

Quantum Multi-Agent Meta Reinforcement Learning

Won Joon Yun¹, Jihong Park², Joongheon Kim¹

¹ School of Electrical Engineering, Korea University, Seoul, Republic of Korea

² School of Information Technology, Deakin University, Geelong, VIC, Australia

ywjoon95@korea.ac.kr, jihong.park@deakin.edu.au, joongheon@korea.ac.kr

Abstract

Although quantum supremacy is yet to come, there has recently been an increasing interest in identifying the potential of quantum machine learning (QML) in the looming era of practical quantum computing. Motivated by this, in this article we re-design multi-agent reinforcement learning (MARL) based on the unique characteristics of quantum neural networks (QNNs) having two separate dimensions of trainable parameters: *angle* parameters affecting the output qubit states, and *pole* parameters associated with the output measurement basis. Exploiting this dyadic trainability as meta-learning capability, we propose *quantum meta MARL (QM2ARL)* that first applies angle training for meta-QNN learning, followed by pole training for few-shot or local-QNN training. To avoid overfitting, we develop an *angle-to-pole regularization* technique injecting noise into the pole domain during angle training. Furthermore, by exploiting the pole as the memory address of each trained QNN, we introduce the concept of *pole memory* allowing one to save and load trained QNNs using only two-parameter pole values. We theoretically prove the convergence of angle training under the angle-to-pole regularization, and by simulation corroborate the effectiveness of QM2ARL in achieving high reward and fast convergence, as well as of the pole memory in fast adaptation to a time-varying environment.

Introduction

Spurred by recent advances in quantum computing hardware and machine learning (ML) algorithms, quantum machine learning (QML) is closer than ever imagined. The noisy intermediate-scale quantum (NISQ) era has already been ushered in, where quantum computers run with up to a few hundred qubits (Cho 2020). Like the neural network (NN) of classical ML, the parameterized quantum circuit (PQC), also known as a quantum NN (QNN), has recently been introduced as the standard architecture for QML (Chen et al. 2020; Jerbi et al. 2021; Lockwood and Si 2020b). According to IBM’s roadmap, it is envisaged to reach the full potential of QML by around 2026 when quantum computers can run with 100k qubits (Gambetta 2022).

Motivated by this trend, recent works have started re-implementing existing ML applications using QNNs, ranging from image classification to reinforcement learning (RL)

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

tasks (Schuld et al. 2020; Chen et al. 2020). Compared to classical ML, QML is still in its infancy and too early to demonstrate quantum supremacy in accuracy and scalability due to the currently limited number of qubits. Instead, the main focus in the current research direction is to identify possible challenges and novel potential in QNN-based QML applications (Schuld and Killoran 2022).

Following this direction, in this article we aim to re-design multi-agent reinforcement learning (MARL) using QML, i.e., *quantum MARL (QMARL)*. The key new element is to leverage the novel aspect of the QNN architecture, having two separate dimensions of trainable parameters. As analogous to training a classical NN by adjusting weight parameters, the standard way of training a QNN is optimizing its circuit parameters, or equivalently the rotation of the *angle* ϕ of the parameter quantum circuit’s output qubit states that are represented on the surface of the Bloch sphere (Bloch 1946). Unlike classical ML, the QNN’s output for the loss calculation is not deterministic, but is measurable after projecting multiple observations with a projector on a Hilbert space (Nielsen and Chuang 2010). According to the quantum kernel theory (Schuld and Killoran 2019), this projector is represented as the *pole* θ of the Bloch sphere, and is tunable by rotating the pole, providing another dimension for QNN training.

Interpreting such dyadic QNN trainability as meta-learning capability (Finn, Abbeel, and Levine 2017), we propose a novel QMARL framework coined *quantum meta MARL (QM2ARL)*. As Fig. 1 visualizes, QM2ARL first trains angle ϕ for meta Q-network learning, followed by training pole θ for few-shot or local training. The latter step is much faster in that θ has only two parameters (i.e., polar and azimuthal angles in spherical coordinates. When the number of environments is limited, angle training may incur overfitting, making the meta Q-network ill-posed. To avoid this while guaranteeing convergence, we develop an *angle-to-pole regularization* method that injects random noise into the pole domain during angle training, and theoretically proves the convergence of angle training under the presence of angle-to-pole regularization.

Furthermore, it is remarkable that each locally-trained QNN in QM2ARL can be uniquely represented only using its pole deviation θ from the origin or equivalently the meta-trained QNN. Inspired by this, we introduce the concept of

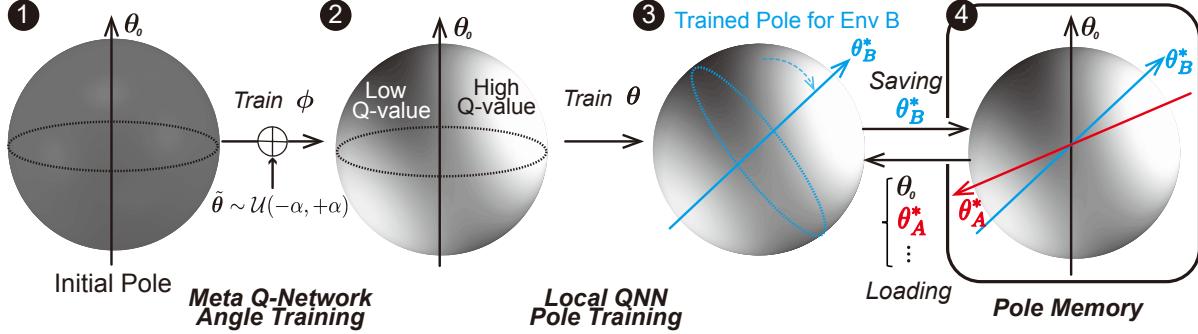


Figure 1: The key concept of quantum multi-agent meta reinforcement learning for a single qubit. A classical data (e.g., state or observation) with the Q -value is mapped onto the Bloch sphere. The lighter and the darker surfaces stand for the high and low Q -value, respectively. The first Bloch sphere presents the untrained Q -value map. From meta-QNN angle training with the noise regularizer, the parameter quantum circuit’s parameters ϕ are trained, and the output is presented as the second Bloch sphere. Then, the pole parameters θ are trained. If catastrophic remembering for adapting to multiple environments is needed, the position of the pole is initialized to a certain pole (e.g., an initial pole or optimal pole of other environments).

pole memory that can flexibly save and load the meta-trained and locally-trained QNNs only using their pole values as the memory addresses. Thanks to the two-parameter space of θ , the pole memory can store K different QNN configurations only using $2K$ parameters, regardless of the QNN sizes, i.e., ϕ dimensions. By simulation, we show the effectiveness of pole memory to cope with time-varying and cyclic environments wherein QM2ARL can swiftly adapt to a revisiting environment by loading the previous training history.

Contributions. The main contribution of this work is summarized as follows. First, we propose QM2ARL, the first QMARL framework that utilizes both the angle and pole domains of a QNN for meta RL. Second, we develop the angle-to-pole regularization technique to avoid the meta-trained QNN’s overfitting. Third, we theoretically prove the bounded convergence of meta-QNN training in the presence of the angle-to-pole regularization, which is non-trivial as the gradient variance during angle training may diverge. Fourth, we introduce the pole memory, and show its effectiveness in fast adaptation to a time-varying environment. Lastly, by simulation we validate that QM2ARL achieves higher rewards and faster convergence than several baselines, including the naïve CTDE-QMARL (Yun et al. 2022) under the MARL environments of a two-step game (Sunezag et al. 2018; Rashid et al. 2020; Son et al. 2019) and a single-hop offloading scenario (Yun et al. 2022).

Preliminaries of QMARL

QMARL Setup

Notation. The bold symbol in this paper denotes the vectorized form of the normal symbol. $|\psi\rangle$, $\phi \triangleq \{\phi_1, \dots, \phi_k, \dots, \phi_{|\phi|}\}$, and $\theta \triangleq \{\theta_1, \dots, \theta_k, \dots, \theta_{|\theta|}\}$ are defined as an entangled quantum state, the parameters of parameter quantum circuit, and the parameters of measurement, respectively. Here, ϕ_k and θ_k are the k -th entries of ϕ and θ . Moreover, \otimes , $(\cdot)^\dagger$, and $(\cdot)^T$ denote Kronecker product, complex conjugate operator, and transpose operator, re-

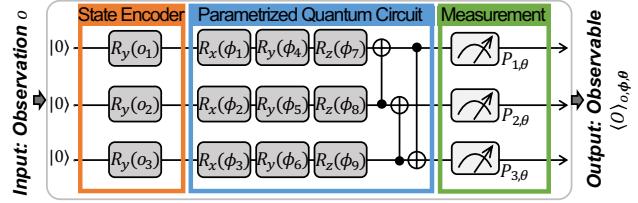


Figure 2: QNN where the PQC and the measurement have tunable parameters ϕ_k (angle) and θ (pole), respectively.

spectively. The terms “*observation*” and “*observable*” are used in this paper where the “*observation*” is information that an agent locally obtained in a multi-agent environment, whereas the “*observable*” is an operator whose property of the quantum state is measured.

Multi-agent Settings. QMARL is modeled with *decentralized partially observable Markov decision process (DecPOMDP)* (Oliehoek and Amato 2016). In DecPOMDP, the local observation, true state on current/next time-step, action, reward, and the number of agents are denoted as $o \in \mathbb{R}^{|o|}$, $s, s' \in \mathbb{R}^{|s|}$, $a \in \mathbb{R}^{|a|}$, $r(s, a)$, and N respectively. The joint observation or joint action is denoted as $\mathbf{o} = \{o^n\}_{n=1}^N$, and $\mathbf{a} = \{a\}_{n=1}^N$, respectively. Each agent has (i) QNN-based policy for execution and (ii) Q -network and target Q -network for training.

Quantum Neural Network

As illustrated in Fig. 2, QNN consists of three components: state encoding, PQC, and measurement (Killoran et al. 2019), as elaborated next.

State Encoding. The state encoder feeds the observation o into the circuit. In one-variable state encoding, the encoding process is written as $|\psi_o\rangle = \otimes_{k=1}^L [R_y(o_k) |0\rangle]$, where o_k and R_y stand for the k -th entry of observation and the rotation gate over the y -axis, respectively. As an extension to our QNN architecture, four-variable state encoding is presented

for feeding state information into the QNN-based state-value network in QMARL (Yun et al. 2022). According to (Lockwood and Si 2020b), the state encoding with fewer qubit variables (*e.g.*, one- or two-variable state encoding) will not suffer from performance degradation. In addition, a classical neural encoder is utilized to prevent dimensional reduction (Lockwood and Si 2020a). This paper uses one-variable state encoding, well-known and widely used in QML.

Parameterized Quantum Circuit. A QNN is designed to emulate the computational procedure of NNs. QNN takes encoded quantum state $|\psi_o\rangle$, which is encoded by the state encoder. Then, PQC consists of unitary gates such as rotation gates (*i.e.*, R_x , R_y , and R_z) and *controlled-NOT* ($CNOT$) gates. Each rotation gate has its trainable parameter ϕ , and the rotation gate transforms and entangles the probability amplitude of the quantum state. This process is expressed with a unitary operator $U(\phi)$, written as $|\psi_{o,\phi}\rangle = U(\phi)|\psi_o\rangle$. The quantum state of PQC’s output $|\psi_{o,\phi}\rangle$ is mapped into a 2^L -dimensional quantum space similar to NNs’ feature/latent space. This paper utilizes a basic operator block as three rotation gates and a $CNOT$ gate for each qubit (O’Brien et al. 2004), which is controlled by the adjacent neighboring qubit circularly.

Measurement. A projective measurement is described by a Hermitian operator $O \in [-1, 1]^{|\mathcal{M}_a|}$, called *observable*. According to the *Born rule*, its spectral decomposition defines the outcomes of this measurement as $O = \sum_{m \in \mathcal{M}_a} \alpha_m P_{m,\theta}$, where $P_{m,\theta}$ and α_m denote the orthogonal projections of the m -th qubit and the eigenvalue of the measured state, respectively. The orthogonal projections are written as $P_{m,\theta} = I^{\otimes m-1} \otimes M_{\theta_m} \otimes I^{\otimes L-m}$, where I and \mathcal{M}_a stand for the 2×2 identity matrix and the $|a|$ -combination of $\mathbb{N}[1, L]$, respectively. Throughout this paper, let $M_{\theta_m} = \begin{bmatrix} \cos \theta_m & -\sin \theta_m \\ -\sin \theta_m & \cos \theta_m \end{bmatrix}$ denote the measurement operator of the m -th qubit, which is decomposed into: $M_{\theta_m} = \begin{bmatrix} \cos \frac{\theta_m}{2} & \sin \frac{\theta_m}{2} \\ -\sin \frac{\theta_m}{2} & \cos \frac{\theta_m}{2} \end{bmatrix} \cdot \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \cdot \begin{bmatrix} \cos \frac{\theta_m}{2} & -\sin \frac{\theta_m}{2} \\ \sin \frac{\theta_m}{2} & \cos \frac{\theta_m}{2} \end{bmatrix}$.

The quantum state gives the outcome α_m and is projected onto the measured state $P_{m,\theta}|\psi_{o,\phi}\rangle/\sqrt{p(m)}$ with probability $p(m) = \langle\psi_{o,\phi}|P_{m,\theta}|\psi_{o,\phi}\rangle = \langle P_m \rangle_{o,\phi}$. The expectation of the observable O concerning $|\psi_{o,\phi}\rangle$ is $\mathbb{E}_{\psi_{o,\phi},\theta}[O] = \sum_m p(m)\alpha_m = \langle O \rangle_{o,\phi,\theta}$.

QNN Implementation to Reinforcement Learning

Suppose that every MARL agent has its own Q -network and policy. Motivated by (Jerbi et al. 2021), we define the Q -network as follows.

Definition 1 (Q-NETWORK). Given QNN acting on L qubits, taking input observation o , the trainable PQC parameters (*i.e.*, angle parameters) $\phi \in [-\pi, \pi]^{|\phi|}$ and their corresponding unitary transformation $U(\phi)$ produces the quantum state $|\psi_{o,\phi}\rangle = U(\phi)|\psi_o\rangle$. The quantum state $|\psi_{o,\phi}\rangle$ and trainable pole parameters $\theta \in [-\pi, \pi]^{|\theta|}$, and a scaling hyperparameter $\beta \in \mathbb{R}$ make an observable with a projection matrix P_m associated with an action a and m -th qubit. The Q -network is defined as $Q(o, a; \phi, \theta) = \beta \langle O_a \rangle_{o,\phi,\theta} = \beta \langle \psi_{o,\phi} | \sum_{m \in \mathcal{M}_a} P_{m,\theta} | \psi_{o,\phi} \rangle$.

In QMARL architecture, the policy can be expressed with the softmax function of the Q -network. Note that the policy proposed by (Jerbi et al. 2021) has utilized measurement as weighted Hermitian that are trained in a classical way. However, we consider that the Q -network/policy utilizes pole parameters, which will be further discussed.

Quantum Multi-Agent Meta Reinforcement Learning

Hereafter, we use the prefix “meta-” for the terms that are related to meta learning, *e.g.*, meta Q -network, and meta agent. We divide QMARL into two stages, *i.e.*, meta-QNN angle training and local QNN pole training. In this section, we first describe the meta-QNN angle training, where the angle parameters are trained in the angle domain $[-\pi, \pi]^{|\phi|}$. Then, we present the local-QNN pole training where the pole parameters are trained in the pole domain $[-\pi, \pi]^{|\theta|}$ with the fixed meta-QNN angle parameters. Lastly, we introduce a QM2ARL application exploiting its pole memory to cope with catastrophic forgetting.

Meta QNN Angle Training with Angle-to-Pole Regularization

Meta-QNN angle training aims to be generalized over various environments. This is, however challenged by the limited size of the Q -network and/or biased environments. Injecting noise as regularization can ameliorate this problem. One possible source of noise is the state noise (Wilson et al. 2021), which is neither controllable nor noticeably large, particularly with a small number of qubits. Alternatively, during the Q -network angle training, we consider injecting an artificial noise into the pole domain as follows.

Definition 2 (ANGLE-TO-POLE REGULARIZATION). Let an angle noise be defined as a multivariate independent random variable, $\tilde{\theta} \sim \mathcal{U}(-\alpha, +\alpha)$, where $\forall \alpha \in [0, \pi]$, where the probability density function $f_{\tilde{\theta}}(x) = \frac{1}{2\alpha}$, $\forall x \in \mathcal{U}$.

The angle-to-pole regularization injects artificial noise on the pole parameters. Thus, it impacts on the projection matrix, the meta Q -network, and finally the loss function of meta Q -network. Since QM2ARL has an independent MARL architecture, following independent deep Q -networks (IDQN) (Tampuu et al. 2017) and the *double deep Q*-networks (DDQN) (Van Hasselt, Guez, and Silver 2016), the loss function is given as a temporal difference of the meta Q -network: $\mathcal{L}(\phi; \theta + \tilde{\theta}, \mathcal{E}) = \frac{1}{n(\mathcal{E})} \sum_{\langle o, a, r, o' \rangle \in \mathcal{E}} [r + Q(o', \arg \max_{a'} a'; \phi', \theta) - Q(o, a; \phi, \theta + \tilde{\theta})]^2$. In contrast to the classical gradient descent, the quantum gradient descent of loss can be obtained by the *parameter shift rule* (Mitarai et al. 2018; Schuld et al. 2019). After calculating the loss gradient, angle parameters are updated.

Note here that while the pole parameters are not yet updated during meta Q -network angle training, the angle-to-pole regularization affects the training of angle parameters in the meta Q -network. It is unclear whether such a random regularization impact obstructs the QM2ARL convergence, calling for convergence analysis. To this end, we focus on analyzing the convergence of the meta-QNN angle training

under the existence of the angle-to-pole regularization while ignoring the measurement noise.

In the recent literature on quantum stochastic optimization, the QNN convergence has been proved under a black-box formalization (Harrow and Napp 2021) and under the existence of the measurement noise with noise-free gates (Gentini et al. 2020). Based on this, in order to focus primarily on the impact of the angle-to-pole regularization, we assume the convergence of meta Q-network angle training without the regularization as below.

Assumption 1 (CONVERGENCE WITHOUT REGULARIZATION). *Without the angle-to-pole regularization, at the t -th epoch, the angle training error of a meta Q-network with its suboptimal parameter ϕ^* is upper bounded by a constant $\epsilon_t \geq 0$, i.e., $\|\phi_t - \phi^*\| \leq \epsilon_t$.*

For the sake of mathematical amenability, we additionally consider the following assumption.

Assumption 2 (BI-LIPSCHITZ CONTINUOUS). *Meta Q-network is bi-Lipschitz continuous, i.e., $L_1 \|\sum_{j=t}^{\infty} \eta_j \nabla_{\phi} \mathcal{L}(\phi, \theta; \mathcal{E}_j)\| \leq \|\phi_t - \phi^*\| \leq L_2 \|\sum_{j=t}^{\infty} \eta_j \nabla_{\phi} \mathcal{L}(\phi, \theta; \mathcal{E}_j)\|$, where $L_2 \geq L_1 > 0$, $\eta_j > 0$ and \mathcal{E}_j are constant learning rate and an episode at epoch j , respectively.*

Then, we are ready to show the desired convergence.

Theorem 1. *With Assumptions 1 and 2, the angle training error of a meta Q-network with angle-to-pole regularization at epoch t is upper bounded by a constant $\mathbb{E}_{\tilde{\theta}} \|\tilde{\phi}_t - \phi^*\| \leq \frac{\sin \alpha}{\alpha} (\epsilon_t + \epsilon'_t)$.*

Proof Sketch. We derive the bound of the expected action value (Lemma 1), the expected derivative of action value (Lemma 2), and the variance of action value (Lemma 3) over the angle-to-pole regularization. According to Assumption 1, the term $\mathbb{E}_{\tilde{\theta}} \|\tilde{\phi}_t - \phi^*\|$ is bounded by ϵ_t . Meanwhile, by applying Lemma 3, the error term $\|\sum_{j=i}^{\infty} \frac{2\beta^2 \eta_j}{|\mathcal{E}_j|} \sum_{\tau \in \mathcal{E}_j} \nabla_{\phi} \langle O \rangle_{o, \phi, \theta}\|$ due to the regularization is upper bounded by a constant ϵ'_t , completing the proof.

Local QNN Pole Training

We design the few-shot learning in the pole domain. The reason is as follows: First of all, if the size of QNN is small, then the QNN suffers from adapting to multiple environments. The simplest way to cope with this problem is that extend the model size. However, scaling up the QNN model size is also challenging. Second, the measurement is an excellent kernel where the entangled quantum state is mapped to the observable (Havlíček et al. 2019; Schuld and Killoran 2022), while using very small parameters. The last one is that the heritage of measurement enables fast catastrophic remembering to be more intuitive, reversible, and memorable than the classical few-shot method (Farquhar and Gal 2018). Motivated by the mentioned above, we propose the *pole memory* where the pole parameters are temporally/permanently stored.

For the local-QNN pole training, and inspired by *VDN* (Sunehag et al. 2018), we assume that the joint

Algorithm 1: Training Procedure

```

1 Initialize parameters,  $\phi \leftarrow \phi_0$ ,  $\phi' \leftarrow \phi_0$ ,  $\theta \leftarrow 0$ ,  

   and  $\forall \theta^n \leftarrow 0$ ;  

2 while Meta-QNN Angle Training do  

3   Generate an episode,  

    $\mathcal{E} \leftarrow \{(\mathbf{o}_0, \mathbf{a}_0, r_1, \dots, \mathbf{o}_{T-1}, \mathbf{a}_{T-1}, r_T)\}$ , s.t.  

    $a \sim \pi_{\phi, \theta + \tilde{\theta}}$ ,  $\mathbf{a} \setminus a \sim (\text{random policy})$ ;  

4   Sampling angle noise for every step,  

    $\tilde{\theta} \sim \mathcal{U}[-\alpha, \alpha]$ ;  

5   Compute temporal difference,  $\mathcal{L}(\phi; \theta + \tilde{\theta}, \mathcal{E})$ ,  

   and its gradient  $\nabla_{\phi} \mathcal{L}(\phi; \theta + \tilde{\theta}, \mathcal{E})$ ;  

6   Update angle parameters,  

    $\phi \leftarrow \phi - \eta \nabla_{\phi} \mathcal{L}(\phi; \theta + \tilde{\theta}, \mathcal{E})$ ;  

7   if Target update period then  $\phi' \leftarrow \phi$ ;  

8 while Local-QNN Pole Training do  

9   Generate an episode  

    $\mathcal{E} \leftarrow \{(\mathbf{o}_0, \mathbf{a}_0, r_1, \dots, \mathbf{o}_{T-1}, \mathbf{a}_{T-1}, r_T)\}$ , s.t.  

    $\forall a^n \sim \pi_{\phi, \theta^n}$ ;  

10  Compute temporal difference  $\mathcal{L}_a(\Theta; \phi, \mathcal{E})$ , and  

   its gradient  $\nabla_{\Theta} \mathcal{L}_a(\Theta; \phi, \mathcal{E})$ ;  

11  Update  $\Theta \leftarrow \Theta - \eta \nabla_{\Theta} \mathcal{L}_a(\Theta; \phi, \mathcal{E})$ ;  

12  if Target update period then  $\Theta' \leftarrow \Theta$ ;
```

action-value can be expressed as the summation of local action-value across N agents, which is written as $Q_{tot}(s, a^1, a^2, \dots, a^N) \approx \frac{1}{N} \sum_{n=1}^N \tilde{Q}(o^n, a^n; \phi, \theta^n)$, where $\tilde{Q}(o^n, a^n; \phi, \theta^n)$ denotes n -th agent's local Q-network which is parametrized with the angle parameters ϕ , and n -th agent's pole parameters θ^n . Note that the local-QNN pole training only focuses on training pole parameters $\Theta \triangleq \{\theta^n\}_{n=1}^N$, and we do not consider angle noise in local-QNN pole training. We expect to maximize joint-action value by training the pole parameters (i.e., rotating the measurement axes). To maximize the cumulative returns, we design a loss function for multi-agent as $\mathcal{L}_a(\Theta; \phi, \mathcal{E}, \Theta') = \frac{1}{|\mathcal{E}|} \sum_{\tau \sim \mathcal{E}} [r + \frac{1}{N} \cdot \sum_{n=1}^N (\max_{a'^n} \tilde{Q}(o'^n, a'^n; \phi, \theta'^n) - \tilde{Q}(o^n, a^n; \phi, \theta^n))^2]$, where $\tau = \langle \mathbf{o}, \mathbf{a}, r, \mathbf{o}' \rangle$ and $\Theta' \triangleq \{\theta'^n\}_{n=1}^N$ stand for the transition sampled from environment and the target trainable pole parameters of whole agents, respectively. Note that the loss gradient of \mathcal{L}_a can be obtained by the parameter shift rule. Since the convergence analysis of \mathcal{L}_a is challenging, we show the effectiveness of the local-QNN pole training via numerical experiments.

QM2ARL Algorithm

The training algorithm is presented in Algorithm 1. All parameters are initialized. Note that the pole parameters are initialized to 0. Thus, the measurement axes of all qubits are the same as the z -axes in the initialization step. From (line 2) to (line 7), the parameters of the meta Q-network are trained with an angle noise $\tilde{\theta} \sim \mathcal{U}[-\alpha, +\alpha]$. In meta-QNN angle training, only one agent (i.e., meta agent) is trained,

Algorithm 2: Learning Procedure for Fast Remem-
bering

- 1 **Notation.** θ_p : the pole from pole memory;
- 2 **Initialization.** $\forall \phi, \phi' \leftarrow \phi_0, \forall \theta, \theta^n \leftarrow 0$;
- 3 **while** Meta-QNN Angle Training **do**
- 4 $\mathcal{E} \leftarrow \emptyset$;
- 5 **for** $env \in \text{set of environment}$ **do**
- 6 Generate an episode, \mathcal{E}_{tmp} ;
- 7 $\mathcal{E} \leftarrow \mathcal{E} \cup \mathcal{E}_{tmp}$;
- 8 $\phi \leftarrow \phi - \eta \nabla_{\phi} \mathcal{L}(\phi; \theta + \tilde{\theta}, \mathcal{E})$;
- 9 **if** Target update period **then** $\phi' \leftarrow \phi$;
- 10 **while** Training **do**
- 11 Initialize $\forall \theta^n \leftarrow \theta_p, \eta \leftarrow \eta_0$;
- 12 Local QNN Pole Training in Algorithm 1;

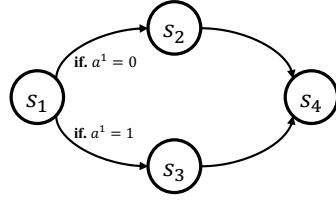
where the meta agent interacts with other agents in a multi-agent environment. The action of the meta-agent is sampled from its meta policy, *i.e.*, $\pi_{\phi, \theta + \tilde{\theta}}$, and the other agents follow other policies. The loss and its gradient are calculated and then updated. The target network is updated at a certain updating period. After the meta-QNN angle training, all pole parameters are trained to estimate the optimal joint action-value function, which is corresponding to the *local QNN pole training*. In local QNN pole training, all actions are sampled with their local Q -networks to achieve a goal of the new task. The rest of the local QNN pole training is similar to the procedure of the meta-QNN angle training except for training pole parameters with a different loss function.

Pole Memory for Fast Remembering Against Catastrophic Forgetting

The pole memory in QM2ARL can be utilized to cope with catastrophic forgetting via fast remembering (Kirkpatrick et al. 2017). To illustrate this, we first consider that a Q -network is trained in one environment. After training, a Q -network is trained in another environment, which may forget the learned experiences in the previous environment (Mirzadeh et al. 2020). To cope with such catastrophic forgetting, existing methods rely on replaying the entire previous experiences (Daniels et al. 2022), incurring a high computational cost. Alternatively, QM2ARL can first reload a meta-model in the pole memory, thereby fast remembering its previous environment with fewer iterations. Precisely, as illustrated in Algorithm 2, the meta Q -network is trained over various environments. Then, local QNN pole training is conducted for fine-tuning the meta Q -network to a specific environment. Finally, when re-encountering a forgotten environment during continual learning, the pole parameters can be re-initialized as the meta Q -network in the pole memory. This achieves fast remembering against catastrophic forgetting as we shall discuss in Fig. 6.

Numerical Experiments

The numerical experiments are conducted to investigate the following four aspects.



(a) Two-step game

$r(s_2, a^1, a^2)$	$a^2 = 0$	$a^2 = 1$
$a^1 = 0$	7	7
$a^1 = 1$	7	7

$r(s_3, a^1, a^2)$	$a^2 = 0$	$a^2 = 1$
$a^1 = 0$	0	1
$a^1 = 1$	1	8

(c) Reward of environment A

$r(s_2, a^1, a^2)$	$a^2 = 0$	$a^2 = 1$
$a^1 = 0$	7	7
$a^1 = 1$	7	7

(b) Reward of two-step game

$r(s_2, a^1, a^2)$	$a^2 = 0$	$a^2 = 1$
$a^1 = 0$	4	4
$a^1 = 1$	4	4

$r(s_3, a^1, a^2)$	$a^2 = 0$	$a^2 = 1$
$a^1 = 0$	8	1
$a^1 = 1$	1	0

(d) Reward of environment B

Figure 3: Two-step game environment.

Impact of Pole Parameters on the Meta Q -Network

To confirm how the action-value distribution of the meta Q -network is determined according to the position of the pole, we probe all pole positions of the meta Q -network when the meta-QNN angle training is finished. Then, we trace the pole position while the local QNN pole training proceeds. The experiment is conducted with the two-step game environment (Son et al. 2019), which is composed of discrete state spaces and discrete action spaces.

Fig. 4 shows the action value regarding the position of the pole. Fig. 4(a) corresponds to when the angle-to-pole regularization is not applied. In addition, the application of the angle-to-pole regularization is shown in Fig. 4(b–d). The angle noise bound was set to $\alpha = \{30^\circ, 60^\circ, 90^\circ\}$. As shown in Fig. 4(a), the action-value distribution has both high and low values. When the angle-to-pole regularization exists, we figure out that the minimum and maximum values are evenly and uniformly distributed as shown in Fig. 4(b–d). In addition, the variance of the action value is large, if angle noise exists. Therefore, the pole parameter is trained in diverse directions, and thus it can be confirmed that the momentum is large in Fig. 4. Finally, it is obvious meta Q -network is affected by angle-to-pole regularization.

Impact of Angle-to-Pole Regularization

The main proof of Theorem 1 suggests that the angle bound α of the angle-to-pole regularization plays a vital role in the convergence bound. Therefore, we conduct an experiment to investigate the role of angle-to-pole regularization. The purpose is to observe the final performance and the angle bound α . Likewise, the experiment is conducted with a two-step game environment.

To deep dive into the impact of the angle noise regularizer, we investigate the optimality corresponding to the loss function and optimal action-value function. We conduct meta-QNN angle training and local-QNN pole training with 3,000 and 20,000 iterations for the simulation,

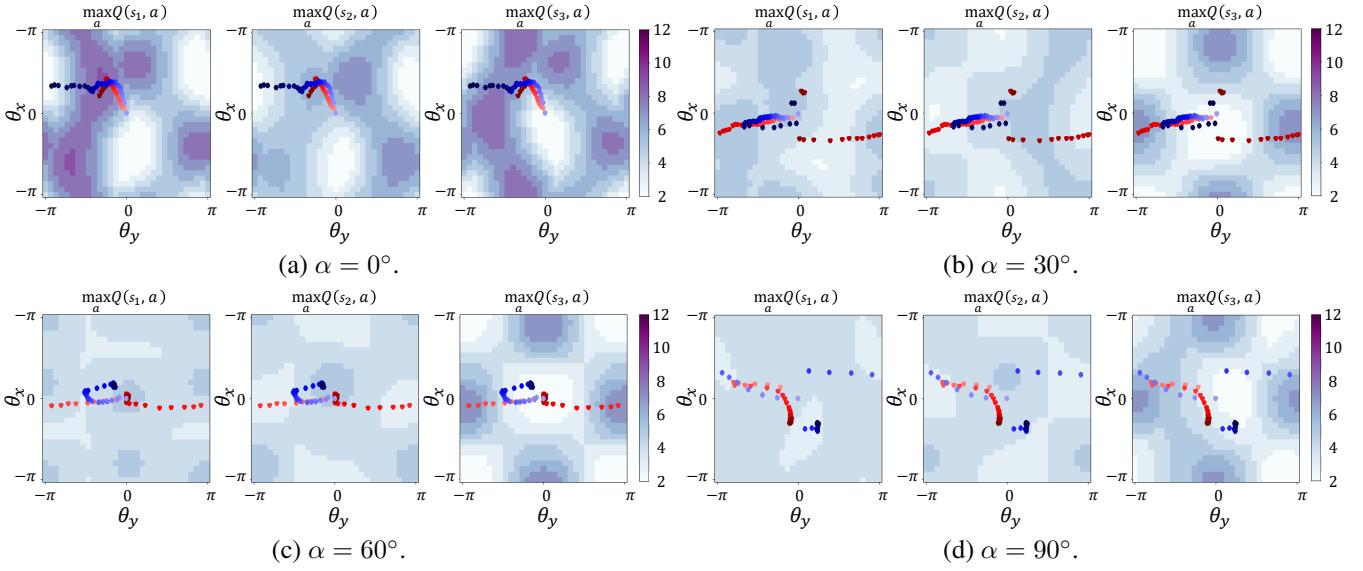


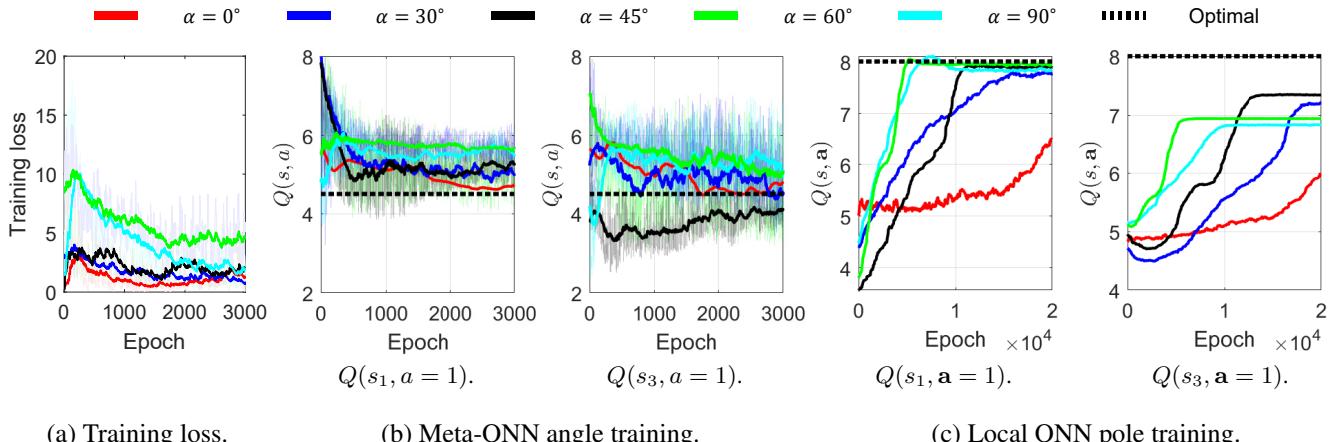
Figure 4: Action values distribution over the pole positions in meta Q -network and the trajectories of two agents' pole positions in the local-QNN pole training. The *darker* color in the grid coordinate or color bar indicates the *higher* the action-value (*i.e.*, $Q(s, a) = 12$), and the *lighter* color is the *lower* the action-value (*i.e.*, $Q(s, a) = 2$). In addition, the *red/blue dot* indicates the pole position of the first and second agent, respectively. In addition, the points of *dark/light color* indicate the pole positions at the beginning/end of local QNN pole training.

respectively. The two agents' pole parameters (*i.e.*, θ^1 and θ^2) are trained in local-QNN pole training. We test under the angle noise bound $\alpha \in \{0^\circ, 30^\circ, 45^\circ, 60^\circ, 90^\circ\}$. We set the criterion of numeric convergence when the action-values given s_1 and s_3 , stop increasing/decreasing. As shown in Fig. 5(a), the training loss is proportioned to the intensity of angle noise. Fig. 5(b)/(c) shows the numerical results of meta-QNN angle training and local QNN pole training. As shown in Fig. 5(b), the larger angle bound, the distance between the action-value of the meta Q -network and the optimal action-value is larger. As shown in Fig. 5(c), the smaller angle bound, QM2ARL converges slowly, and when the angle bound is large, it shows faster convergence. In summary, despite the angle noise regularizer making meta-QNN angle training slow convergence, the angle noise regularizer makes the local-QNN pole training fast convergence.

Effectiveness of Pole Memory in Fast Remembering

Compared to the Q -network used in the existing reinforcement learning, the structural advantages of QM2ARL exist corresponding to fast remembering. QM2ARL can return the pole position to the desired position. This structural property does not require a lot of memory costs, and it only requires a small number of qubits. Therefore, we investigate the advantage of QM2ARL corresponding to fast remembering. We consider a two-step game scenario under two different environments, Env A and Env B as shown in Fig. 3(c)/(d). To intentionally incur catastrophic forgetting, in Algorithm 2, the original reward function (*i.e.*, Env A) is changed into a wrong reward function (Env B so that

the training fails to achieve high Q -value. Then, the reward function (Env B) is rolled back to that of Env A, under which conventional training may take longer time or fail to achieve the original Q -value. The meta Q -network is trained for 5,000 epochs. Then, we conduct local-QNN pole training for 10,000 epochs per environment. To avoid catastrophic forgetting, we train the pole parameters in *Env A* (Phase I), then *Env B* (Phase II), and finally *Env A* (Phase III). From this experiment, we compare QM2ARL with pole memory (denoted as w. PM.) and without pole memory (denoted as w/o. PM.) for two angles bound $\alpha \in \{0^\circ, 30^\circ\}$. For performance comparison, we adopt optimal Q -value distance. Fig. 6 shows the results of the fast remembering scheme. In the dotted curve in Fig. 6, if the value of y -axis is closer to zero, QM2ARL achieves adaptation to the environment. In common, all frameworks shows the better adaptation to *Env A* then *Env B*. As shown in phase I of Fig. 6, the initial tangent of the optimal Q -value distance shows steep when the pole memory is utilized. In phase II, the comparison (*i.e.*, no pole memory) cannot adapt to *Env B*, while the proposed framework (denoted as $\alpha = 30^\circ$ w. PM) adapts to *Env B*. In addition, utilizing a pole memory shows faster adaptation. The result of phase III shows faster adaptation, which is similar to the results of phase I. In summary, the pole memory enables QM2ARL to achieve faster adaptation, and the proposed framework (*i.e.*, QM2ARL leveraging both angle-to-pole regularization and pole memory) shows faster and better adaptation.



(a) Training loss.

(b) Meta-QNN angle training.

(c) Local QNN pole training.

Figure 5: The results of the meta-QNN angle training process in two-step game: (a) the learning curve of QM2ARL corresponding to training loss of meta-QNN angle training process, (b) the action-value of meta Q -network given state, and (c) the joint action-value given state.

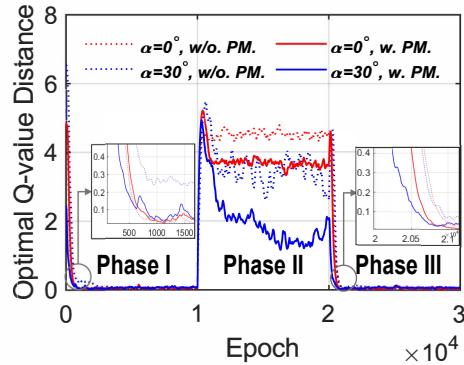


Figure 6: Fast remembering.

QM2ARL vs. QMARL and Classical MARL

We compare QM2ARL with the existing QMARL algorithms and the performance benefits compared to a classical deep Q -network based MARL with the same number of parameters (Yun et al. 2022; Sunehag et al. 2018).

An experiment has been conducted on a more difficult task than a two-step game, a *single-hop environment*. In Fig. 7(a), there was no significant difference in performance when angle noise exists (*i.e.*, $\alpha = 30^\circ$) and when it does not exist. However, in the local-QNN pole training, the performance difference between the existence of angle noise is large, as shown in *red line* and *blue line* of Fig. 7(b). When comparing *blue line* and *green line*, the convergence of the proposed scheme is faster than that of the conventional CTDE QMARL technique. When compared with the local-QNN pole training without pretraining, the performance deterioration is significant. In summary, like *single-hop environment*, the performance of the method proposed in this paper is superior to multi-agent learning and inference over a finite horizon.

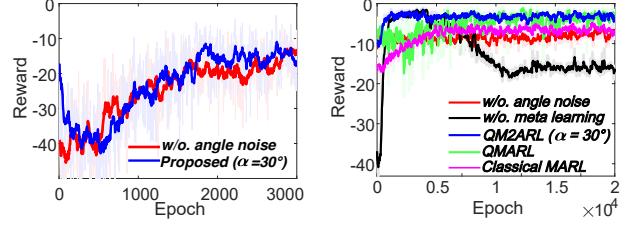


Figure 7: Performance in the more complex environment.

Concluding Remarks

Inspired from meta learning and the unique measurement nature in QML, we proposed a novel quantum MARL framework, dubbed Q2MARL. By exploiting the two trainable dimension in QML, the meta and local training processes of Q2MARL were separated in the PQC angle and its measurement pole domains, *i.e.*, meta-QNN angle training followed by local-QNN pole training. By reflecting this sequential nature of angle-to-pole training operations, we developed a new angle-to-pole regularization technique that injects noise into the pole domain during angle training. Furthermore, by exploiting the angle-pole domain separation and the small pole dimension, we introduced a concept of pole memory that can save all meta-QNN and local-QNN training outcomes in the pole domain and load them only using two parameters per each. Simulation results corroborated that Q2MARL achieves higher reward with faster convergence than an QMARL baseline and a classical MARL with the same number of parameters. The results also showed that the proposed angle-to-pole regularization is effective in generalizing the meta-QNN training, yet at the cost of compromising convergence speed. It is therefore worth optimizing this trade-off between generalization and convergence speed in future research.

Acknowledgments

This work was partly supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No.2022-0-00907, Development of AI Bots Collaboration Platform and Self-organizing AI) and National Research Foundation of Korea (NRF-2022R1A2C2004869). Jihong Park and Joongheon Kim are the corresponding authors of this paper.

References

- Bloch, F. 1946. Nuclear induction. *Physical Review*, 70(7-8): 460.
- Chen, S. Y.-C.; Yang, C.-H. H.; Qi, J.; Chen, P.-Y.; Ma, X.; and Goan, H.-S. 2020. Variational quantum circuits for deep reinforcement learning. *IEEE Access*, 8: 141007–141024.
- Cho, A. 2020. IBM promises 1000-qubit quantum computer—a milestone—by 2023. *Science*, 15.
- Daniels, Z.; Raghavan, A.; Hostetler, J.; Rahman, A.; Sur, I.; Piacentino, M.; and Divakaran, A. 2022. Model-Free Generative Replay for Lifelong Reinforcement Learning: Application to Starcraft-2. In *Proc. of the Conference on Lifelong Learning Agents (CoLLAs)*. Montreal, Canada.
- Farquhar, S.; and Gal, Y. 2018. Towards robust evaluations of continual learning. *arXiv preprint arXiv:1805.09733*.
- Finn, C.; Abbeel, P.; and Levine, S. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proc. of the International Conference on Machine Learning (ICML)*, 1126–1135. Sydney, Australia: PMLR.
- Gambetta, J. 2022. Our new 2022 development roadmap. *IBM Quantum Computing*.
- Gentini, L.; Cuccoli, A.; Pirandola, S.; Verrucchi, P.; and Banchi, L. 2020. Noise-resilient variational hybrid quantum-classical optimization. *Physical Review A*, 102(5): 052414.
- Harrow, A. W.; and Napp, J. C. 2021. Low-depth gradient measurements can improve convergence in variational hybrid quantum-classical algorithms. *Physical Review Letters*, 126(14): 140502.
- Havlíček, V.; Córcoles, A. D.; Temme, K.; Harrow, A. W.; Kandala, A.; Chow, J. M.; and Gambetta, J. M. 2019. Supervised learning with quantum-enhanced feature spaces. *Nature*, 567(7747): 209–212.
- Jerbi, S.; Gyurik, C.; Marshall, S.; Briegel, H. J.; and Dunjko, V. 2021. Variational quantum policies for reinforcement learning. In *Proc. of Neural Information Processing Systems (NeurIPS)*. Virtual.
- Killoran, N.; Bromley, T. R.; Arrazola, J. M.; Schuld, M.; Quesada, N.; and Lloyd, S. 2019. Continuous-variable quantum neural networks. *Physical Review Research*, 1(3): 033063.
- Kirkpatrick, J.; Pascanu, R.; Rabinowitz, N.; Veness, J.; Desjardins, G.; Rusu, A. A.; Milan, K.; Quan, J.; Ramalho, T.; Grabska-Barwinska, A.; Hassabis, D.; Clopath, C.; Kumaran, D.; and Hadsell, R. 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13): 3521–3526.
- Lockwood, O.; and Si, M. 2020a. Playing atari with hybrid quantum-classical reinforcement learning. In *Proc. NeurIPS Workshop on Pre-registration in Machine Learning*, 285–301. PMLR.
- Lockwood, O.; and Si, M. 2020b. Reinforcement learning with quantum variational circuit. In *Proc. of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 16, 245–251.
- Mirzadeh, S. I.; Farajtabar, M.; Pascanu, R.; and Ghasemzadeh, H. 2020. Understanding the role of training regimes in continual learning. *Advances in Neural Information Processing Systems*, 33: 7308–7320.
- Mitarai, K.; Negoro, M.; Kitagawa, M.; and Fujii, K. 2018. Quantum circuit learning. *Physical Review A*, 98(3): 032309.
- Nielsen, M. A.; and Chuang, I. L. 2010. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press.
- O’Brien, J. L.; Pryde, G.; Gilchrist, A.; James, D.; Langford, N. K.; Ralph, T.; and White, A. 2004. Quantum process tomography of a controlled-NOT gate. *Physical Review Letters*, 93(8): 080502.
- Oliehoek, F. A.; and Amato, C. 2016. *A Concise Introduction to Decentralized POMDPs*. Springer Publishing Company, Incorporated.
- Rashid, T.; Samvelyan, M.; de Witt, C. S.; Farquhar, G.; Foerster, J. N.; and Whiteson, S. 2020. Monotonic Value Function Factorisation for Deep Multi-Agent Reinforcement Learning. *Journal of Machine Learning Research*, 21: 178:1–178:51.
- Schuld, M.; Bergholm, V.; Gogolin, C.; Izaac, J.; and Killoran, N. 2019. Evaluating analytic gradients on quantum hardware. *Physical Review A*, 99(3): 032331.
- Schuld, M.; Bocharov, A.; Svore, K. M.; and Wiebe, N. 2020. Circuit-centric quantum classifiers. *Physical Review A*, 101(3): 032308.
- Schuld, M.; and Killoran, N. 2019. Quantum machine learning in feature Hilbert spaces. *Physical review letters*, 122(4): 040504.
- Schuld, M.; and Killoran, N. 2022. Is Quantum Advantage the Right Goal for Quantum Machine Learning? *PRX Quantum*, 3: 030101.
- Son, K.; Kim, D.; Kang, W. J.; Hostallero, D.; and Yi, Y. 2019. QTRAN: Learning to Factorize with Transformation for Cooperative Multi-Agent Reinforcement Learning. In Chaudhuri, K.; and Salakhutdinov, R., eds., *Proc. of the International Conference on Machine Learning (ICML)*, volume 97, 5887–5896. Long Beach, CA, USA.
- Sunehag, P.; Lever, G.; Gruslys, A.; Czarnecki, W. M.; Zambaldi, V. F.; Jaderberg, M.; Lanctot, M.; Sonnerat, N.; Leibo, J. Z.; Tuyls, K.; and Graepel, T. 2018. Value-Decomposition Networks For Cooperative Multi-Agent Learning Based On Team Reward. In *Proc. of the International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 2085–2087. Stockholm, Sweden.
- Tampuu, A.; Matiisen, T.; Kodelja, D.; Kuzovkin, I.; Korjus, K.; Aru, J.; Aru, J.; and Vicente, R. 2017. Multiagent cooperation and competition with deep reinforcement learning. *Plos one*, 12(4): e0172395.

Van Hasselt, H.; Guez, A.; and Silver, D. 2016. Deep reinforcement learning with double q-learning. In *Proc. of the AAAI Conference on Artificial Intelligence*, volume 30. Arizona, USA.

Wilson, M.; Stromswold, R.; Wudarski, F.; Hadfield, S.; Tuman, N. M.; and Rieffel, E. G. 2021. Optimizing quantum heuristics with meta-learning. *Quantum Machine Intelligence*, 3(1): 1–14.

Yun, W. J.; Kwak, Y.; Kim, J. P.; Cho, H.; Jung, S.; Park, J.; and Kim, J. 2022. Quantum Multi-Agent Reinforcement Learning via Variational Quantum Circuit Design. In *Proc. of the IEEE International Conference on Distributed Computing Systems (ICDCS)*. Bologna, Italy.