- [39] H. A. Rabitz, M. M. Hsieh, and C. M. Rosenthal, Quantum optimally controlled transition landscapes, *Science* **303**, 1998 (2004).
- [40] Re-Bing Wu, Qiuyang Sun, Tak-san Ho, and Herschel Rabitz, Inherently trap-free convex landscapes for fully quantum optimal control, *J. Math. Chem.* **57**, 2154 (2019).