

Received June 17, 2020, accepted July 6, 2020, date of publication July 20, 2020, date of current version August 12, 2020.

Digital Object Identifier 10.1109/ACCESS.2020.3010470

# Variational Quantum Circuits for Deep Reinforcement Learning

SAMUEL YEN-CHI CHEN<sup>1,2</sup>, CHAO-HAN HUCK YANG<sup>3</sup>, JUN QI<sup>3</sup>, (Member, IEEE), PIN-YU CHEN<sup>4</sup>, (Member, IEEE), XIAOLI MA<sup>3</sup>, (Fellow, IEEE), AND HSI-SHENG GOAN<sup>1,2,5</sup>

<sup>1</sup>Department of Physics, National Taiwan University, Taipei City 10617, Taiwan

<sup>2</sup>Center for Theoretical Physics, National Taiwan University, Taipei City 10617, Taiwan

<sup>3</sup>School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA

<sup>4</sup>IBM Thomas J. Watson Research Center, Yorktown Heights, NY 10598, USA

<sup>5</sup>Center for Quantum Science and Engineering, National Taiwan University, Taipei City 10617, Taiwan

Corresponding authors: Samuel Yen-Chi Chen (ycchen1989@gmail.com) and Hsi-Sheng Goan (goan@phys.ntu.edu.tw)

This work was supported in part by the Ministry of Science and Technology (MOST) of Taiwan under Grant MOST 106-2112-M-002-013-MY3, Grant MOST 108-2627-E-002-001, Grant MOST 108-2622-8-002-016, and Grant MOST 107-2627-E-002-001-MY3; and in part by the National Taiwan University (NTU) under Grant NTU-CC-108L893202 and Grant NTU-CC-109L892002, and by the thematic group program of the National Center for Theoretical Sciences, Taiwan. The use of the quantum computer machines of the International Business Machines Corporation (IBM) in this work was supported through the IBM Q Hub at NTU under Grant MOST 107-2627-E-002-001-MY3.

**ABSTRACT** The state-of-the-art machine learning approaches are based on classical von Neumann computing architectures and have been widely used in many industrial and academic domains. With the recent development of quantum computing, researchers and tech-giants have attempted new quantum circuits for machine learning tasks. However, the existing quantum computing platforms are hard to simulate classical deep learning models or problems because of the intractability of deep quantum circuits. Thus, it is necessary to design feasible quantum algorithms for quantum machine learning for noisy intermediate scale quantum (NISQ) devices. This work explores variational quantum circuits for deep reinforcement learning. Specifically, we reshape classical deep reinforcement learning algorithms like experience replay and target network into a representation of variational quantum circuits. Moreover, we use a quantum information encoding scheme to reduce the number of model parameters compared to classical neural networks. To the best of our knowledge, this work is the first proof-of-principle demonstration of variational quantum circuits to approximate the deep  $Q$ -value function for decision-making and policy-selection reinforcement learning with experience replay and target network. Besides, our variational quantum circuits can be deployed in many near-term NISQ machines.

**INDEX TERMS** Communication network, deep reinforcement learning, quantum machine learning, quantum information processing, variational quantum circuits, noisy intermediate scale quantum, quantum computing.

## I. INTRODUCTION

Deep Learning (DL) [1] has been widely used in many machine learning domains, such as computer vision [2]–[4], natural language processing [5], communication network congestion control [6], and mastering the game of Go [7]. The successful deployment of DL is primarily attributed to the improvement of new computer architectures associated with powerful computing capabilities in the past decades.

The associate editor coordinating the review of this manuscript and approving it for publication was Minh Jo.

Many researchers also utilized DL-based data analysis methods on fundamental physics researches such as quantum many-body physics [8]–[10], phase-transitions [11], quantum control [12], [13], and quantum error correction [14], [15]. In the meantime, great efforts from both the physics and machine learning community have dedicated to and empowered quantum computation. Quantum computing machines have been brought to the market (e.g., IBM's and D-Wave's hardware solutions [16], [17]), but a large-scale quantum circuits cannot be faithfully employed upon the quantum computing platforms due to the lack of quantum error

correction [18], [19]. Therefore, Mitarai *et al.* design approximate quantum algorithms, circuits and encoding schemes [20] on the devices with noise tolerance. More specifically, the work takes the advantages of quantum entanglement [20], [21] in quantum computing to reduce the model size into an essentially small number and take advantage of the iterative optimization to reduce the quantum circuit depth to a practically low value such that hybrid quantum-classical algorithms can be realized on the available quantum platforms which are named as noisy intermediate-scale quantum (NISQ) machines [21], [22].

By taking the strengths of quantum computing with significantly fewer parameters [21], variational quantum circuits on NISQ have succeeded in implementing standard classification and clustering algorithms on classical benchmark datasets [20], [23], [24]. Besides, it is also possible to employ quantum circuits for implementing new DL algorithms like generative adversarial networks [25] (GAN) on NISQ machines. These frameworks and development pave the way towards applications of near-term quantum devices for quantum machine learning. However, to the best of our knowledge, variational circuits on current NISQ computing for deep neural network based decision making and policy selection problems have not been discussed, which constrains the application of NISQ in many machine learning scenarios with sequential decision making.

Since reinforcement learning (RL) and deep reinforcement learning (DRL) are two paradigms of complex sequential decision-making systems and satisfy the requirements of automatic policy learning under uncertainty, our work focuses on the empowerment of DRL on NISQ computation, which refers to an agent interacting with the environment to gain knowledge of backgrounds and deriving the policy of decision making accordingly [26], [27]. We propose a novel variational quantum circuit feasible on the current NISQ platform hybridized with iterative parameter optimization on a classical computer to resolve the circuit-depth challenges. Furthermore, we generalize variational quantum circuits to standard DRL based *action-value* function approximation [27], [28]. Finally, we analyze the policy reward and the memory cost for performance of our variational quantum circuits (VQC) based DRL in comparison with standard RL and DRL approaches in the context of *frozen-lake* [29] and *cognitive-radio* [30] environments. The frozen lake is a simple maze environment in openAI Gym [29] and is a typical and simple example that is demonstrated in standard RL. Cognitive radio is a wireless technology that enables for optimizing the use of available communication channels between users and has been studied by the standard machine learning technique [6], [30]. Under current limitations on the scale of quantum machines and the capabilities of quantum simulations, we select the frozen-lake and cognitive-radio environments for the proof-of-principle quantum machine learning study. To the best of our knowledge, this work is the first demonstration of variational quantum circuits to the DRL-based decision-making and policy-selection problems.

## II. REINFORCEMENT LEARNING

*Reinforcement learning* is a machine learning paradigm in which an *agent* interacts with an *environment*  $\mathcal{E}$  over a number of discrete time steps [26]. At each time step  $t$ , the agent receives a *state* or *observation*  $s_t$  and then chooses an *action*  $a_t$  from a set of possible actions  $\mathcal{A}$  according to its *policy*  $\pi$ . The policy is a function mapping the state  $s_t$  to action  $a_t$ . In general, the policy can be stochastic, which means that given a state  $s$ , the action output can be a probability distribution. After executing the action  $a_t$ , the agent receives the state of the next time step  $s_{t+1}$  and a scalar *reward*  $r_t$ . The process continues until the agent reaches the terminal state. An *episode* is defined as an agent starting from a randomly selected initial state and following the aforementioned process all the way through the terminal state.

Define  $R_t = \sum_{t'=t}^T \gamma^{t'-t} r_{t'}$  as the total discounted return from time step  $t$ , where  $\gamma$  is the discount factor that lies in  $(0, 1]$ . In principle,  $\gamma$  is provided by the investigator to control how future rewards are given to the decision making function. When a large  $\gamma$  is considered, the agent takes into account future rewards no matter what a discount rate is. As to a small  $\gamma$ , an agent can quickly ignore future rewards within a few time steps. The goal of the agent is to maximize the expected return from each state  $s_t$  in the training process. The *action-value function* or *Q-value function*  $Q^\pi(s, a) = \mathbb{E}[R_t | s_t = s, a]$  is the expected return for selecting an action  $a$  in state  $s$  based on policy  $\pi$ . The optimal action value function  $Q^*(s, a) = \max_\pi Q^\pi(s, a)$  gives a maximal action-value across all possible policies. The value of state  $s$  under policy  $\pi$ ,  $V^\pi(s) = \mathbb{E}[R_t | s_t = s]$ , is the agent's expected return by following policy  $\pi$  from the state  $s$ . The classical temporal difference (TD) error [26] is used to update value function in reinforcement learning tasks.

### A. Q-LEARNING

Q-learning [26] is a model-free RL algorithm. Before the learning process begins,  $Q$  is initially assigned to an arbitrary fixed value (chosen by the programmer). Then, at each time, the agent selects an action  $a_t$  (using, e.g.,  $\epsilon$ -greedy policy derived from  $Q$ ), observes a reward  $r_t$ , and enters a new state  $s_{t+1}$  (that may depend on both the previous state  $s_t$  and the selected action), and then  $Q$  is updated with the learning rate  $\alpha$ . The Q-learning is an *off-policy* learner since it updates its Q-value using the observed reward  $r_t$  and the maximum reward  $\max_a Q(s_{t+1}, a)$  for the next state  $s_{t+1}$  over all possible actions  $a$ . The updating is done according to the benchmark formula:<sup>1</sup>

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha \left[ r_t + \gamma \max_a Q(s_{t+1}, a) - Q(s_t, a_t) \right]. \quad (1)$$

### B. STATE-ACTION-REWARD-STATE-ACTION (SARSA)

An SARSA [26] agent interacts with the environment and updates the policy based on the undertaking actions. The Q

<sup>1</sup>The formula and loss are from the original DQN work Mnih et al. [28].

value represents the possible reward received in the next time step for taking action  $a_t$  in state  $s_t$ , plus the discounted future reward received from the next state-action observation, and is updated by temporal difference with transitions from state-action pair  $(s_t, a_t)$  to state-action pair  $(s_{t+1}, a_{t+1})$ , adjusted by the learning rate  $\alpha$  as:<sup>2</sup>

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha [r_t + \gamma Q(s_{t+1}, a_{t+1}) - Q(s_t, a_t)]. \quad (2)$$

### C. DEEP Q-LEARNING

The action-value function  $Q(s, a)$  can be explicitly represented by a two-dimensional table with a total number of entries  $s \times a$ , that is, the number of possible states times the number of possible actions. However, when the state space or the action space is large or even continuous the tabular method is unfeasible. In such a situation, the action-value function is represented with function approximators such as neural networks [27], [28]. This neural-networks-based reinforcement learning is called *deep reinforcement learning* (DRL).


The employment of neural networks for function approximators to represent the  $Q$ -value function has been studied extensively [27], [28] and succeeded in many tasks like playing video games. In this setting, the action-value function  $Q(s, a; \theta)$  is parameterized by  $\theta$ , which can be derived by a series of iterations from a variety of optimization methods adopted from other machine learning tasks. The simplest form is the  $Q$ -learning. In this method, the goal is to directly approximate the optimal action-value function  $Q^*(s, a)$  by minimizing the mean square error (MSE) loss function:

$$L(\theta) = \mathbb{E}[(r_t + \gamma \max_{a'} Q(s_{t+1}, a'; \theta^-) - Q(s_t, a_t; \theta))^2]. \quad (3)$$

Here, the prediction is  $Q(s_t, a_t; \theta)$ , where  $\theta$  is the parameter of the policy network, and the target is  $r_t + \gamma \max_{a'} Q(s_{t+1}, a'; \theta^-)$ , where  $\theta^-$  is the parameter of the target network and  $s_{t+1}$  is the state encountered after playing action  $a_t$  at state  $s_t$ . The loss function in DRL is normally hard to converge and is likely to get divergent when a nonlinear approximator like a neural network is used to represent the action-value function [28]. There are several possible culprits. When the states or observations are serially correlated with each other along the trajectory, thereby violating the assumption that the sample needs to be independent and identically distributed (IID), the  $Q$  function changes dramatically and changes the policy at a relatively large scale. In addition, the correlation between the action-value  $Q$  and the target values  $r_t + \gamma \max_{a'} Q(s_{t+1}, a')$  can be large. Unlike the supervised learning where the targets are given and invariant, the setting of DRL allows targets to vary with  $Q(s, a)$ , causing  $Q(s, a)$  to chase a nonstationary target.

The *deep Q-learning* (DQL) or *deep Q-network* (DQN) presented in the work [28] addressed these issues through two mechanisms:

<sup>2</sup>We follow the classical SARSA definition from Sutton et. al. [26].



**FIGURE 1.** Frozen-Lake environment for the variational quantum DRL agent. In this frozen-lake environment, the RL agent is expected to go from the start location (S) to the goal location (G). There are several holes (H) on the way, and the agent should learn to avoid stepping into these hole locations. Furthermore, we set a negative reward for each step the agent takes. The agent is expected to learn the policy that going from S to G with the shortest path possible. In this work, we train the agents on three configurations of the frozen-lake environment shown in (a), (b) and (c) separately.

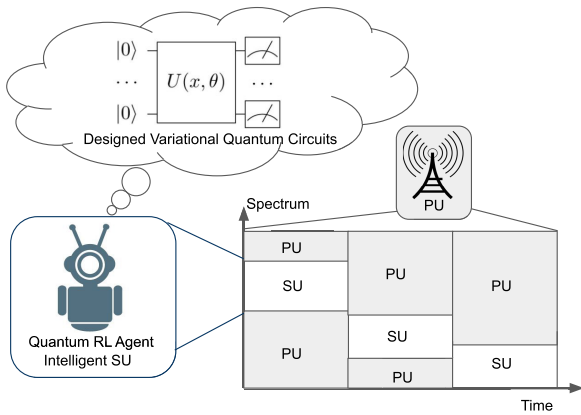
- *Experience replay*: To perform experience replay, one stores each transition the agent encounters. The transition is stored as a tuple in the following form:  $(s_t, a_t, r_t, s_{t+1})$  at each time step  $t$ . To update the  $Q$ -learning parameters, one randomly samples a batch of experiences from the replay memory and then performs gradient descent with the following MSE loss function:  $L(\theta) = \mathbb{E}[(r_t + \gamma \max_{a'} Q(s_{t+1}, a'; \theta^-) - Q(s_t, a_t; \theta))^2]$ , where the loss function is calculated over the batch sampled from the replay memory. The key importance of experience replay is to lower the correlation of inputs for training the  $Q$ -function.
- *Target Network*:  $\theta^-$  is the parameter of the target network and these parameters are only updated at every finite time steps. This setting helps to stabilize the  $Q$ -value function training since the target is relatively stationary compared to the action-value function.

### III. TESTING ENVIRONMENTS

To study the performance of a reinforcement learning agent, we need to specify the *environment* for the test. We will consider the *frozen-lake* [29] and *cognitive-radio* [30] environments. The reason why we choose the frozen-lake environment is two-fold. First is that it is a fairly simple and commonly tested example in standard RL, and if the dimension of the problem size is not too large, the simulation is feasible with available quantum simulators and NISQ devices, and the time consumption for the experiment is reasonable too. The second is that we want to demonstrate that quantum circuits are capable of learning sequential decision making process (also called *policy*). The choice of cognitive-radio environment is that we want to demonstrate some kinds of real-world applications and the complexity of this environment is comparable to the frozen-lake environment.

#### A. FROZEN LAKE

The first testing environment we consider in this work is the frozen lake, a simple maze environment in openAI Gym [29].



**FIGURE 2.** Cognitive-Radio environment for the variational quantum DRL agent. In the cognitive-radio environment, the agent is expected to select a channel that is free of interference in each time step. For example, there is a *primary user* (PU) that will occupy a specific channel in each time step periodically. Our agent, which is the *secondary user* (SU), can only select the channels that are not occupied without the knowledge of the PU in advance. The agent is expected to learn the policy through the interaction with the environment.


In this environment, the agent standing on a frozen lake is expected to go from the start location (S) to the goal location (G) (see Fig. 1). Since the lake is not all frozen, there are several holes (H’s) on the way, and the agent should learn to avoid stepping into these hole locations, otherwise the agent will get a large negative reward and the episode will terminate. Furthermore, the agent is also expected to take the shortest possible path. In order to accomplish this, we set a little negative reward on each move. Here we demonstrate three different configurations of the frozen-lake environment, as shown in Fig. 1, for the training.

**The frozen-lake environment mapping is:**


- Observation: observed records of all time steps.
- Action: there are four actions LEFT, DOWN, RIGHT, UP in the action space. How to choose the action in a variational quantum circuit will be described in Sec. VI-A.
- Reward: The rewards in this environment are +1.0 for successfully achieving the goal, -0.2 for failing the task, which is stepping into one of the holes. Moreover, to encourage the agent to take the shortest path, there is also a -0.01 reward for each step taken.

**B. COGNITIVE RADIO**

In the second testing environment we study the proposed variational quantum-DQN or -DQL (VQ-DQN; VQ-DQL) agent in a real-world application. We consider the cognitive-radio experiment. In this setting, the agent is expected to select a channel that is not occupied or interfered by a primary user (see Fig. 2). If the agent succeeds, then it will get +1 reward, otherwise it will get -1 reward. Note that the episode will terminate if the agent collects three failed selections or the agent plays more than 100 steps. This task is crucial for the modern wireless multi-channel environment



(a)



(b)



(c)

**FIGURE 3.** Cognitive-Radio environment with periodical channel-changing pattern for the variational quantum DRL agent. We provide three configurations for the cognitive-radio experiment; the first setting (a) is the main configuration for experiments on a different number of channels and experiments in noisy situations. The other two configurations in (b) and (c) are only tested in the case of 4 channels, the purpose of these additional experiments is to demonstrate that the proposed framework is generally applicable in different scenarios.

since channels are possibly occupied or under interference. For the demonstration in this work, we consider that there are *n* possible channels for the agent to select and the channel-changing by the primary user follows a simple periodic pattern with *n* time-steps in a full cycle. Three different configurations of the cognitive-radio environment in the case of four channels for the training considered here are illustrated in Fig. 3.

**The cognitive-radio environment mapping is:**

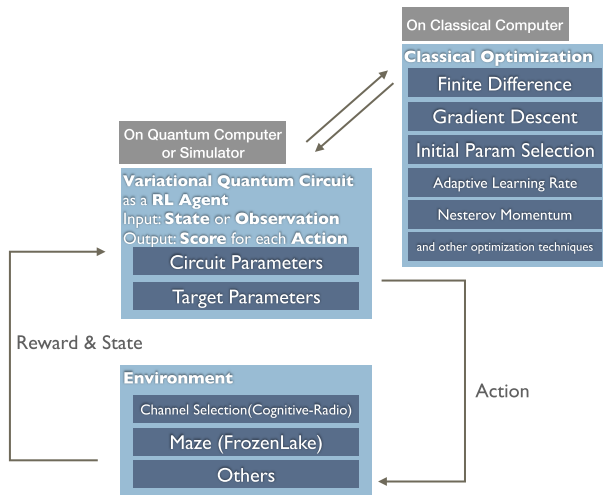
- Observation: *ns*3 [30] statistics with the radio channels capacity, with a customized channel number = *n*. (e.g., a state of [1 0 0 0] represents for *n* = 4 channels and a primary user on the 1st channel.)
- Action: selecting one channel for the secondary user accessing a radio channel out of *n* channels. How to choose the action in a variational quantum circuit for the cognitive-radio scenario will be described in Sec. VI-A.
- Reward: -1 for a collision with the primary user; +1 for no collision. The list of score presenting rewards in the testing environments is shown in Table 1. Agent can achieve a maximum score of 100.



**TABLE 1.** List of rewards in our frozen-lake and cognitive-radio testing environments. In the frozen lake, the environment is non-slippery. This setting encourages the agent not only to achieve the goal but also to select the shortest path.

(a) Frozen-Lake		(b) Cognitive-Radio	
Location	Reward	Location	Reward
HOLE	-0.2	Occupied Channel	-1.0
GOAL	+1.0	Available Channel	+1.0
OTHER	-0.01		

### C. VARIATIONAL QUANTUM DEEP Q-LEARNING



**FIGURE 4.** Overview of variational quantum circuits for DRL. In this work, we study the capability of variational quantum circuits in performing DRL tasks. This DRL agent includes a quantum part and a classical part. Under current limitations on the scale of quantum machines and the capabilities of quantum simulations, we select frozen-lake and cognitive-radio environments for the proof-of-principle study. The proposed framework is rather general and is expected to solve complicated tasks when larger-scale quantum machines are available.


### IV. VARIATIONAL QUANTUM CIRCUITS AND DEEP Q-LEARNING

The variational quantum circuit is a hybrid quantum-classical approach which leverages the strengths of quantum and classical computation. It is one type of quantum circuits with tunable parameters which are optimized in an iterative manner by a classical computer. These parameters can be seen as the *weights* in artificial neural networks. The variational quantum circuit approach has been shown to be flexible in circuit depth and somewhat resistant to noise [31]–[33]. Therefore, even though there is still lack of quantum error correction and fault-tolerant quantum computation in the NISQ devices, the quantum machine learning algorithms powered by variational quantum circuits can circumvent the complex quantum errors which exist in the available quantum devices. Previous results in [20], [23], [34] have demonstrated that the variational quantum circuits can model any function approximators, classifiers and even quantum-many-body physics that are intractable on classical computers. For example, the work in [20] shows that a variational quantum circuits can approximate an analytical function  $f(x)$ .

It is hard to simulate quantum circuits of a large number of qubits via classical computers. For example, a quantum circuit with 100 qubits corresponds to a computational state space of dimensions  $2^{100}$ . This huge number of computational state-space dimensions exceeds the storage and thus the computational capability of classical computers. Recently, Google demonstrated that a 53-qubit quantum computer can successfully sample, by quantum measurement, one instance of the probability distribution of a quantum circuit a million times in around 200 seconds while it is estimated that such calculation to generate such amount of large size of random numbers will take 10000 years on state-of-the-art classical supercomputer [35]. In this study, we consider a small-scale simulation to demonstrate the possibility of running DRL applications on a quantum computer. The scale we consider here (several qubits) is still simulable by a quantum simulator on a classical computer. The hope is that when a larger scale, for example, 100-qubit quantum computer is available, we may be, with some small changes to the variational quantum circuits, able to implement a VQ-DQL or VQ-QDN agent that is impossible to be simulated on a classical computer. For a review on the advantages of quantum computers over classical computers, see [36]. In this work, we attempt to expand the expressive power of variational quantum circuits for the action-value function of DRL. In certain cases, the variational quantum circuits require fewer parameters than a conventional neural network [21], making them promising for modeling complex environments. Consider a physical system with a size of 100 qubits; it is basically impossible to simulate this system on any currently available supercomputer, or this system requires a significant amount of classical computing resources beyond what is currently available to simulate. Therefore, the expressive power we consider here is that potentially some applications of a very large size may be represented either by a quantum circuit of an intermediate size of qubits or in theory by a classical neural network while the quantum circuit would require a fewer number of parameters than the classical neural network.

In Fig. 4, we present the overview of our proposed variational quantum circuit based DRL and its relevant components. The RL agent includes a quantum part and a classical part. The quantum part of the variational quantum circuit takes two sets of parameters and outputs measurement results that determine possible actions to take. The classical part of a classical computer performs the optimization procedure and calculates what new sets of parameters should be. Fig. 5 shows a generic quantum circuit architecture for DRL (the detailed description of the quantum circuit will be presented later), and the algorithm for the VQ-DQL or VQ-QDN is presented in Algorithm 1. We construct two sets of circuit parameters with the same circuit architecture. The main circuit parameters are updated every step, while the target circuit parameters are updated per 20 steps. For *experience replay*, the replay memory is set for the length of 80 to adapt to the frozen-lake testing environment and the length of 1000 for the cognitive-radio testing environment, and the size of

training batch is 5 for all of the environments. The process of optimization needs to calculate gradients of expectation values of quantum measurements, which can be conducted by the same circuit architecture and slightly different parameters, respectively [37]. Further, we encode the state with *computational basis encoding*. In the frozen-lake environment [29] we consider, there are totally 16 states. Thus, it requires 4 qubits to represent all states (see Fig. 5). In the cognitive-radio experiments, we apply similar method and circuit architectures with different number of qubits to match the number of possible channels (see Fig. 6). Besides, [38] provides a general discussion about the different encoding schemes. We discuss next the concept of computational basis encoding and the quantum circuits for the frozen-lake and cognitive-radio problems.



**FIGURE 5.** Generic variational quantum circuit architecture for the deep Q network (VQ-DQN). The single-qubit gates  $R_x(\theta)$  and  $R_z(\phi)$  represent rotations along x-axis and z-axis by the given angle  $\theta$ , respectively. The CNOT gates are used to entangle quantum states from each qubit and  $R(\alpha, \beta, \gamma)$  represents the general single qubit unitary gate with three parameters. The parameters labeled  $\theta_i$  and  $\phi_i$  are for state preparation and are not subject to iterative optimization. Parameters labeled  $\alpha_i, \beta_i$  and  $\gamma_i$  are the ones for iterative optimization. Note that the number of qubits can be adjusted to fit the problem of interest and the grouped box may repeat several times to increase the number of parameters, subject to the capacity and capability of the quantum machines used for the experiments.

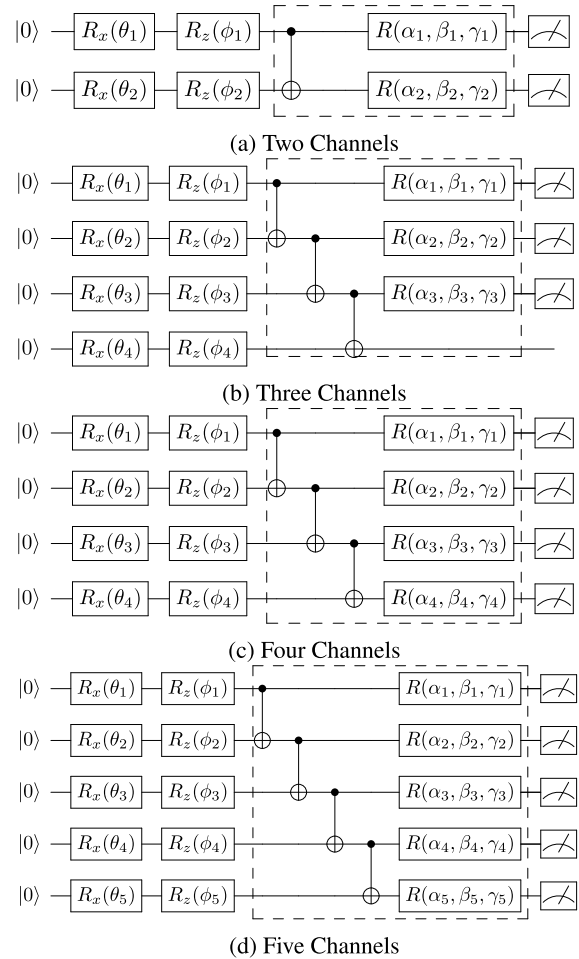
It has been shown that artificial neural networks (ANN) are *universal approximators* [39], meaning that in theory, a single hidden layer neural network can approximate any computable function. However, the number of neurons in this hidden layer may be very large, which means that this model contains so many parameters. In machine learning applications, in addition to the capability of approximation, one needs to consider the amount of resources the model consumes.

### A. COMPUTATIONAL BASIS ENCODING AND QUANTUM CIRCUIT FOR THE FROZEN LAKE PROBLEM

A general  $n$ -qubit state can be represented as:

$$|\psi\rangle = \sum_{(q_1, q_2, \dots, q_n) \in \{0,1\}^n} c_{q_1, \dots, q_n} |q_1\rangle \otimes |q_2\rangle \otimes |q_3\rangle \otimes \dots \otimes |q_n\rangle, \quad (4)$$

where  $c_{q_1, \dots, q_n} \in \mathbb{C}$  is the *amplitude* of each quantum state and each  $q_n \in \{0, 1\}$ . The square of the amplitude  $c_{q_1, \dots, q_n}$  is the *probability* of measurement with the post-measurement state in  $|q_1\rangle \otimes |q_2\rangle \otimes |q_3\rangle \otimes \dots \otimes |q_n\rangle$ , and the total probability



**FIGURE 6.** Variational quantum circuits for the cognitive-radio experiments. The basic architecture is the same as the circuit used in the frozen-lake experiment. The main difference is that we have different number of qubits in order to fit the number of channels. In cognitive-radio experiments, all circuits are with two layers (grouped box repeated twice), regardless of the number of channels. In the 3-channel case, since the number of possible state is  $3^2$ , which is greater than the number of computational basis of a 3-qubit system (which is  $2^3 = 8$ ). We use a 4-qubit system to accommodate the all possible states while the number of parameters is not increased.

should sum to 1, i.e.,

$$\sum_{(q_1, q_2, \dots, q_n) \in \{0,1\}^n} \|c_{q_1, \dots, q_n}\|^2 = 1. \quad (5)$$

We discuss the procedure, adopted here, to encode classical states to the quantum register of the variational quantum circuit. Let us take the frozen-lake environment [29] shown in Fig. 1 as an example. There are  $4 \times 4 = 16$  possible states and we label each possible state with an integer in the range from 0 to 15. We need a 4-qubit system to include all possibilities of 16 states. The *decimal* number is first converted into a *binary* number and then encoded into the quantum states through single-qubit unitary rotations. In other words, each classical state can be denoted by a four-digit binary number  $b_1b_2b_3b_4$ , where  $b_1, b_2, b_3, b_4$  can only take the value of 0 or 1, and then the corresponding encoded quantum state will be  $|b_1\rangle \otimes |b_2\rangle \otimes |b_3\rangle \otimes |b_4\rangle$ .

**Algorithm 1** Variational Quantum Deep Q Learning

---

```

Initialize replay memory  $\mathcal{D}$  to capacity  $N$ 
Initialize action-value function quantum circuit  $Q$  with random parameters
for episode = 1, 2, ...,  $M$  do
  Initialise state  $s_1$  and encode into the quantum state
  for  $t = 1, 2, \dots, T$  do
    With probability  $\epsilon$  select a random action  $a_t$ 
    otherwise select  $a_t = \max_a Q^*(s_t, a; \theta)$  from the output of the quantum circuit
    Execute action  $a_t$  in emulator and observe reward  $r_t$  and next state  $s_{t+1}$ 
    Store transition  $(s_t, a_t, r_t, s_{t+1})$  in  $\mathcal{D}$ 
    Sample random minibatch of transitions  $(s_j, a_j, r_j, s_{j+1})$  from  $\mathcal{D}$ 
    Set  $y_j = \begin{cases} r_j & \text{for terminal } s_{j+1} \\ r_j + \gamma \max_{a'} Q(s_{j+1}, a'; \theta) & \text{for non-terminal } s_{j+1} \end{cases}$ 
    Perform a gradient descent step on  $(y_j - Q(s_j, a_j; \theta))^2$ 
  end for
end for

```

---

We propose the following single-qubit unitary rotation method to encode the classical input states from the testing environment into the quantum circuit of Fig. 5. In quantum computing, the single-qubit gate with rotation along the  $j$ -axis by angle  $\theta$  is given by

$$R_j(\theta) = e^{-i\theta\sigma_j/2} = \cos \frac{\theta}{2} I - i \sin \frac{\theta}{2} \sigma_j, \quad (6)$$

where  $I$  is an identity matrix and  $\sigma_j$  is the Pauli matrix with  $j = x, y, z$ . The rotation angles for  $R_x(\theta_i)$  and  $R_z(\phi_i)$  in Fig. 5 are chosen to be

$$\theta_i = \pi \times b_i, \quad (7)$$

$$\phi_i = \pi \times b_i, \quad (8)$$

where  $i$  represents the index of qubit  $i$  and  $\pi$  here is the radian. In the quantum circuit with four input qubits, the index is the set  $\{1, 2, 3, 4\}$ . The rotational angle parameters  $\theta_i$  and  $\phi_i$  are for state preparation and are not subject to iterative optimization.

Take the state labeled 11 observed by the agent as an example. The decimal number 11 of the state is first converted to the binary number 1011, and then this classical state will be encoded into a quantum state  $|1\rangle \otimes |0\rangle \otimes |1\rangle \otimes |1\rangle = |1011\rangle$ . The detailed procedure is as follows. In this case, the binary digits  $b_1, b_2, b_3, b_4$  are 1, 0, 1, 1 respectively. Then according to Eqs. (7) and (8), One can obtain the values of  $\theta_i$  and  $\phi_i$  as  $(\theta_1, \theta_2, \theta_3, \theta_4) = (\pi, 0, \pi, \pi)$  and  $(\phi_1, \phi_2, \phi_3, \phi_4) = (\pi, 0, \pi, \pi)$ . One can furthermore obtain from Eq. (6)

$$R_x(\pi) = -i\sigma_x, \quad (9)$$

$$R_z(\pi) = -i\sigma_z, \quad (10)$$

and

$$R_x(0) = I, \quad (11)$$

$$R_z(0) = I. \quad (12)$$

When the two quantum gates for encoding,  $R_z(\theta_i)R_x(\phi_i)$ , operate on each qubit in initial state  $|0\rangle$  as shown in Fig. 5,

the resultant qubit state becomes either

$$R_z(\pi)R_x(\pi)|0\rangle = (-i\sigma_z)(-i\sigma_x)|0\rangle = |1\rangle \quad (13)$$

or

$$R_x(0)R_x(0)|0\rangle = II|0\rangle = |0\rangle. \quad (14)$$

Thus one obtains for  $b_1, b_2, b_3, b_4$  being 1, 0, 1, 1, respectively, a quantum state  $|1\rangle \otimes |0\rangle \otimes |1\rangle \otimes |1\rangle$ . Other classical states can be encoded into their corresponding quantum states in the same way. This procedure is applicable for all the experiments in this work, including cognitive-radio experiments, regardless of the number of qubits.

In the quantum circuit, the controlled-NOT (CNOT) gates are used to entangle quantum states from each qubit.

$$R(\alpha_i, \beta_i, \gamma_i) = R_z(\alpha_i)R_y(\beta_i)R_z(\gamma_i) \quad (15)$$

represents the general single qubit unitary gate with three parameters. Parameters labeled  $\alpha_i, \beta_i$  and  $\gamma_i$  are the ones for iterative optimization.

The variation quantum circuit is flexible in circuit depth. A shallow circuit that well represents the solution space can still achieve approximately the goal of certain tasks although a more sophisticated and deeper circuit may have better performance. But it remains a challenge to choose the right effective circuit that can parametrize and represent the solution space well for a general task while maintaining a low circuit depth and a low number of parameters [40]. It has been empirically demonstrated that a strongly entangling low-depth circuit has the potential powers and advantages to efficiently represent the solution space for some specific problems [23], [24], [31], [40]. Thus we design the strongly entangling circuit by appending a layer (i.e., the grouped box in dashed lines in Fig.5) comprised of two-qubit CNOT gates and parametrized general single-qubit unitary gates. Note that the layer or grouped box in dashed lines may repeat several times to increase the expressibility, entangling capability and also the number of parameters [40]. To accommodate both the

**TABLE 2.** Comparison of the number of parameters in classical Q-learning and variational quantum deep Q network (VQ-DQN).

Environment	2-Channels	3-Channels	4-Channels	5-Channels	Frozen-Lake
Q-Learning	$2 \times 2 \times 2$	$3 \times 3 \times 3$	$4 \times 4 \times 4$	$5 \times 5 \times 5$	$4 \times 4 \times 4$
VQ-DQN	$2 \times (3 \times 2 + 1)$	$3 \times (3 \times 2 + 1)$	$4 \times (3 \times 2 + 1)$	$5 \times (3 \times 2 + 1)$	$4 \times (3 \times 2 + 1)$

use of NISQ machines and the performance of the variational quantum circuits, in this work, the grouped box repeats two times regardless of the testing environments. The number of qubits can be adjusted to fit the problem of interest and the capacity of the simulators or quantum machines. For example, in the frozen-lake experiments and also the four-channel cognitive-radio experiments that will be discussed later, there are four input qubits and the grouped circuit repeats twice. Therefore the total number of circuit parameters subject to optimization is  $4 \times 3 \times 2 = 24$ . It is often to add a bias after the quantum measurement, the length of the bias vector is the same as the number of qubits. The bias vector is also subject to optimization. Therefore, the total number of parameters in this example is  $24 + 4 = 28$  which is also listed in Table 2.

## B. QUANTUM CIRCUIT FOR COGNITIVE RADIO NETWORKS

In the experiments on the *cognitive radio* [30], the total number of channels  $n$  that can be selected by the agent at each time-step is known in advance. Since the occupied channel changes from time to time, it is necessary to include not only the *channel* but also the *temporal* information into the observation. The *observation* is in the following form: (*channel*, *time*). In our experimental setting, we consider that the channel-changing by the primary user follows a simple periodic pattern with  $n$  time steps in a full cycle (see Fig. 3 for the  $n = 4$  case). Therefore, the number of possible states is  $n^2$ . Similar to the circuit architecture of the frozen-lake experiments in Fig. 5, the variational quantum circuits for cognitive-radio experiments are shown in Fig. 6. We choose different qubit numbers to accommodate possible  $n^2$  states of the  $n$ -channel cognitive-radio environment. Normally,  $n$  qubits are used for the encoding and action (channel) selection of an  $n$ -channel cognitive-radio environment. However, the 3-channel case is special since the number of possible state is  $3^2$ , which is greater than the number of computational basis states of 3-qubit system (which is  $2^3 = 8$ ). We thus use a 4-qubit system to accommodate the all possible states. The scheme to encode the classical  $n^2$  states into their corresponding quantum states in the cognitive-radio environment is the same as that in the frozen-lake experiment introduced in Sec. IV-A except that some quantum states are not used when  $2^{n_q} > n^2$ , where  $n_q$  is the qubit number and  $n$  is the channel number.

In addition, at each time step, the agent can select one of the channels from the set of all possible channels, which is of number  $n$ . This corresponds exactly to the  $n$  possible actions that the agent can select. The action selection in our VQ-DQN scheme is determined through the expectation values of  $n$  qubits, which will be discussed in Sec. VI-A. Since  $n_q > n$

for channel number  $n = 3$ , we use the repeated quantum measurements of only the first three qubits for the estimation of the expectation values for channel or action selection. At the same time, we would like to keep the number of parameters to scale as  $n \times (3 \times 2 + 1)$ . So for the special case of  $n = 3$ , there is no single-qubit unitary operation with optimization parameters acting on the fourth qubit in the grouped box in Fig. 6(b). We note here that since the number of possible qubit states  $2^n$ , which is greater than  $n^2$  for  $n > 5$ , grows exponentially, leaving our variational quantum circuit the possibilities to deal with more complex scenarios with more possible states than the simple periodical channel-changing pattern with  $n^2$  possible states when the number of channels  $n$  becomes large.

## V. DATA AVAILABILITY

The data that support the findings of this study are available in the GitHub repository, <https://github.com/ycchen1989/VarQuantumCircuits-DeepRL>

## VI. EXPERIMENTS AND RESULTS

### A. ENVIRONMENT SETUP AND ACTION SELECTION

The frozen-lake testing environment is depicted in Fig. 1. We set up the experiment following the circuit architecture in Fig. 5. For the cognitive-radio environment, we consider the case where the external interference of the primary user follows a periodic pattern, i.e. sweeping over all channels from 1 to  $n$  in the same order (see, e.g., Fig. 3 for the  $n = 4$  case). We set up the experiment following the circuit architectures in Fig. 6 according to the number of channels  $n$  in the cognitive-radio environment. This cognitive-radio environment offers a feasible test-bed for quantum DQL (DQN) with a desirable self-defined environment with lower action and space complexity working in the current NISQ machines.

Next, we describe how to select an action in the variational quantum circuit. The measurement output of the expectation (ensemble average) values of the  $n$ -qubit quantum circuit is an  $n$ -tuple. The index counting of the measurement qubit output ports or wires is from zero to  $n - 1$ . The choice of the action is just the index of the measurement qubit output port or wire that has the largest expectation value. Taking the four-qubit setting as an example, one then has the output of  $[a, b, c, d]$  with  $a, b, c, d$  being numerical expectation values obtained from a set of measurements. If the numerical value  $b$  is the largest value among the four expectation values, then the action choice is 1. This action 1 will then be sent to the testing environment. To be more specific, for the frozen-lake environment, there are four actions LEFT, DOWN, RIGHT, UP in the action space. The output ports or wires of the four-qubit quantum circuit are labeled 0, 1, 2, 3



and they correspond to the action LEFT, DOWN, RIGHT, UP, respectively. If output wire 1 has the largest expectation value, then the action that will be selected by the agent is to go DOWN one step from the current state in the frozen-lake environment. Similarly, there are  $n$  possible action choices, Channel 1 to Channel  $n$ , in the cognitive-radio environment, and they correspond to the output qubit ports or wires labeled from 0 to  $n - 1$ , respectively. If the output qubit wire 1 has the largest measured expectation value, the agent (the secondary user) will select Channel 2 as the action.

As described above, the next action that the agent selects is determined by the expectation value of each qubit, not by the random outcome of 0 or 1 of each qubit in each single run of measurement. The expectation value can be calculated analytically if we use the quantum simulator (for example, PennyLane [42] or IBM Qiskit) on a classical computer, and in this case the result is deterministic. If the agent is on a real quantum computer, the expectation value is estimated with a large number of measurement samplings, which should be close to the value calculated theoretically by the quantum simulator. Let us give a brief summary of the whole procedure on a real machine with an example. Suppose the RL agent receives a classical binary number for the state 1011. First, this binary number will be converted into the quantum state  $|1\rangle \otimes |0\rangle \otimes |1\rangle \otimes |1\rangle$  by the encoding single-qubit gates and then go through the quantum circuit blocks before the measurements. If we measure the first qubit, for example, we will get either 0 or 1, but which of these two possible outcomes will appear is truly stochastic. Thus this single-shot random measurement outcome is not enough for the agent to decide the next move or action, and instead, the expectation value is used. To find the expectation value, one needs to measure an ensemble of identically prepared systems, or perform an ensemble of measurements on a single quantum system repeatedly prepared in the identical quantum state  $|1\rangle \otimes |0\rangle \otimes |1\rangle \otimes |1\rangle$ , which is our case here. So the agent will prepare the qubits in state  $|1\rangle \otimes |0\rangle \otimes |1\rangle \otimes |1\rangle$  and let it go through the quantum circuit blocks and then perform a measurement on each qubit. This process will be repeated for a large number of times. The measurement outcome of being in state  $|0\rangle$  or state  $|1\rangle$  in each measurement is purely random and unpredictable. However, the probability of being in state  $|0\rangle$  and the probability of being in state  $|1\rangle$  can be predicted. The average value of all these repeated measurements is the expectation value. For example, if the agent repeats the process 1000 times and get 600 times of 1 and 400 times of 0 on the first qubit, then the expectation value of the first qubit will be about 0.6. The procedure is the same for other qubits when we consider multi-qubit systems.

Normally, the number of observations or measurements needed to learn a discrete probability distribution on an  $n$ -qubit system of size  $2^n$  is, in the worst case, linear in  $2^n$ . This means that learning the probabilities may require a number of measurements that scales exponentially with the number of qubits  $n$ . However, what is necessary for our algorithm to determine the next action by the QRL agent

here is to find the expectation value of each qubit in the  $n$ -qubit system, rather than the expectation value on the  $n$ -qubit computational state space. As a result, we only need to measure discrete probability distribution of size  $n$  rather than size  $2^n$ . Moreover, the measurements on each qubit of the  $n$  qubits can be performed simultaneously in parallel. Thus the number of repeated experiments (measurements) to obtain the expectation values can be chosen to be a fixed number, rather independent of  $n$ . Furthermore, our algorithm does not require to find the exact expectation value of each qubit, and only the qubit that has the largest expectation value is concerned, making our DRL algorithm relatively robust against noise and errors in the real quantum machines. We note here that in the inference experiment with our trained QRL model running on a real quantum computer discussed in Sec. VI-E, the number of repeated measurements is set to 1024, and this number of measurement shots can already give a fairly stable result.

## B. NUMERICAL SIMULATION

The quantum circuits as constructed in Fig. 5 and Fig. 6 are numerically simulated with the software package PennyLane [42]. We use the standard package PyTorch [43] to help the linear algebraic manipulation to accelerate the simulation. OpenAI Gym [29] provides the testing environment. In this work, we choose an environment with low computational complexity, the frozen-lake environment, to implement the proof-of-concept experiments and choose the cognitive-radio environment to study the possible real-world applications.

The optimization method is chosen to be *RMSprop* [44] in which parameters are set to be *learning rate* = 0.01, *alpha* = 0.99 and *eps* =  $10^{-8}$ , widely used in DRL. Note that the learning rate, alpha and eps here are only for the gradient descent optimization. Please do not be confused with the aforementioned DRL hyperparameters. The batch-size for the experience replay is 5. The  $\epsilon$ -greedy strategy used in the frozen-lake environment is the following:

$$\epsilon \leftarrow \frac{\epsilon}{\frac{\text{episode}}{100} + 1}, \quad (16)$$

and in the cognitive-radio environment  $\epsilon$  is updated in every single step as:

$$\epsilon \leftarrow 0.99\epsilon \quad (17)$$

with initial  $\epsilon = 1.0$  for encouraging more exploration in early episodes and shifting to more exploitation in later episodes.

## C. SIMULATION WITH NOISE

To investigate the robustness of our proposed variational quantum circuit-based DRL against the noise from current and possible near-term NISQ devices, we perform additional simulations which include the noise from a real quantum computer. The experiment setting is the same as the previous experiments, except that the simulation backend is replaced with the Qiskit-Aer simulation software, which has the capability to incorporate the noise models

from the IBM quantum computers. We perform the noise-model simulations for the 2-channel and 4-channel settings in cognitive-radio experiments. The qubit properties, such as relaxation time  $T_1$ , dephasing time  $T_2$ , qubit frequency, gate error and gate length, of the IBM Q 20-qubit machine `ibmq-poughkeepsie`, from which we download noise data for the above mentioned noise-model simulations, are listed in Tables 6, 9, and 8 in Appendix A.

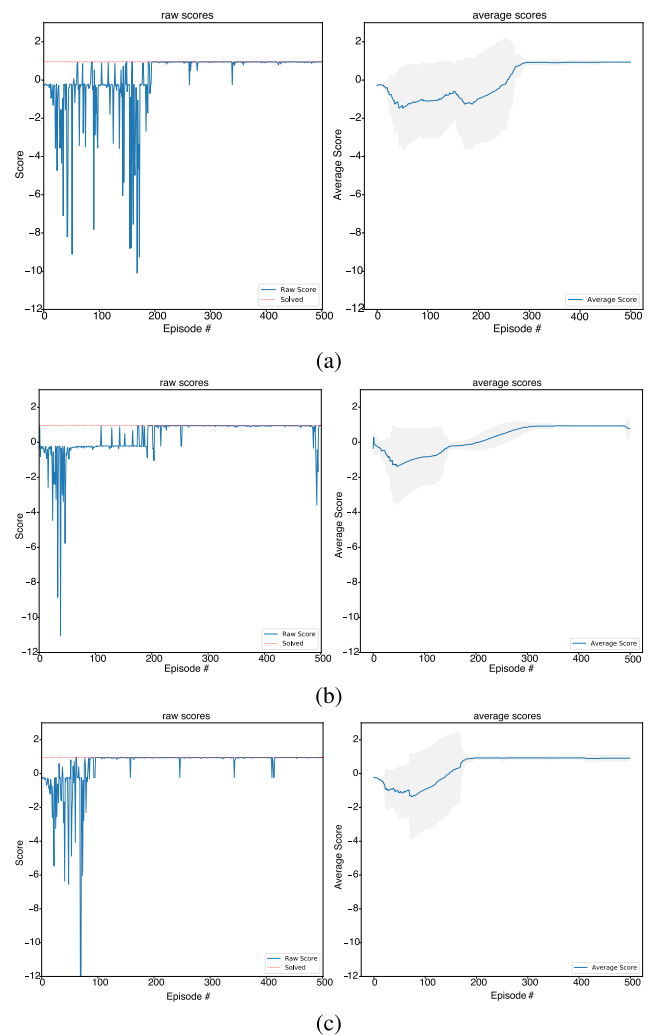
The variational quantum circuits can be relatively robust against noises because they involve a classical optimization step and the related deviations can be absorbed by the parameters during the iterative optimization process. Another appealing feature of our quantum variational DRL algorithm is the additional ability to tolerate errors and noises. The next action that the agent selects in our algorithm is determined by which qubit has the largest expectation value among all the  $n$  qubits. In other words, although quantum gate operations and the measurement fidelity on NISQ devices are degraded by noises and errors, as long as the qubit that has the largest expectation value is the same as that on the ideal quantum simulator, the DRL agent will select the same action.

#### D. PERFORMANCE ANALYSIS

In the frozen-lake experiment, we run 500 episodes on all three configurations. Subfigures (a), (b), and (c) in Fig. 7 correspond to the results of the environment configurations (a), (b), and (c) in Fig. 1, respectively. Take subfigure (a) for example. The agent converges to the total reward 0.9 after the 198th episode. The results of the other two configurations shown in subfigures (b) and (c) in Fig. 7 also demonstrate the similar pattern. It is noted that, however, several sub-optimal results occur. This phenomenon is probably due to the  $\epsilon$ -greedy policy selection.

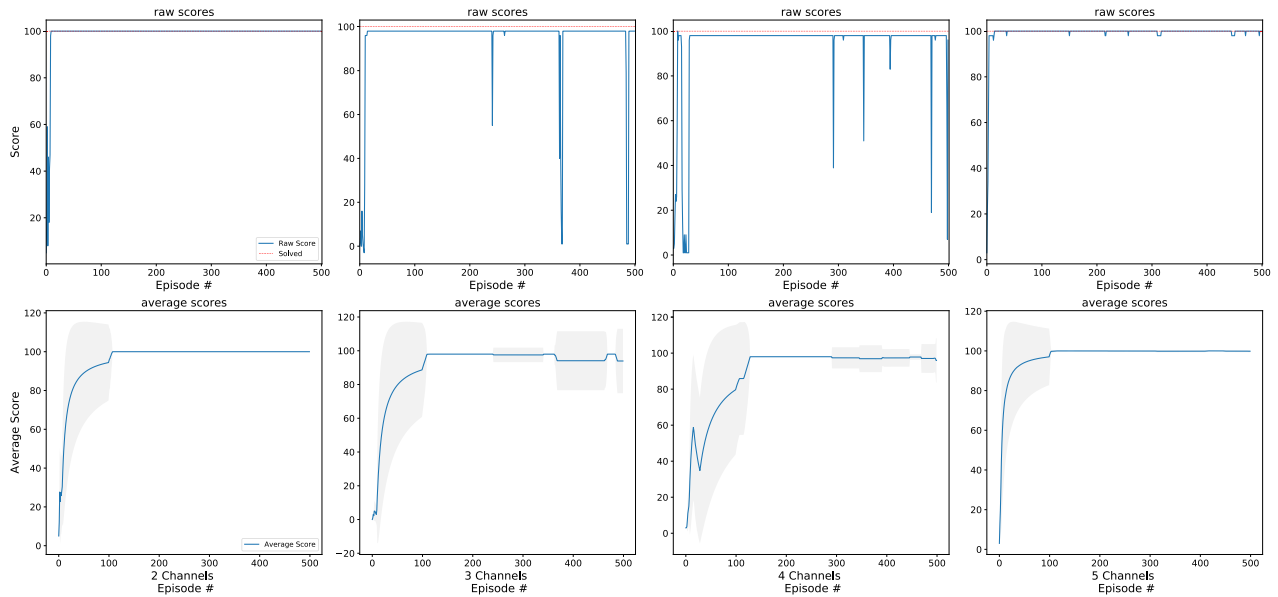
To demonstrate the stability of our quantum RL agents' training process, we calculate the temporary mean value and the standard-deviation boundary of the total rewards. The mean and standard-deviation values are calculated from the last 100 episodes. In the case that there are fewer than 100 episodes, we include all the episodes. The standard deviation values are shown in gray color and the mean values are in blue color in the right panel of Fig. 7. In our simulations, the quantum RL agents' average total reward (average score) are with small standard-deviation values, meaning that they are stable after training.

In the cognitive-radio experiment, we tested situations where there are 2, 3, 4 and 5 possible channels with the environment configuration shown in (a) of Fig. 3 for the agent to choose, respectively. In all the four situations, we run 500 episodes and the agent converges to the optimal total reward value at around 100 iterations or episodes (see Fig. 8). We further perform the simulations on the other two 4-channel environment configurations shown in (b) and (c) of Fig. 3, the training result shown in Fig. 9 is comparable to the 4-channel training result on environment configuration (a) shown in Fig. 8. In the simulation with the noise model from the IBM quantum machine (with the environment



**FIGURE 7.** Performance of the variational quantum circuits for DQL on the frozen-lake experiment. Subfigures (a), (b), and (c) correspond to the results of the environment configurations (a), (b), and (c) in Fig. 1, respectively. Take subfigure (a) for example. The left panel of subfigure (a) shows that our proposed variational quantum circuits based DQL-agent reaches the optimal policy in the frozen-lake environment with a total reward of 0.9 at the 198th iteration. The gray area in the right panel of the subfigure represents the standard deviation of reward in each iteration during exploration with the standard reinforcement learning reproducible setting for stability. The mean and the standard deviation values of the average score (reward) are calculated from the scores (rewards) in the past 100 episodes. The policy learned via quantum circuit becomes more stable after the 301st iteration. The other two experiments in subfigures (b) and (c) demonstrate the similar pattern.

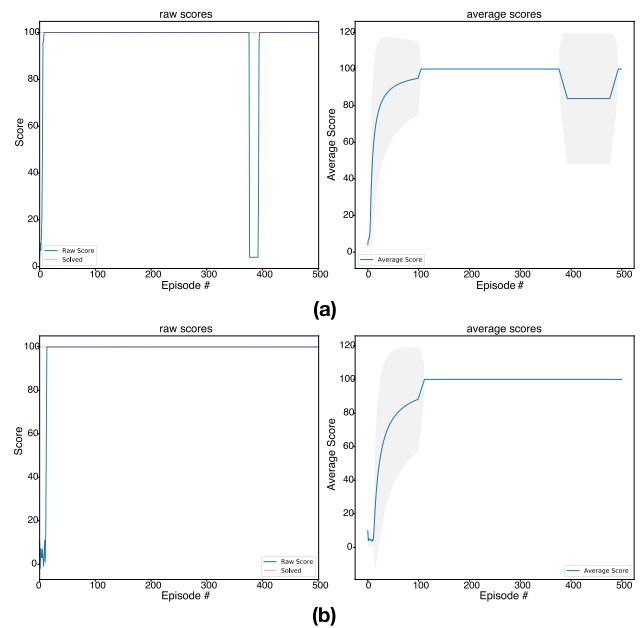
configuration shown in (a) of Fig. 3), the agent converges to the optimal total reward value at around 110 iterations in the 2-channel and 4-channel settings (see Fig. 10), which are comparable to previous ideal and noiseless simulations in Fig. 8, showing that our proposed quantum circuit based DRL is robust against noise in the current machines. The mean and standard-deviation values of the total rewards for cognitive-radio experiments are shown in the right panel of Fig. 9 and the bottom panels of Fig. 8 and Fig. 10. The small standard-deviation values indicate that they are stable after training.



**FIGURE 8.** Performance of the Variational quantum circuits for DQL on the cognitive-radio experiment. In the cognitive-radio experiments, we limit the maximum steps an agent can run to be 100, and the reward scheme is that for each correct choice of the channel, the agent will get a +1 reward and  $-1$  for incorrect selection. The maximum reward an agent can achieve under this setting is 100. The top panels of the figure show that our proposed variational quantum circuits based DQL-agent reaches the optimal policy with a total reward of 100 in the 2-channel and 5-channel cases at only several iterations. For the cases of 3-channel and 4-channel, the agent also reaches near-optimal policy at only several iterations. The gray area in the bottom panels of the figure represents the standard deviation of reward in each iteration during exploration with the standard reinforcement learning reproducible setting for stability [41]. The mean and the standard deviation at each episode are calculated from the rewards(scores) in the past 100 episodes. The policy learned via variational quantum circuits becomes more stable after around 100 iterations for all the four cases. This part of experiment is tested on the configuration (a) shown in Fig. 3.

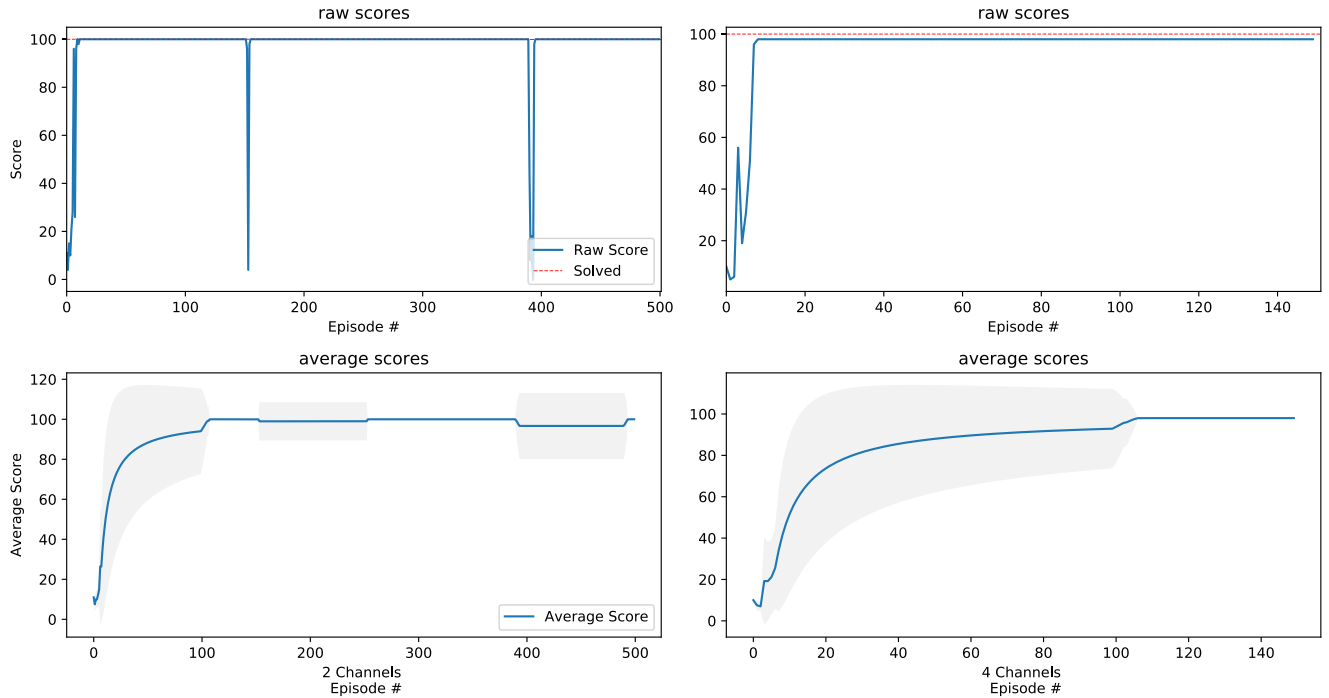
### E. RUNNING ON A QUANTUM COMPUTER

We further upload our trained VQ-DQL (VQ-DQN) models to the IBM Q cloud-based quantum computing platform to investigate whether the models are feasible on a real quantum computer. In the cognitive-radio experiment, we upload the trained model parameters of the variational quantum circuit of configuration (c) in the 4-channel case described in Fig. 3. Due to the limited resource available on the cloud-based quantum computing platform, we exclusively carry out five episodes of this specific experiment on the IBM Q backend machine *ibmq-valencia*. The result of the experiment conducted on the IBM Q machine listed in Table 3 has almost the same total reward as that obtained by running on the PennyLane or Qiskit quantum simulator. This demonstrates that even if the training is on the simulation software without the noise, the trained model of the variational quantum circuit for DRL still performs well on the real quantum computer. The reason for this is that our quantum DRL algorithm does not require to find the exact expectation value of each qubit, and only cares which qubit that has the largest expectation value. Thus, our algorithm is relatively robust against errors and noises. In the frozen-lake experiment, we test the trained model parameters for the environment configuration (c) described in Fig. 1 also on the IBM Q machine *ibmq-valencia* and carry out seven episodes of this specific experiment. The result listed in Table 4 for the quantum DRL model experiment conducted on the IBM Q quantum computer is also comparable to that obtained by running on



**FIGURE 9.** Performance of the Variational quantum circuits for DQL on the cognitive-radio experiment. This part of experiment is tested on configurations (b) and (c) shown in Fig. 3, and the results are comparable to the results in Fig. 8, which are tested on configuration (a) in Fig. 3. This demonstrates that our proposed quantum DRL (DQL) framework can be trained on different scenarios.

the PennyLane or Qiskit quantum simulator. In these real-machine experiments, the number of shots of quantum measurements for the calculation or estimation of the expectation



**FIGURE 10.** Performance of the variational quantum circuit for DQL on cognitive-radio experiment conducted on the IBM Qiskit Aer quantum simulator with noise from the real quantum machine, `ibmq-poughkeepsie`. In this experiment, we investigated the robustness of the VQ-DQN against the noise in quantum machines. The device noise is downloaded from remote real quantum machines and then incorporated with the simulator software Qiskit Aer provided by IBM. This noisy backend then replaces the noiseless backend in previous experiments. The result of the experiment shows that our proposed variational quantum circuits based DQL or DQN is robust against noise in the current quantum machines.

**TABLE 3.** Results of the trained VQ-DQL (VQ-DQN) for the cognitive-radio experiment conducted on the IBM Q quantum computer, `ibmq-valencia`. In this experiment, we test the trained quantum DRL model in the cognitive-radio experiment of configuration (c) in the 4-channel case described in Fig. 3. Even if the training is on the simulation software without the quantum noise, the trained model still performs well on the real quantum computer.

Episodes	1	2	3	4	5
Total Steps	100	100	100	100	100
Total Reward	100	100	100	100	98

**TABLE 4.** Results of the trained VQ-DQL (VQ-DQN) for the frozen-lake experiment conducted on the IBM Q quantum computer, `ibmq-valencia`. In this experiment, we test the trained quantum DRL model in the frozen-lake experiment of configuration (c) in Fig. 1. Even if the training is on the simulation software without the quantum noise, the trained model still performs well on the real quantum computer.

Episode	1	2	3	4	5	6	7
Total Steps	6	6	6	7	7	7	6
Total Reward	0.95	0.95	0.95	0.94	0.94	0.94	0.95

values is 1024. The qubit properties, such as relaxation time  $T_1$ , dephasing time  $T_2$  and qubit frequency, gate error and gate length, of the IBM Q 5-qubit machine `ibmq-valencia` we use for the above mentioned two experiments are listed in Tables 11, 10, and 7 in Appendix B.

### F. QUANTUM ADVANTAGE ON MEMORY CONSUMPTION

Current NISQ machines are not suitable for deep quantum circuit architectures due to the lack of quantum error correction. Therefore, to utilize these near-term NISQ machines

for more complex situations, it is urgent to develop quantum circuits which are not so deep in quantum gates as there will be more errors when the number of quantum gates increases. With this constraint, the proposed variational quantum circuit cannot have many parameters compared to classical neural networks. In this work, we show that for the cases we consider, variational quantum circuits with fewer iterative parameters can achieve comparable performances to classical neural networks.

To compare the performance of our VQ-DQL (VQ-DQN) models with classical counterparts, we need at least one classical candidate. In the frozen-lake environment, the number of parameters in our proposed VQ-DQL is 28, while in tabular Q-Learning, the table that needs to store all the state-action information is of the size  $16 \times 4 = 64$ . Thus our VQ-DQL method reduces the number of parameters by  $(64 - 28)/64 = 56.25\%$ . In the cognitive-radio environment, the classical counterpart that we compare is from the original ns3-gym work [30], and it is a fairly simple neural network consisting of a single hidden layer. From the results of Fig. 6 and Fig. 8 in that original paper, we can see that our agents converge faster than the classical ones. For example, even in the presence of noise, our agents can reach score very close to 100 at less than 10 iterations as shown in Fig. 10 while in the original paper [30], the agents reach score 100 in more than 20 iterations. Moreover, in this example, our proposed quantum circuit uses fewer parameters than its classical counterpart. For a given problem, there are many possible neural network



**TABLE 5.** Comparison of classical reinforcement learning algorithms with discrete action space and variational quantum deep Q networks (VQ-DQN).

Algorithm	Policy	Action Space	State Space	Operator	Complexity of Parameters
Monte Carlo	Off-policy	Discrete	Discrete	Sample-means	$\mathcal{O}(n^3)$
Q-Learning	Off-policy	Discrete	Discrete	Q-value	$\mathcal{O}(n^3)$
SARSA	On-policy	Discrete	Discrete	Q-value	$\mathcal{O}(n^3)$
DQN	Off-policy	Discrete	Continuous	Q-value	$\mathcal{O}(n^2)$
VQ-DQN <sup>a</sup>	Off-policy	Quantum	Quantum	Q-value	$\mathcal{O}(n)$
VQ-DQN <sup>b</sup>	Off-policy	Quantum	Quantum	Q-value	$\mathcal{O}(\log n)$

<sup>a</sup> The number of parameters in VQ-DQN with computational basis encoding grows only linearly with the dimension of the input vector  $n$ .

<sup>b</sup> VQ-DQN with amplitude encoding can harvest full logarithmic less parameters compared with classical models.


architectures and it is impossible to exhaust all the possible solutions. In our case, we select the neural network architecture from the article which demonstrates the ns3 simulation framework [30]. For the experiments we conduct and test in the cognitive-radio environments with  $n$  possible channels, the number of parameters is  $n \times (3 \times 2 + 1)$  in our variational-quantum-circuit-based DQL (DQN), while it is  $2 \times n^2 + 2 \times n$  in the neural network based RL [30]. In general, there is no guarantee that a given neural network model or a variational quantum circuit can scale as we wish when the complexity of the problem increases. Therefore, in this work, the comparison in scaling is restricted to the cases we test. For example, in the cognitive-radio experiment, we consider only the cases with number of channels  $n \leq 5$ .

In summary, the quantum advantage in our proposed method refers to the less memory consumption, which means there are less parameters in the quantum circuits. The quantum advantage in terms of the parameters of the quantum circuit relies on the data encoding schemes. For example, the number of parameters in an amplitude coding may be  $\text{poly}(\log n)$  in contrast to  $\text{poly}(n)$  in a standard neural network. However, it is hard to implement the amplitude encoding scheme as there is no known efficient algorithm to prepare classical vectors into quantum registers [45]. The computational basis encoding in our proposed variational quantum DQL (DQN) involves  $n$  parameters, but there are  $n^2$  parameters in the neural network based RL [30], and  $n^3$  in the tabular Q-learning, where  $n$  is the dimension of input vectors. A comparison of classical reinforcement learning algorithms with discrete action space and variational quantum deep Q networks (VQ-DQN) is given in Table 5. The blue dotted lines in Fig. 11 shows explicitly that the number of parameters of our proposed VQ-DQN method grows linearly with the dimension of the input vector, at least for the cases we test, i.e., the number of channels  $n \leq 5$  in the cognitive-radio experiments.

## VII. DISCUSSIONS

### A. OVERVIEW OF QUANTUM MACHINE LEARNING

A general review of the field of *quantum machine learning* can refer to [46], [47]. As for *quantum reinforcement learning*, the early work can refer to [48], and in this work, the authors use the amplitude amplification method to perform the action selection. However, such operations are hard to implement on NISQ machines. In addition, in their



**FIGURE 11.** Comparison of memory consumption in different learning schemes. The figure shows our proposed variational quantum circuits based DQL-agent has the quantum advantage in memory consumption compared to classical Q-learning and DQL (DQN). Specifically, in our cognitive-radio channel selection experiment, we set up four different testing environments with 2,3,4,5 possible channels, respectively. With the classical Q-learning, the number of parameters grow with  $n^3$ , and in DQL (DQN), the number of parameters grow with  $n^2$ . However, with our proposed variational quantum circuits based DQL (DQN), the number of parameters grows as  $n \times (3 \times 2 + 1)$  only.

scheme [48], there is a need to encode the environment in the quantum superposition state, which is not always possible when the agent interacts with classical environments. To study the effect of computational agents coupled to environments which are quantum-mechanical, refer to [49]. For a review of recent developments in quantum reinforcement learning, refer to this good review article [50]. For example, recent work [51], [52] have proposed a framework called *projective simulation*, and the key concept in this setting is that the *agent keeps memory* of the transition history. Before executing each action, the agent will *simulate* several possible outcomes according to historical data stored in the memory. It is conceptually related to the well-known Monte-Carlo Tree Search [7], [53] and it is interesting to investigate the quantum counterparts and possible quantum advantages.

### B. MORE COMPLEX TESTING ENVIRONMENTS

We have applied our VQ-DQN to a simple maze problem, the frozen lake environment in OpenAI Gym [29], and to a classical spectrum control problem in cognitive radio with the ns3-gym [30] environment. We choose the frozen-lake environment as an environment with low computational complexity to implement the proof-of-principle quantum DQL (DQN)

**TABLE 6.** Qubit relaxation time  $T_1$ , dephasing time  $T_2$  and qubit frequency data of the 20-qubit machine `ibmq-poughkeepsie` downloaded from the IBM Q system service at the time when the numerical simulation in Sec.VI-C performed.

Qubit	$T_1$ ( $\mu$ s)	$T_2$ ( $\mu$ s)	Frequency (GHz)
0	61.985008	63.458383	4.919976
1	85.949810	117.651095	4.831989
2	51.816944	67.602274	4.939780
3	66.607063	67.884742	4.515083
4	100.442250	100.323516	4.66276
5	66.118221	65.421503	4.957158
6	67.734257	86.704533	4.995404
7	76.736639	17.822148	4.811046
8	66.743110	64.9755788	5.013321
9	73.998165	82.741817	5.056846
10	68.916861	12.902928	4.719699
11	68.614596	86.433312	4.900151
12	45.072575	9.751853	4.772410
13	57.462386	23.218794	5.110509
14	72.111045	83.851059	4.990636
15	92.118025	41.604544	4.806460
16	89.207166	61.459798	4.955924
17	106.508278	18.273605	4.599708
18	93.291790	94.709020	4.828349
19	77.475090	99.860137	4.938456

**TABLE 7.** Two-qubit coupling and CNOT gate error and gate length data of the 5-qubit machine `ibmq-valencia` downloaded from the IBM Q system service at the time when the experiment in Sec.VI-E performed.

Index	Coupling pair	CNOT gate error	CNOT gate length (ns)
0	[0, 1]	0.008928	320.0000
1	[1, 0]	0.008928	284.4444
2	[1, 2]	0.007558	341.3333
3	[1, 3]	0.011670	661.3333
4	[2, 1]	0.007558	376.8889
5	[3, 1]	0.011670	696.8889
6	[3, 4]	0.009659	305.7778
7	[4, 3]	0.009659	341.3333

experiments. We consider the cognitive-radio problem in a wireless multi-channel environment, e.g. 802.11 networks with external interference. The objective of the agent is to select a channel free of interference in the next time slot. We create a scalable reinforcement learning environment sim-radio-spectrum (SRS) with a customized state and an action echo in a real multi-channel spectrum scenario for the quantum DQL (DQN) demonstration of a real-world application.

Different from the benchmark environments of complex RL like Atari games, we use simple testing environments to study the feasibility of quantum circuits-based DQL (DQN). Although the mainstream benchmark environments can be encoded with computational basis encoding or amplitude encoding, the number of qubits needed is intractable for numerical simulations on classical computers and also exceeds currently commercially available quantum devices. However, we could further investigate these complex RL environments with the same setting proposed in this work when large quantum machines are released.

### C. SCALING UP THE ARCHITECTURE

In the proposed architecture, we require the observation of the expectation value of each qubit in order to determine the next action, which can be evaluated analytically by a

**TABLE 8.** Two-qubit coupling and CNOT gate error and gate length data of the 20-qubit machine `ibmq-poughkeepsie` downloaded from the IBM Q system service at the time when the numerical simulation in Sec.VI-C performed. The two-qubit CNOT gate error and lengths are usually different for different two-qubit coupling pairs due to different coupling strength and different qubit frequencies in the qubit pairs.

Index	Coupling pair	CNOT gate error	CNOT gate length (ns)
0	[0, 1]	0.019174	455.1111
1	[0, 5]	0.022823	433.7778
2	[1, 0]	0.019174	455.1111
3	[1, 2]	0.018289	455.1111
4	[2, 1]	0.018289	455.1111
5	[2, 3]	0.027451	839.1111
6	[3, 2]	0.027451	839.1111
7	[3, 4]	0.023150	832.0000
8	[4, 3]	0.023150	832.0000
9	[4, 9]	0.123852	888.8889
10	[5, 0]	0.022823	433.7778
11	[5, 6]	0.017515	455.1111
12	[5, 10]	0.022059	476.4444
13	[6, 5]	0.017515	455.1111
14	[6, 7]	0.029664	618.6667
15	[7, 6]	0.029664	618.6667
16	[7, 8]	0.024846	398.2222
17	[7, 12]	0.024576	554.6667
18	[8, 7]	0.024846	398.2222
19	[8, 9]	0.021136	426.6667
20	[9, 4]	0.123852	888.8889
21	[9, 8]	0.021136	426.6667
22	[9, 14]	0.032180	568.8889
23	[10, 5]	0.022059	476.4444
24	[10, 11]	0.020104	661.3333
25	[10, 15]	0.014944	476.4444
26	[11, 10]	0.020104	661.3333
27	[11, 12]	0.016619	448.0000
28	[12, 7]	0.024576	554.6667
29	[12, 11]	0.016619	448.0000
30	[12, 13]	0.132743	618.6667
31	[13, 12]	0.132743	618.6667
32	[13, 14]	0.021946	398.2222
33	[14, 9]	0.032180	568.8889
34	[14, 13]	0.021946	398.2222
35	[14, 19]	0.017132	469.3333
36	[15, 10]	0.014944	476.4444
37	[15, 16]	0.021673	789.3333
38	[16, 15]	0.021673	789.3333
39	[16, 17]	0.019825	704.0000
40	[17, 16]	0.019825	704.0000
41	[17, 18]	0.020812	704.0000
42	[18, 17]	0.020812	704.0000
43	[18, 19]	0.015828	419.5556
44	[19, 14]	0.017132	469.3333
45	[19, 18]	0.015828	419.5556

quantum simulation software on a classical computer. While in a real quantum computer, it is needed to perform multiple samplings of repeated measurements. For example, in a  $n$ -qubit circuit with the number of output actions to be  $n$ , the number of samplings in the inference stage could be a fixed number of 1024 as the the measurements on each qubit of the  $n$  qubits can be performed simultaneously in parallel. Thus number of samplings needed could be pretty much fixed or may increase only slightly with the number of qubits  $n$  or the number of possible output actions. Furthermore, in the proposed quantum DRL algorithm, we only care which qubit or action is the most probable with the largest expectation

**TABLE 9.** Single-qubit U1, U2 and U3 gate errors, and readout error data of the 20-qubit machine `ibmq-poughkeepsie` downloaded from the IBM Q system service at the time when the numerical simulation in Sec. VI-C performed. Different single-qubit gates have different gate lengths, but the gate length of the same single-qubit gate is the same for every qubit. The gate lengths of the Identity (Id) gate, U1 gate, U2 gate, and U3 gate, are 113.7778 ns, 0.0 ns, 103.1111 ns, and 206.2222 ns, respectively.

Qubit	Id gate error	U1 gate error	U2 gate error	U3 gate error	Readout error
0	0.001207	0.0	0.001207	0.002413	0.032
1	0.000927	0.0	0.000927	0.001854	0.031
2	0.001200	0.0	0.001200	0.002401	0.030
3	0.001287	0.0	0.001287	0.002574	0.043
4	0.001383	0.0	0.001383	0.002767	0.067
5	0.001075	0.0	0.001075	0.002150	0.031
6	0.001545	0.0	0.001545	0.003091	0.040
7	0.003174	0.0	0.003174	0.006349	0.070
8	0.001416	0.0	0.001416	0.002832	0.031
9	0.001986	0.0	0.001986	0.003972	0.028
10	0.002024	0.0	0.002024	0.004047	0.061
11	0.000570	0.0	0.000570	0.001140	0.025
12	0.002100	0.0	0.002100	0.004200	0.059
13	0.003205	0.0	0.003205	0.006410	0.025
14	0.001408	0.0	0.001408	0.002816	0.029
15	0.000900	0.0	0.000900	0.001800	0.023
16	0.000901	0.0	0.000901	0.001801	0.018
17	0.001404	0.0	0.001404	0.002807	0.080
18	0.000769	0.0	0.000769	0.001538	0.023
19	0.000943	0.0	0.000943	0.001886	0.023

value and do not need the exact probability. Thus, even though repeated experiments (measurements) are needed for our quantum DRL algorithm performed on a real quantum computer, the probabilistic decision rules of choosing the action corresponding to the highest frequency outcome with a fixed number of measurements can give better performances in terms of number of operations, such that the classical simulation is still not favorable when larger architectures or qubit numbers are considered.

#### D. SIMULABILITY ON A CLASSICAL COMPUTER

The hardness of simulability of quantum circuits does not depend only on the number of qubits, but also on the structure of the circuit. In the proposed circuit architecture with fixed number of variational layers, the number of CNOT gates scales linearly with the number of qubits  $n$ . Therefore, the circuit is not constant-depth. In addition, the circuit includes quantum gates beyond the Clifford group, meaning that it is not in the family of Clifford circuits. Thus, it is not obviously classically simulable when the number of qubits is large. The hardness of simulability of this model is also worth studying theoretically.

#### E. FUTURE WORK: AMPLITUDE ENCODING SCHEME

Unlike the amplitude encoding, the computational basis encoding has not fully employed the quantum advantages. Although in a constraint condition of quantum simulators, we can verify the feasibility of applying quantum circuits for resolving DRL problems. The related empirical results suggest that the quantum advantages outperform both tabular  $Q$ -learning and neural network based RL. To obtain the ideal quantum circuits with significantly fewer parameters, one can apply amplitude encoding to reduce the complexity of parameters as small as  $\text{poly}(\log n)$  in contrast to a standard neural

network with  $\text{poly}(n)$  parameters. Future work includes an investigation of applying amplitude encoding scheme to more complex input data and variational quantum circuits to solving more sophisticated problems.

## VIII. CONCLUSION

This is the first demonstration of variational quantum circuits to approximate the deep  $Q$ -value function with experience replay and target network. From the results obtained from the testing environments we consider, our proposed framework shows the quantum advantage in terms of less memory consumption and the reduction of model parameters. Specifically for the considered testing environments in this paper, the variational quantum deep  $Q$ -learning involves parameters as small as  $\mathcal{O}(n)$ , but the tabular  $Q$ -learning and neural network based deep  $Q$ -learning have  $\mathcal{O}(n^3)$  and  $\mathcal{O}(n^2)$  parameters, respectively.

## APPENDIX NOISE INFORMATION OF THE IBM Q MACHINES USED IN THE EXPERIMENTS

### A. DEVICE PROPERTIES FOR THE NOISY SIMULATION

In Sec. VI-C, we perform the numerical simulation using the IBM Qiskit simulator with the noise data downloaded from the IBM Q 20-qubit machine `ibmq-poughkeepsie` at the date 2019-07-28. The qubit relaxation time  $T_1$ , dephasing time  $T_2$  and qubit frequency data are listed in Table 6, the single-qubit gate error, gate length and readout error data are in Table 9, and the two-qubit coupling and corresponding CNOT gate error and gate length data are listed in Table 8. Note that the IBM Q system service is still under active development, meaning that the noise data will change gradually.

**TABLE 10.** Single-qubit U1, U2 and U3 gate errors, and readout error data of the 20-qubit machine *ibmq-valencia* downloaded from the IBM Q system service at the time when the experiment in Sec. VI-E performed. The gate lengths of the Identity (Id) gate, U1 gate, U2 gate, and U3 gate, are 35.55556 ns, 0.0 ns, 35.55556 ns, and 71.11111 ns, respectively.

Qubit	Id gate error	U1 gate error	U2 gate error	U3 gate error	Readout error
0	0.000361	0.0	0.000361	0.000721	0.01750
1	0.000327	0.0	0.000327	0.000654	0.01500
2	0.001065	0.0	0.001065	0.002129	0.12250
3	0.000295	0.0	0.000295	0.000590	0.01500
4	0.000294	0.0	0.000294	0.000588	0.01375

**TABLE 11.** Qubit relaxation time  $T_1$ , dephasing time  $T_2$  and qubit frequency data of the 5-qubit machine *ibmq-valencia* downloaded from the IBM Q system service at the time when the experiment in Sec. VI-E performed.

Qubit	$T_1$ ( $\mu$ s)	$T_2$ ( $\mu$ s)	Frequency (GHz)
0	119.527115	86.931245	4.744506
1	112.397617	83.546189	4.650691
2	123.174284	16.005230	4.792272
3	135.673765	66.236028	4.834118
4	97.472783	103.480782	4.959371

The single-qubit gates implemented on the IBM Q systems are defined through a general single-qubit unitary gate

$$U(\theta, \phi, \lambda) = \begin{pmatrix} \cos(\theta/2) & -e^{i\lambda} \sin(\theta/2) \\ -e^{i\phi} \sin(\theta/2) & e^{i\lambda+i\phi} \cos(\theta/2) \end{pmatrix}, \quad (18)$$

with  $u1(\lambda) = U(0, 0, \lambda)$ ,  $u2(\phi, \lambda) = U(\pi/2, \phi, \lambda)$ , and  $u3(\theta, \phi, \lambda) = U(\pi/2, \phi, \lambda)$ . The single-qubit gates listed in Table 9 are defined as  $U1 = u1(\pi/2)$ ,  $U2 = u2(\pi/2, \pi/2)$ , and  $U3 = u3(\pi/2, \pi/2, \pi/2)$ , respectively. In reality, the U1 gate is done using a frame change which means it is done in software (gate time is zero) and thus its gate error is also set to zero in Table 9.

## B. DEVICE PROPERTIES FOR THE QUANTUM INFERENCE

The real IBM Q system we use in the experiments described in Sec. VI-E is the 5-qubit machine of *ibmq-valencia*. The qubit relaxation time  $T_1$ , dephasing time  $T_2$  and qubit frequency of each qubit are listed in Table 11, the single-qubit gate error, gate length and readout error data are listed in Table 10, and the two-qubit coupling and corresponding CNOT gate error and gate length data are listed in Table 7.

## ACKNOWLEDGMENT

Samuel Yen-Chi Chen would like to thank Yun-Cheng Tsai for constructive discussion and great support, and thanks Cheng-Lin Hong for helping on setting computational clusters.

## REFERENCES

- [1] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, pp. 436–444, May 2015.
- [2] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *Proc. ICLR*, 2015, pp. 1–14.
- [3] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 1–9.
- [4] A. Voulodimos, N. Doulamis, A. Doulamis, and E. Protopapadakis, "Deep learning for computer vision: A brief review," *Comput. Intell. Neurosci.*, vol. 2018, pp. 1–13, Feb. 2018.
- [5] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 3104–3112.
- [6] S.-C. Kao, C.-H.-H. Yang, P.-Y. Chen, X. Ma, and T. Krishna, "Reinforcement learning based interconnection routing for adaptive traffic optimization," in *Proc. 13th IEEE/ACM Int. Symp. Netw.-Chip*, Oct. 2019, pp. 1–2.
- [7] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hassabis, "Mastering the game of go with deep neural networks and tree search," *Nature*, vol. 529, no. 7587, pp. 484–489, Jan. 2016.
- [8] A. Borin and D. A. Abanin, "Approximating power of machine-learning ansatz for quantum many-body states," 2019, *arXiv:1901.08615*. [Online]. Available: <http://arxiv.org/abs/1901.08615>
- [9] G. Carleo, K. Choo, D. Hofmann, J. E. T. Smith, T. Westerhout, F. Alet, E. J. Davis, S. Efthymiou, I. Glasser, S.-H. Lin, M. Mauri, G. Mazzola, C. B. Mendl, E. van Nieuwenburg, O. O'Reilly, H. Théveniaut, G. Torlai, F. Vicentini, and A. Wietek, "NetKet: A machine learning toolkit for many-body quantum systems," *SoftwareX*, vol. 10, Jul. 2019, Art. no. 100311.
- [10] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová, "Machine learning and the physical sciences," *Rev. Mod. Phys.*, vol. 91, Dec. 2019, Art. no. 045002.
- [11] A. Canabarro, F. F. Fanchini, A. L. Malvezzi, R. Pereira, and R. Chaves, "Unveiling phase transitions with machine learning," *Phys. Rev. B: Condens. Matter*, vol. 100, no. 4, Jul. 2019, Art. no. 045129.
- [12] Z. An and D. L. Zhou, "Deep reinforcement learning for quantum gate control," *EPL (Europhys. Lett.)*, vol. 126, no. 6, p. 60002, Jul. 2019.
- [13] E. Flurin, L. S. Martin, S. Hacoen-Gourgy, and I. Siddiqi, "Using a recurrent neural network to reconstruct quantum dynamics of a superconducting qubit from physical observations," *Phys. Rev. X*, vol. 10, no. 1, Jan. 2020, Art. no. 011006.
- [14] P. Andreasson, J. Johansson, S. Liljestrand, and M. Granath, "Quantum error correction for the toric code using deep reinforcement learning," *Quantum*, vol. 3, p. 183, Sep. 2019.
- [15] H. Poulsen Nautrup, N. Delfosse, V. Dunjko, H. J. Briegel, and N. Friis, "Optimizing quantum error correction codes with reinforcement learning," *Quantum*, vol. 3, p. 215, Dec. 2019.
- [16] A. Cross, "The IBM q experience and QISKit open-source quantum computing software," in *Proc. APS Meeting*, Mar. 2018.
- [17] T. Lanting *et al.*, "Entanglement in a quantum annealing processor," *Phys. Rev. X*, vol. 4, May 2014, Art. no. 021041.
- [18] D. Gottesman, "Stabilizer codes and quantum error correction," 1997, *arXiv:quant-ph/9705052*. [Online]. Available: <https://arxiv.org/abs/quant-ph/9705052>
- [19] D. Gottesman, "Theory of fault-tolerant quantum computation," *Phys. Rev. A, Gen. Phys.*, vol. 57, no. 1, pp. 127–137, Jan. 1998.
- [20] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, "Quantum circuit learning," *Phys. Rev. A, Gen. Phys.*, vol. 98, no. 3, Sep. 2018, Art. no. 032309.
- [21] Y. Du, M.-H. Hsieh, T. Liu, and D. Tao, "The expressive power of parameterized quantum circuits," 2018, *arXiv:1810.11922*. [Online]. Available: <http://arxiv.org/abs/1810.11922>
- [22] J. Preskill, "Quantum computing in the NISQ era and beyond," *Quantum*, vol. 2, p. 79, Aug. 2018.
- [23] M. Schuld, A. Bocharov, K. Svore, and N. Wiebe, "Circuit-centric quantum classifiers," 2018, *arXiv:1804.00633*. [Online]. Available: <http://arxiv.org/abs/1804.00633>
- [24] V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta, "Supervised learning with quantum-enhanced feature spaces," *Nature*, vol. 567, no. 7747, pp. 209–212, Mar. 2019.



- [25] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 2672–2680.
- [26] R. S. Sutton and A. G. Barto, *Reinforcement Learning*. New York, NY, USA: Springer, 1992.
- [27] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu, "Asynchronous methods for deep reinforcement learning," in *Proc. ICML*, 2016, pp. 1928–1937.
- [28] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, Feb. 2015.
- [29] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, "OpenAI gym," 2016, *arXiv:1606.01540*. [Online]. Available: <http://arxiv.org/abs/1606.01540>
- [30] P. Gawłowicz and A. Zubov, "ns-3 meets OpenAI gym: The playground for machine learning in networking research," in *Proc. 22nd Int. ACM Conf. Modeling, Anal. Simulation Wireless Mobile Syst. (MSWIM)*, Nov. 2019, pp. 113–120.
- [31] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, "Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets," *Nature*, vol. 549, no. 7671, pp. 242–246, Sep. 2017.
- [32] E. Farhi, J. Goldstone, and S. Gutmann, "A quantum approximate optimization algorithm," 2014, *arXiv:1411.4028*. [Online]. Available: <http://arxiv.org/abs/1411.4028>
- [33] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, "The theory of variational hybrid quantum-classical algorithms," *New J. Phys.*, vol. 18, no. 2, Feb. 2016, Art. no. 023023.
- [34] E. Farhi and H. Neven, "Classification with quantum neural networks on near term processors," 2018, *arXiv:1802.06002*. [Online]. Available: <http://arxiv.org/abs/1802.06002>
- [35] F. Arute et al., "Quantum supremacy using a programmable superconducting processor," *Nature*, vol. 574, no. 7779, pp. 505–510, 2019.
- [36] A. W. Harrow and A. Montanaro, "Quantum computational supremacy," *Nature*, vol. 549, no. 7671, pp. 203–209, Sep. 2017.
- [37] M. Schuld, V. Bergholm, C. Gogolin, J. Izaac, and N. Killoran, "Evaluating analytic gradients on quantum hardware," *Phys. Rev. A, Gen. Phys.*, vol. 99, no. 3, Mar. 2019, Art. no. 032331.
- [38] M. Schuld and F. Petruccione, "Information encoding," in *Information Encoding*. Cham, Switzerland: Springer, 2018, pp. 139–171.
- [39] K. Hornik, M. Stinchcombe, and H. White, "Multilayer feedforward networks are universal approximators," *Neural Netw.*, vol. 2, no. 5, pp. 359–366, Jan. 1989.
- [40] S. Sim, P. D. Johnson, and A. Aspuru-Guzik, "Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms," *Adv. Quantum Technol.*, vol. 2, no. 12, Oct. 2019, Art. no. 1900070.
- [41] C.-H.-H. Yang, J. Qi, P.-Y. Chen, Y. Ouyang, I.-T.-D. Hung, C.-H. Lee, and X. Ma, "Enhanced adversarial strategically-timed attacks against deep reinforcement learning," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, May 2020, pp. 3407–3411.
- [42] V. Bergholm, J. Izaac, M. Schuld, C. Gogolin, M. S. Alam, S. Ahmed, J. M. Arrazola, C. Blank, A. Delgado, S. Jahangiri, K. McKiernan, J. J. Meyer, Z. Niu, A. Száva, and N. Killoran, "PennyLane: Automatic differentiation of hybrid quantum-classical computations," 2018, *arXiv:1811.04968*. [Online]. Available: <http://arxiv.org/abs/1811.04968>
- [43] S. Zagoruyko, A. Lerer, T.-Y. Lin, P. Pinheiro, S. Gross, S. Chintala, and P. Dollar, "A MultiPath network for object detection," in *Proc. Brit. Mach. Vis. Conf.*, 2016, pp. 1–14.
- [44] T. Tieleman and G. Hinton, "Lecture 6.5—RmsProp: Divide the gradient by a running average of its recent magnitude," *COURSERA, Neural Netw. Mach. Learn.*, vol. 4, no. 2, pp. 26–31, 2012.
- [45] M. Möttönen, J. J. Vartiainen, V. Bergholm, and M. M. Salomaa, "Transformation of quantum states using uniformly controlled rotations," *Quantum Inf. Comp.*, vol. 5, no. 6, pp. 467–473, 2005.
- [46] V. Dunjko and H. J. Briegel, "Machine learning & artificial intelligence in the quantum domain: A review of recent progress," *Rep. Prog. Phys.*, vol. 81, no. 7, Jun. 2018, Art. no. 074001.
- [47] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, "Quantum machine learning," *Nature*, vol. 549, pp. 195–202, Sep. 2017.
- [48] D. Dong, C. Chen, H. Li, and T. J. Tarn, "Quantum reinforcement learning," *IEEE Trans. Syst., Man, B. Cybern.*, vol. 38, no. 5, pp. 1207–1220, Oct. 2008.
- [49] V. Dunjko, J. M. Taylor, and H. J. Briegel, "Framework for learning agents in quantum environments," 2015, *arXiv:1507.08482*. [Online]. Available: <http://arxiv.org/abs/1507.08482>
- [50] V. Dunjko, J. M. Taylor, and H. J. Briegel, "Advances in quantum reinforcement learning," in *Proc. IEEE Int. Conf. Syst., Man, Cybern. (SMC)*, Oct. 2017, pp. 282–287.
- [51] H. J. Briegel and G. De las Cuevas, "Projective simulation for artificial intelligence," *Sci. Rep.*, vol. 2, no. 1, p. 400, May 2012.
- [52] J. Mautner, A. Makmal, D. Manzano, M. Tiersch, and H. J. Briegel, "Projective simulation for classical learning agents: A comprehensive investigation," *New Gener. Comput.*, vol. 33, no. 1, pp. 69–114, Jan. 2015.
- [53] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, Y. Chen, T. Lillicrap, F. Hui, L. Sifre, G. van den Driessche, T. Graepel, and D. Hassabis, "Mastering the game of go without human knowledge," *Nature*, vol. 550, no. 7676, pp. 354–359, Oct. 2017.



**SAMUEL YEN-CHI CHEN** received the B.S. degree in physics and the M.D. degree in medicine from National Taiwan University, Taipei City, Taiwan, in 2016, where he is currently pursuing the Ph.D. degree in physics. His research interests include combining quantum computing and machine learning. He was a recipient of the Theoretical High-Energy Physics Fellowship from the Chen Cheng Foundation, in 2014, the Theoretical Physics Fellowship from the National Taiwan University Center for Theoretical Physics, in 2015, and the First Prize in the Software Competition (Research Category) from Xanadu Quantum Technologies, in 2019.



**CHAO-HAN HUCK YANG** received the B.S. degree from National Taiwan University, Taipei City, Taiwan, in 2016. He is currently pursuing the Ph.D. degree from the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA, USA. His recent research interests include adversarial robustness of deep neural networks and reinforcement learning with real-world applications on speech processing, communication networks, and audio-visual processing. He is a Student Member of IEEE Society. He was a recipient of the Wallace H. Coulter Fellowship from the Georgia Institute of Technology, in 2017 and 2018. He received IEEE SPS Travel Grant for ICIP 2019, 1st Prize on the Research Track of Xanadu Quantum Software Global Competition, and DeepMind Travel Award for NeurIPS 2019. He took research interns in the Image and Visual Representation Laboratory (IVRL), the Ecole Polytechnique Federale de Lausanne (EPFL), Switzerland, in 2018, KAUST, and the Amazon Alexa Research.



**JUN QI** (Member, IEEE) received the master's degree in electrical engineering from the University of Washington, Seattle, WA, USA, in 2016, and the master's degree in electrical engineering from Tsinghua University, in 2013. He is currently pursuing the Ph.D. degree with the School of Electrical and Computer Engineering, Georgia Institute of Technology. Previously, he was a Researcher at Microsoft Research, in 2017. His research interests include non-convex optimization and statistical learning for understanding deep learning systems, tensor decomposition and applications in machine learning, deep learning, submodular optimization, and speech processing.



**PIN-YU CHEN** (Member, IEEE) received the B.S. degree (Hons.) in electrical engineering and computer science from National Chiao Tung University, Taiwan, in 2009, the M.S. degree in communication engineering from National Taiwan University, Taiwan, in 2011, and the M.A. degree in Statistics and the Ph.D. degree in electrical engineering and computer science from the University of Michigan, Ann Arbor, USA, in 2016. He is currently a Research Staff Member with the IBM

Thomas J. Watson Research Center, Yorktown Heights, NY, USA. He is also the Chief Scientist of RPI-IBM AI Research Collaboration and a PI of ongoing MIT-IBM Watson AI Lab projects. His recent research is on adversarial machine learning and robustness of neural networks. His long-term research vision is building trustworthy machine learning systems. He has published more than 20 articles on trustworthy machine learning at major AI and machine learning conferences, given tutorials at CVPR'20, ECCV'20, ICASSP'20, KDD'19 and Big Data'18, and co-organized several workshops on adversarial learning for machine learning and data mining. His research interests also include graph and network data analytics and their applications to data mining, machine learning, signal processing, and cyber security. At IBM Research, he has co-invented more than 20 U.S. patents. In 2019, he received two Outstanding Research Accomplishments on research in adversarial robustness and trusted AI, and one Research Accomplishment on research in graph learning and analysis. He was a recipient of the Chia-Lun Lo Fellowship from the University of Michigan, Ann Arbor. He received the NIPS 2017 Best Reviewer Award, and was also a recipient of the IEEE GLOBECOM 2010 GOLD Best Paper Award. He is also on the editorial board of *PLOS One*.



**XIAOLI MA** (Fellow, IEEE) is currently a Professor with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA, USA. Her research interests are in the areas of signal processing for communications and networks, signal estimation algorithms, coding theory, wireless communication theory, and sensor and ad hoc networks. Her recent research interests include intelligent wireless communication, decentralized networks, and sensor

networks. She is an IEEE Fellow for her contributes to block transmissions over wireless fading channels. She has served as the Publication Chair for the IEEE GLOBECOM 2013, local arrangements chair for IEEE GlobeSIP 2014, and a General Chair for the ACM International Conference on Underwater Networks and Systems 2015 and 2019. She is a Senior Area Editor of IEEE SIGNAL PROCESSING LETTERS and Elsevier *Digital Signal Processing*. She has been an Associate Editor of the IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS AND SIGNAL PROCESSING LETTERS.



**HSI-SHENG GOAN** received the Ph.D. degree in physics from the University of Maryland, College Park, USA, in 1999. He then worked as a Postdoctoral Research Fellow at the University of Queensland, Brisbane, Australia, from 1999 to 2001. From 2002 to 2004, he was a Senior Research Fellow awarded the Hewlett-Packard Fellowship with the Center for Quantum Computer Technology, University of New South Wales, Sydney, Australia, before he took up a faculty position at

the Department of Physics, National Taiwan University (NTU), in 2005. He is currently a Professor of physics with NTU working in the fields of Quantum Computing and Quantum Information, Quantum Control, Mesoscopic (Nano) Physics, Quantum Optics, and Quantum Optomechanical and Electromechanical Systems. He has served as a member of the Editorial Boards of several international scientific journals, such as the *International Journal of Quantum Information*, *European Physical Journal: Quantum Technology*, the *Chinese Journal of Physics*, and *Frontiers in ICT: Quantum Computing*.

• • •