D3PG: Dirichlet DDGP for Task Partitioning and Offloading with Constrained Hybrid Action Space in Mobile Edge Computing

This paper was downloaded from TechRxiv (<https://www.techrxiv.org>).

LICENSE

CC BY 4.0

SUBMISSION DATE / POSTED DATE

14-12-2021 / 16-12-2021

CITATION

Ale, Laha; King, Scott; Zhang, Ning; Sattar, Abdul; Skandaraniyam, Janahan (2021): D3PG: Dirichlet DDGP for Task Partitioning and Offloading with Constrained Hybrid Action Space in Mobile Edge Computing. TechRxiv. Preprint. <https://doi.org/10.36227/techrxiv.17203607.v1>

DOI

[10.36227/techrxiv.17203607.v1](https://doi.org/10.36227/techrxiv.17203607.v1)

D3PG: Dirichlet DDGP for Task Partitioning and Offloading with Constrained Hybrid Action Space in Mobile Edge Computing

Laha Ale, Scott A. King, Ning Zhang, Abdul Rahman Sattar, Janahan Skandaraniyam

Abstract—Mobile Edge Computing (MEC) has been regarded as a promising paradigm to reduce service latency for data processing in Internet of Things, by provisioning computing resources at network edge. In this work, we jointly optimize the task partitioning and computational power allocation for computation offloading in a dynamic environment with multiple IoT devices and multiple edge servers. We formulate the problem as a Markov decision process with constrained hybrid action space, which cannot be well handled by existing deep reinforcement learning (DRL) algorithms. Therefore, we develop a novel Deep Reinforcement Learning called Dirichlet Deep Deterministic Policy Gradient (D3PG), which is built on Deep Deterministic Policy Gradient (DDPG) to solve the problem. The developed model can learn to solve multi-objective optimization, including maximizing the number of tasks processed before expiration and minimizing the energy cost and service latency. More importantly, D3PG can effectively deal with constrained distribution-continuous hybrid action space, where the distribution variables are for the task partitioning and offloading, while the continuous variables are for computational frequency control. Moreover, the D3PG can address many similar issues in MEC and general reinforcement learning problems. Extensive simulation results show that the proposed D3PG outperforms the state-of-art methods.

Index Terms—Mobile Edge Computing, Task Partition, Deep Reinforcement Learning, Computation Offloading, Energy efficiency, TD3, DDPG, Dirichlet

I. INTRODUCTION

Internet of Things (IoTs) [1, 2] is considered as the foundation for a wide range of applications, including self-driving cars, smart cities, and environment monitoring. Although IoTs can address small tasks with a reasonable amount of energy consumption, many computational-intensive tasks are beyond the capacity of the IoTs. Moreover, many applications, such as self-driving cars and smart factory robots, require real-time responses, and the IoTs struggle to respond to users when they require a relatively large amount of computational resources. Furthermore, most of the IoTs are extremely sensitive to energy consumption when these devices are running wireless. Since IoTs have limited computational resources and energy support, they offload computational- and delay-intensive tasks to online servers to process the tasks. However, it is challenging to offload a large number of tasks through the core network and process them on remote servers because the networks would be congested and therefore increase the delay time. Mobile Edge Computing (MEC) is proposed to address the tasks in proximity and reduce the burden of the core network [3]. Unfortunately, MEC is not a panacea for the above problem because MEC servers are equipped with

much less computational resources than central cloud servers; therefore, task offloading and scheduling optimization are vital to exploit the limited resources and improve service quality and reduce costs.

Various methods have been proposed to optimize the MEC resource usage to fully utilize the limited computational resource in MEC servers. To reduce the idle time of edge servers and respond to IoT user timely, an offloading task can be sliced into small sub-tasks, and the sub-tasks can be processed in heterogeneous edge servers [4]. Therefore, slicing task [5] into small sub-tasks and offloading to edge servers so that the limited computational power on the edge servers can be fully utilized. Conventional optimization methods (eg. CVX and MIP) [6], machine learning [7], deep learning [8], and reinforcement learning [9, 10] methods have been introduced to address computation offloading challenges. It is very challenging to adopt the aforementioned methods into MEC tasks partitioning and scheduling. First, it is considerably challenging to describe the practical MEC network into mathematical forms that can fit into conventional optimization methods. Second, machine learning and deep learning usually require labeled data to train the models, which can be extremely difficult for humans to manually compute and label slicing and scheduling datasets.

Deep reinforcement learning (DRL) methods can mitigate the issues mentioned above. To optimize the task partitioning, offloading, and computing power allocation, the existing methods typically use DRL and optimization techniques to deal with those decision variables separately, which can lead to a poor overall system performance, instead of addressing the joint optimization problem in an end-to-end manner. To jointly optimize those decision variables, we need to deal with a hybrid action space. Moreover, it is even more complex as there are some constraints on the action space, because the sum of the percentage of all sub-tasks for offloading from a given task should be equal to one. The majority of existing DRL models can only address discrete action [10] or continuous action space [11]. Several works try to address hybrid action with approximation or relaxation of continuous action space [12]. Hausknecht *et al.* [13]. However, they cannot address constrained action space in edge computing. Wu *et al.* [14] use *softmax* to capture the task partitioning actions to satisfy the constraints of the action space, which is a proportional action space, the sum of which should be one. However, *softmax* does not have an exploration mechanism to explore all the possible actions and derive the optimal

policies.

In this work, we propose a novel deep reinforcement learning approach called Dirichlet Deep Deterministic Policy Gradient (D3PG), based on Deep Deterministic Policy Gradient (DDPG), to jointly optimize the task partitioning, computation offloading, and computational frequency control. The developed model can decide to partition the tasks flexibly, offload the sub-tasks to the edge servers, and select the computational frequencies of edge servers to execute the sub-tasks. The goal of the model is to make those complex decisions based on the observation to maximize the number of tasks completed before expiring and minimize the energy consumption and time cost. The model can make the decisions and jointly optimize multiple objectives. The main contributions of this work are:

- We developed a novel D3PG model to optimize MEC resource allocation and improve service quality. The proposed model generates a distribution-continuous hybrid action space to address various issues flexibly. Specifically, each action includes a distribution formulated as a Dirichlet distribution for partitioning and offloading tasks, and continuous components for frequency control.
- A configurable optimization target is proposed to address multiple joint optimization problems. The model optimizes multiple objectives in an end-to-end manner, and it does not require further optimization like existing methods.
- We have tested the developed method by extensive simulations, and results show our method outperforms the state-of-art methods.

The rest of the paper is organized as follows. Section II investigates related works. Section III presents the system modeling and problem formulation. Section IV introduces the proposed method in details. Section V provides the simulation results and Section VI concludes this work.

II. RELATED WORKS

In the literature, many methods have been proposed for task scheduling and offloading. A joint computation offloading and system resource allocation for MEC has been formulated as a mixed-integer non-linear programming format in [15]. Then the authors transformed the non-linear programming to linear programming to reduce the complexity and tackle the challenges. Liang *et al.* [16], adopt linear-fractional programming (LFP), a generalization of linear programming (LP), and a greedy algorithm to optimize the offloading rate and energy consumption. The tasks can also be divided into sub-task processes and processed on local devices or edge servers simultaneously [17, 18] based on optimization. Similarly, Gao *et al.* [19] propose a method to find the optimal ratio of task partition into two sub-tasks for edge servers and local IoTs. He *et al.* [20] propose an optimization method to partition deep learning inference tasks and offload the partitions to edge servers to find optimal-delay for computing resource allocation. Those standard optimization methods are relatively straightforward to develop and have addressed many optimization problems in task offloading. However, MEC network

environments are far complex to describe with mathematical forms, and it is considerably challenging to extend conventional optimization to high-dimensional observations.

Machine learning and deep learning [8] models can learn from historical labeled data and predict future computational offloading so that the system can make plausible decisions for MEC. Shen *et al.* [21] surveyed machine learning methods for resource slice and planning for the next-generation network. They have also summarized various machine learning and deep learning methods adopted in computational offloading and content offloading for MEC. Lyu *et al.* [22] use stochastic gradient descent, a popular machine learning training method, to learn and partition data to offload spatially distributed edge servers; the authors argue that the proposed method can make optimal decisions for data partitioning with respect to time delay. Yang *et al.* [23] proposed a statistical machine learning method to minimize the energy consumption for edge inference [24, 25], where deep learning inference tasks are processed on MEC. Ale *et al.* [26] introduced a deep recurrent neural network to capture and predict the user requests so that they can make decisions for the content offloading and allocate resources based on the prediction. Summarily, the computation offloading problems can be formulated as a supervised classification problem and minimize the cost using deep learning [7]. The deep learning models can also adopt bandwidth allocation optimization and maximize the system utility [6]. Machine learning and deep learning methods can learn and predict the offloading decisions optimized with respect to exploiting resources and reducing the cost for the MEC. However, machine learning and deep learning require training data with labels, which require enormous effort to collect and label data. Moreover, it is challenging for humans to label data from optimization problems because it is difficult to take optimal actions in such a complex system.

To mitigate the above issues, reinforcement learning [9] and Deep Reinforcement Learning [10, 27] methods are extensively adopted for resource planning and optimization in MEC. In the reinforcement learning framework, we do not need to provide labeled data to train the models; instead, the learning agents interact with the environment (i.e., the MEC networks) to learn and find the optimal policies with respect to the objective function (e.g., minimize energy consumption). A Q-learning (a typical reinforcement learning algorithm) based method has been proposed [5] to make decisions for task offloading; the agents learn to decide whether the current tasks are offloaded to edge servers or processed on local devices to minimize the delay time. Similarly, deep Q-learning has been adopted to decide task offloading and select targeted edge servers for smart vehicles [28]. Cheng *et al.* [29] adopted the DRL method to minimize the time delay for computation offloading to Unmanned Aerial Vehicle (UAV) based edge servers and the DRL model outperforms brute-force and the greedy algorithm. Similarly, Baek *et al.* [30] proposed a Deep Q-Network (DQN) [10] variant by replacing the convolutional neural network with a recurrent neural network to control and select task offloading to edge servers; the offloading actions generated by policies to exploit the limited MEC resources and maximize number of tasks been processed before expiring.

Another DQN variant model [31] is proposed to resource allocation optimization by incorporating Bayesian learning and Long Short-Term Memory (LSTM) [32]. In [4], the authors have adopted DRL models to optimize task partitioning and scheduling for vehicular networks that allow two edge servers to process a task collaboratively. Yu *et al.* [33] proposed a DRL model to optimize the task partitioning and offloading; the model makes decisions based on the profiles of sub-tasks and chooses the local devices or edge servers to process the sequentially depended sub-tasks.

Although reinforcement learning and DRL methods are adopted to address many task offloading and resource allocation problems in MEC, the currently existing methods can only deal with relatively small action spaces and are inflexible to partition large tasks. For example, the most of the DRL models for MEC can only take binary action that decides a task to offload to the MEC server or process on local devices. Another type of DRL model can find an optimal proportional (percentage) task to offload to edge servers to process. In addition, the majority of reinforcement learning models deal with discrete action [10] or continuous action space [11]. Several works try to handle hybrid action with approximation or relaxation of continuous action space [12]. Hausknecht *et al.* [13] relaxed the action space to have to support the hybrid action. Masson *et al.* [34] handled discrete action with Q-learning and policy search for continuous action. Similarly, Khamassi *et al.* [35] use Q-learning and policy gradient to achieve the same results. Those methods assume on-policy and handle discrete and continuous actions separately. Xiong *et al.* [36] and Fu *et al.* [37] use a hierarchical structure [38] to deal with the discrete actions and generate the continuous action based on the discrete actions. Neunert *et al.* [39] provided mixed policy to handle the discrete-continuous action space. However, none of those methods can handle constrained action space well in edge computing.

III. SYSTEM MODEL AND PROBLEM FORMULATION

A. System Model

As shown in Fig. 1, there is a group of N IoT users $\mathcal{U} = \{u_1, \dots, u_N\}$ during a given time slot t . The IoT devices rely on a set of K MEC servers $\mathcal{M} = \{m_1, \dots, m_K\}$ to process tasks through computation offloading. A task from the i^{th} user can be defined as: $\Omega_i = \{\mathcal{D}_i, \mathcal{C}_i, \Delta_{max}\}$, where $\mathcal{D}_i$, $\mathcal{C}_i$, and Δ_{max} denote the data size of the task, required CPU cycles to compute the task, and maximum tolerant latency (expired time) of the task Ω_i , respectively. Further, the tasks can be partitioned into smaller sub-tasks and offloaded to different MEC servers for parallel processing. Specifically, a task Ω_i can be partitioned into several smaller sub-tasks, $\xi_i = (\xi_{i,1}, \dots, \xi_{i,j}, \dots, \xi_{i,K})$, where j is the index of the sub-task, and the number of sub-tasks is no greater than the number of MEC servers K . It takes $\delta_{i,j}$ to complete each sub-task $\xi_{i,j}$. In other words, the time cost of all the sub-tasks of Ω_i can be denoted as a vector $\delta_i = (\delta_{i,1}, \dots, \delta_{i,j}, \dots, \delta_{i,K})$, $i \leq N$, and $j \leq K$, and the time cost $\delta_{i,j}$ for sub-task $\xi_{i,j}$ can be computed by:

$$\delta_{i,j} = \delta_{i,j}^T + \delta_j^R + \delta_j^Q + \delta_{i,j}^C, \quad (1)$$

Fig. 1: System Model

where $\delta_{i,j}^T$, δ_j^Q , $\delta_{i,j}^C$, and δ_j^R , are the transmission time, queuing time, the computing time of the sub-task $\xi_{i,j}$, and the remaining running time for the task being processed at the server, respectively.

The transmission time $\delta_{i,j}^T$ can be given as below

$$\delta_{i,j}^T = \frac{D_{i,j}}{\zeta_{i,j}}, \quad (2)$$

where $D_{i,j}$ is the data size of the j^{th} sub-task, and $\zeta_{i,j}$ is the current transmission rate from the i^{th} user to the j^{th} MEC server. The transmission rate $\zeta_{i,j}$ can be given by

$$\zeta_{i,j} = \mathcal{B}_j \log_2 \left(1 + \frac{P_{i,j} h_{i,j} L_{i,j}}{N_0} \right), \quad (3)$$

where $\mathcal{B}_j$ is the bandwidth, $P_{i,j}$ is the transmission power, and $h_{i,j}$, $L_{i,j}$, N_0 are the Rayleigh fading, path loss, and noise power, respectively.

Once the sub-task $\xi_{i,j}$ arrives at the edge server after transmission, it needs to wait for processing before all the sub-tasks in the queue and the server are completed. δ_j^R is the reminding computing time of the sub-task which is being processed by the j^{th} MEC server, and it is the difference between the estimated completion time δ^C and the starting time δ_s :

$$\delta_j^R = \delta^C - \delta_s. \quad (4)$$

where $\delta^C = \frac{\mathcal{C}_k}{f_k}$, and $\mathcal{C}_k$ is the required CPU cycles to compute the k^{th} offloaded sub-task in the current edge server and f_k is the allocated computing frequency to process this sub-task. Further, assume there are J tasks or sub-tasks offloaded to the j^{th} edge server before the current sub-task is assigned to the

j^{th} edge server. Then, the waiting time in the queue δ_j^Q for the j^{th} sub-task on the j^{th} MEC server can be computed by:

$$\delta_j^Q = \sum_{k=1}^J \delta_j^C, \quad (5)$$

After all the sub-tasks before the current task are completed, the sub-task δ_j^C can be processed now. The required computation time δ_j^C for the j^{th} sub-task on the j^{th} MEC server is given by:

$$\delta_j^C = \frac{C_{i,j}}{f_{i,j}}, \quad C_{i,j} \in \mathcal{C}_i, \quad (6)$$

where $C_{i,j}$ is the required CPU cycles to compute the j^{th} sub-task of Ω_i , and f_j is the frequency of j^{th} MEC server for processing the sub-task.

A task Ω_i is considered to be completed before the corresponding deadline if all sub-tasks are completed no later than the maximum tolerant latency. In other words, if the last sub-task has been processed before the expiration time, then the task is processed successfully; otherwise, it is considered as expired and a failure to respond to the user. We can denote a positive flag (+1) to the task partitioning and offloading coordinator when the task is completed before the corresponding deadline, and set this flag 0 when it fails to timely respond to the user, as defined:

$$\Lambda = \begin{cases} +1, & \text{if } \max(\delta_i) \leq \Delta_{\max} \\ 0, & \text{otherwise.} \end{cases} \quad (7)$$

Note that the agent gets a punishment (energy cost) whenever it takes action for the partitioning and offloading; therefore, the agent gets negative feedback when it fails to respond timely to the users.

Finally, the energy consumption to process task Ω_i is given by

$$E_i = E_i^T + E_i^C. \quad (8)$$

The energy consumption due to transmission, E_i^T is the sum of the energy consumption for transmission of the sub-tasks, and is given by

$$E_{i,j}^T = \sum_{j=0}^K \delta_{i,j}^T P_{i,j} = \sum_{j=0}^K \frac{D_{i,j}}{\zeta_{i,j}} P_{i,j}. \quad (9)$$

As shown in [40, 41], the computation energy consumption can be calculated by:

$$E_{i,j}^C = \sum_{j=0}^K c(f_{i,j}^k)^2 C_{i,j}, \quad (10)$$

where $c = 10^{-26}$ and $f_{i,j}^k$ is the frequency used to compute the j^{th} sub-task.

B. Problem Formulation

Theoretically, we can define the objective function as:

$$\begin{aligned} \max_{a_i \in \pi} \quad & \sum_i^N \beta_1 \Lambda - \beta_2 (E_i^T + E_i^C) - \beta_3 \max(\delta_i) \\ \text{s.t.} \quad & \delta_j^R + \delta_j^Q + \delta_{i,j}^T + \delta_{i,j}^C \leq \Delta_{\max}, \\ & f_{i,j}^k \leq f_j^{\max}, f_{i,j}^k \in \mathcal{F}_i, \\ & a_i = \{\Phi_i, \mathcal{F}_i\}, \\ & \Phi_i = \{\phi_0, \dots, \phi_j, \dots, \phi_K\}, \\ & \sum_{j=0}^K \phi_j = 1, 0 \leq \phi_j \leq 1 \end{aligned} \quad (11)$$

where β_1 , β_2 , and β_3 are the normalization factors. The optimization objective indicates maximizing the number of processed tasks before expiring and minimizing energy consumption. Each action contains two vectors, Φ_i for task partitioning and $\mathcal{F}_i$ for frequency control. Specifically, $\Phi_i = \{\phi_0, \dots, \phi_j, \dots, \phi_K\}$ for slicing the tasks according to the edge servers, and ϕ_j denotes the percentage of the task offloaded to j^{th} server; $\mathcal{F}_i = \{f_0, \dots, f_j, \dots, f_K\}$ presents the recommended frequencies, and f_j denotes the recommended frequency for the j^{th} sub-tasks. The formulation is relatively straightforward in the mathematical definition. However, it is inflexible to balance reducing energy consumption and to increase the number of completed tasks. Therefore, we formulate the objective function so as to maximize the expected accumulated rewards given by:

$$\max_{a_i \in \pi} \mathbb{E} \left[\sum_i R_i(s_i, a_i) \right] \quad (12)$$

where s_i is the current system observation and π denotes a policy; a policy maps observation states to actions. Each action a_i taken by the coordinator and corresponding reward is defined as

$$\begin{aligned} R_i(s_i, a_i) = & \alpha w_1 \Lambda - (1 - \alpha) w_2 \log(E_i) \\ & - w_3 \log(\max(\delta_i)) + \mathcal{C}, \end{aligned} \quad (13)$$

where α is the weight that allows the network providers to adjust the reward function based on their interests; w_1, w_2 and w_3 are normalization terms to convert Λ , $\log(E_i)$, and $\log(\max(\delta_i))$ into the same scale, and $\mathcal{C}$ is a small incentive to encourage agents to maintain the stability if the MEC servers.

IV. D3PG: DIRICHLET DEEP DETERMINISTIC POLICYGRADIENT

In this section, we provide a brief introduction to DRL and introduce the developed model with extensive details.

A. Background

The DRL settings are very similar to standard reinforcement learning. The learning agents interact with the environment to learn and make decisions. For reinforcement learning, we need a description of the MEC network, which we call the environment. We assume the environment represents the MEC

network and provides an interface to the agent to interact with it. In other words, reinforcement learning agents make decisions based on the observation provided by the environment, and the decisions are optimized with respect to expected long-term rewards.

To formulate the MEC network environment as an MDP, we need to specify the components of the MDP, including state space, action space, and a reward denoted as (S, A, P, R) , and the transition function $p(s'|s, a)$ of the MDP can be given as

$$p(s'|s, a) \doteq \Pr\{s_{t+1} = s' | s_t = s, a_t = a\}. \quad (14)$$

The transition probability function shows the transition probability of transit from current state s to next state s' , where s_{t+1} is state of the $(t+1)^{th}$ time step; s_t and a_t are the current state and action of the t^{th} time step assigned with current state and action, respectively. However, the transition function of the MEC network is unknown. Therefore, we prefer to design a model-free method to overcome the challenges.

The long-term accumulated rewards are also known as the return function, and it can be given as

$$G_i \doteq R_t + \gamma R_{t+1} + \gamma^2 R_{t+2} + \dots = \sum_{k=0}^{\infty} \gamma^k R_{t+k}, \quad (15)$$

where R_t is the immediate reward and the rest of the terms denote the estimated future rewards discounted by $\gamma \in [0, 1]$. Note that the definition of $R_i(s_i, a_i)$ is almost identical with $\beta_1 \Lambda - \beta_2 (E_i^T + E_i^C) - \beta_3 \max(\delta_i)$ in the target optimization problem (Eq. 11); therefore, the optimization problem can be solved by maximizing long-term rewards (Eq. 15).

The goal of the RL models is to find optimal policies to maximize long-term expected rewards. A policy π can be considered as a function mapping states to action, and there is no policy can collect more rewards than an optimal policy π^* . Note that it is possible than we can find multiple optimal policies. To find the optimal policy π^* that maximizes the expected long-term reward, the standard RL methods store all of the possible state and action pairs in tabular data structure, and each pair of state-action has attached a expected long-term reward. Tabular data structure called Q-table is designed to hold state-action pairs and corresponded reward values. The Q-table also can be formulated as an action-value function

$$Q^*(s, a) = \max_{\pi} \mathbb{E}[R_{\tau} + \gamma G_{\tau+1} | s_{\tau} = s, a_{\tau} = a, \pi]. \quad (16)$$

Intuitively, the RL models save the learned policies, mapping states to actions, into the Q-table; the agent then searches the best actions from the table with the states. However, the Q-table can quickly get an explosion and it is incredibly challenging to search the optimal policies when dealing with high-dimensional or continuous state space. Therefore, Deep Neural Networks (DNNs) have employed to capture the high-dimensional observations and generate plausible policies, which maximize expected long-term rewards.

The essential ideal of incorporating DNNs into reinforcement learning is to employ deep learning to process the complex observation and reinforcement learning to take complex actions. Although DNNs can be considered non-linear

approximators to capture the high-dimensional states, DDNs are notoriously unstable in reinforcement learning because of noisy feedback and other attributes of reinforcement learning presented later of this section. Mnih *et al.* [10] adopted DNNs as the approximators to capture the high-dimensional states and extract the feature maps without knowing the domain knowledge, called Deep Q-Network (DQN). Moreover, they introduce a method to delay parameter updates to stabilize the learning process. Specifically, they copy the network as target network, which is frozen for updating during most of the training episodes and its parameters are updated once after every $\mathcal{N}$ episode. In DQN, the authors also adopted experience reply buffer [42] to decouple the correlation of the sequential interaction tracks. During the learning process, the agents draw training samples from the replay buffer $U(D)$ to train the model, and Temporal Difference (TD) [43] learning method is adopted to training the model. Specifically, the agents try to minimize the target value and current value given by the following loss

$$L_{\tau}(\theta_{\tau}) = \mathbb{E}_{(s,a,r,s') \sim U(D)} \left[\left(r + \gamma \max_{a'} Q(s', a'; \theta_{\tau}^{-}) - Q(s, a; \theta_{\tau}) \right)^2 \right], \quad (17)$$

where $Q(s', a'; \theta_{\tau}^{-})$ and $Q(s, a; \theta_{\tau})$ denote the target and current Q-value, respectively. The s, a, θ_{τ} are the current state, action and local network weights, respectively. Similarly, $s', a', \theta_{\tau}^{-}$ are the next state, next action and target network weights, respectively. Further, the weights θ of the learning network can be updated by

$$\theta \leftarrow \theta - \eta \nabla_{\theta_{\tau}} L(\theta_{\tau}), \quad (18)$$

and the target network weights are fixed and only update every $\mathcal{N}$ step by $\theta^{-} \leftarrow \theta$.

In this work, both the state space and action space have continuous ranges. Therefore, we have to develop a DRL model that can address continuous state space and action space; ideally, it is a policy-based model so that it does not rely on value functions to derive optimal policies. The policy-based reinforcement learning methods are easy to extend to high-dimensional or continuous action space because they do not have to reason individual actions as the objective function is parameterized policy π . In the policy-based DRL method, the models derive the optimal police directly by maximizing the expected long-term utility $\pi \approx \pi^*$ maximize expected return

$$U(\theta) = \sum_{\tau} P(\tau; \theta) R(\tau), \quad (19)$$

where $P(\tau; \theta)$ is the probability selection actions based on the policy. The learning process (policy optimization) can be implemented with many methods such as hill-climbing, genetic algorithms, and Policy Gradient Algorithms [44]. The popular ways to derive the optimal policies is to use Policy Gradient as

$$\theta \leftarrow \theta + \eta \nabla_{\theta} U(\theta) \quad (20)$$

where learning rate η decays over the time steps to avoid overshooting which an result in non-optimal policies. Policy-based methods have many advantages such as good converge

properties, easily extend to high-dimensional and continuous action space, and achieve true stochastic policy. However, policy-based RL methods also have some drawbacks. First, policy-based methods are susceptible to convergence to local optima, especially with non-linear function approximators. The issue is shared with value-based methods, but it is more challenging in the policy-based method. Second, the obtained knowledge is specific and does not always generalize well because it only captures what the agent wants to optimize the policy and includes no other information. Third, it ignores much information in the data.

B. Proposed Model

Considering the advantages and disadvantages of the value-based and policy-based DRL, we prefer to develop an Actor-Critic DRL model because it has the advantages of value-based and policy-based reinforcement learning. Specifically, we develop the model called Dirichlet Deep Deterministic Policy Gradient (D3PG), which builds on Deep Deterministic Policy Gradient (DDPG) [11]. The developed model has to address continuous action space and meet the constraints of MEC task partitioning. Specifically, a task can be partitioned as K sub-tasks, and the size of sub-tasks can be represented by $\Phi_i = \{\phi_1, \dots, \phi_j, \dots, \phi_K\}$, ϕ_j denotes what percentage of the full task is contained in the j^{th} sub-task. Thus the sum is constrained by $\sum \Phi_i = 1$, since the total percentage is one. To satisfy action space constraints, we employ the Dirichlet distribution to capture the constrained action.

Fig. 2 shows the work process of task partitioning and offloading as well as frequency control with the developed DRL model. The system has three components: the MEC network, the MEC environment, and the DRL agent. The learning DRL agent has no control over the MEC network directly; instead, the environment serves as a coordinator to bridge the MEC network and the DRL learning agent. We assume the environment presents the MEC network for the simplicity of argument, and the DRL interacts with the environment and learns from the trials.

The learning process can be considered as the following steps. First, the agent takes action given the observation from the environment. Second, the environment provides feedback and the next state to the DRL agent; the agent then stores the current interaction data into the experience replay buffer for training the model. Each record of interaction includes current state, action, reward, and next states denoted as a tuple $\langle s_t, a_t, r_t, s_{t+1} \rangle$. The DRL agent keeps interacting with the environment to generate the training data sets. Third, the agent then draws training data from the experience replay buffer to train the learning networks inside the DRL model. Each network has a backup copy called target network, and the target networks are for stabilizing the training. The details of the elements of the MDP and DRL training process are presented in the following subsections.

1) *Action*: Although the DDPG models can address many DRL challenges with continuous action space, part of the action space in this work has a further constraint. Specifically, each action has two vectors, one vector $\Phi_i =$

$\{\phi_1, \dots, \phi_j, \dots, \phi_n\}$ for partitioning the tasks into sub-tasks according to the edge servers, and the other vector $\mathcal{F}_i = \{f_0, \dots, f_j, \dots, f_n\}$, where n is the number of edge servers. In other words, the DRL specifies p_j percentage of the task to offloads to the j^{th} edge server; further, it recommends the server process the sub-task by using f_j percentage of maximum CPU frequency of the j^{th} edge server. All of the elements (sub-component of the action) are continuous, and in the range $[0, 1]$. In other words, we can define the constraints as $\phi_j \in [0, 1]$ and $f_j \in [0, 1]$, and the j^{th} edge server should not receive a sub-task when $\phi_j = 0$. Moreover, the sum of proportion of the sliced tasks has to equal 1. Therefore, a specific action a_i can be given,

$$a_i = \{\Phi_i, \mathcal{F}_i\} \quad \text{s.t.} \quad 0 \leq \phi_j \leq 1, \sum_j \phi_j = 1. \quad (21)$$

Indeed, the *softmax* function can satisfy the constraint that $\sum_j \phi_j = 1$. However, the *softmax* function does not have an exploration mechanism, which probably leads the model to a local optimum. In DRL, the learning agent has to explore the environment because the feedback (rewards and punishments) is not labeled data as supervised machine learning; it is an evaluation score of the actions and policies. Therefore, we cannot regard the feedback as the actual label data as the deep Learning training. A possible way to mitigate this problem is to add an exploration method such as $\epsilon - greedy$ by drawing Φ_i from a random distribution (e.g., uniform distribution) or add a noise vector for each action. However, those methods are unstable to explore the continuous action space. Therefore, we use a Dirichlet distribution to characterize the Φ_i , which means $\Phi_i \sim Dir(\phi)$. Dirichlet distribution can not only satisfy the constraint of the Φ_i but naturally explores the possible actions to find the optimal policies by sampling from the Dirichlet distribution. Given the random process of Dirichlet sampling, the agent has achieved stochastic policy without saving specific actions. Therefore, the $\Phi_i \sim Dir(\phi)$ is defined as

$$Dir(\phi) = \frac{1}{B(\Psi_i)} \prod_{j=1}^n \phi_j^{\psi_j-1}, \phi_j \geq 0, \quad (22)$$

$$\text{where } B(\Psi_i) = \frac{\prod_{j=1}^K \Gamma(\psi_j)}{\Gamma(\sum_{j=1}^K \psi_j)}, \psi_j > 0$$

where $\Gamma(\cdot)$ is a standard Gamma function given by

$$\Gamma(z) = \int_0^\infty x^{z-1} e^{-x} dx, \quad \Re(z) > 0 \quad (23)$$

Moreover, to meet the condition $\psi_j > 0$ of the Dirichlet function, we use an exponential to process the actor-network outputs for the slicing action $\Psi_i = e^z + \epsilon$, where z is post-activation (outputs) parts of the previous layer, and ϵ is a very small positive number to avoid zero values. Theoretically, e^z is always greater than 0; however, it could be extremely close to 0 in some cases which would raise issues in the real-world implementation, and ϵ is to avoid such implementation errors. The other parts, post-activation from the previous layer,

Fig. 2: Developed Model

Fig. 3: Actor Network.

are inputted to the Ornstein-Uhlenbeck process. As shown in Fig. 3, the results of the Dirichlet distribution and Ornstein-Uhlenbeck process are concatenated as a complete action which can be given,

$$\mu'(s_t) = \text{Dir}(\theta_t^\phi) \oplus \left[\mu(s_t | \theta_t^f) + \mathcal{N} \right] \quad (24)$$

where $\oplus$ denotes the concatenation of two vectors, and $\text{Dir}(\theta_t^\phi)$ sub-actions for task partitioning, and the rest of the elements are for the frequency control. The Dirichlet distribution and Ornstein-Uhlenbeck process, denoted $\mu(s_t | \theta_t^f) + \mathcal{N}$, can address the exploration challenges during the learning phase. Therefore, the developed model can keep exploring the environment and is unlikely stacked to non-optimal policies because the DRL model keeps exploring by sampling actions from Dirichlet distribution and adding noise to actions with Ornstein-Uhlenbeck process.

2) *State Space*: Although the MEC environment is not fully observable to the DRL agent, which means the observation is not equivalent to the state, we assume the observation is the same as the other reinforcement learning methods. The state space at time slot t is denoted as $s_t = \langle \mathcal{M}, \zeta, \Omega \rangle$, where $\mathcal{M}$ denotes the status of MEC servers, ζ is transmission rate matrix, and Ω is the set of the current tasks ready to offload. Each of the components of the state has its own sub-components as defined in section.III.

3) *Transition Probability and Reward Function*: As the standard reinforcement learning methods, we assume the environment fits in MDP. However, we do not know the transition function (Eq. 14) in the environment. Finally, the reward function for a specific time step is defined by Eq. 13.

C. Loss Function and Gradient

Constructing the loss function is one of the critical steps of the training DRL models, and we need two loss functions: one for the actor-network and another for the critic-network. However, it is unnecessary to provide an explicit form loss function for the actor-network because the actor-network is optimized with respect to the critic value. The actor-network is policy-based, and the critic network is value-based; therefore, we can consider the critic as the Q-Network as in the DQN model [10]. The critical value is the utility for the policy-based actor-network. The learning process of the critic is very similar to Q-learning and DQN. The training process of Q-learning uses Temporal Difference (TD) [43] to update the Q-values so that the agents can search policies from the Q-table. The Q-value $Q(s, a)$ updating can be accomplished with

$$Q(s, a) \leftarrow Q(s, a) + \eta \delta_t, \quad (25)$$

$$\delta_t = \left[r_t + \gamma \max_a Q(s', a') - Q(s, a) \right],$$

where η is the learning rate, δ_t is the TD error, and $\gamma \in [0, 1]$ is the discount factor of the expected feature values. Theoretically, it has been proven [9, 45] that the near optimal Q-value can be obtained by iterating the above steps until $|Q'(s, a) - Q(s, a)| < \epsilon$, where ϵ is a very small positive number.

Similarly, the action-value function of actor-critic can be given as the Bellman equation

$$Q^\pi(s_t, a_t) = \mathbb{E}_{r_t, s_{t+1} \sim E} \left[r(s_t, a_t) + \gamma \mathbb{E}_{a_{t+1} \sim \pi} [Q^\pi(s_{t+1}, a_{t+1})] \right], \quad (26)$$

where $r(s_t, a_t)$ is the immediate reward when the agent takes action a_t based on the given state s_t , and the remaining terms are estimated future values based on policy π discounted by γ . In actor-critic DRL, the actions are made by the actor-network, which is parameterized by $\mu(s|\theta)$. Assuming the policies are deterministic, we can derive the Q-value function when the actions are generated by an actor-network $\mu(s_{t+1})$ given by

$$Q^\mu(s_t, a_t) = \mathbb{E}_{r_t, s_{t+1} \sim E} \left[r(s_t, a_t) + \gamma Q^\mu(s_{t+1}, \mu(s_{t+1})) \right]. \quad (27)$$

Therefore, the loss function based on Q-learning [45] or DQN [10], and the targeted minimize function can be given,

$$L(\theta^Q) = \mathbb{E}_{s_t \sim \rho, a_t \sim \beta, r_t \sim E} \left[(Q(s_t, a_t | \theta^Q) - y_t)^2 \right], \quad (28)$$

s.t. $y_t = r(s_t, a_t) + \gamma Q(s_{t+1}, \mu(s_{t+1}) | \theta^Q),$

where ρ^β is the state transition probability given the action distribution β . Note that y_t is parameterized by the actor-network $\mu(s_{t+1})$. Further, the gradient of the loss function can be derived with the chain rule as

$$\begin{aligned} \nabla_{\theta^\mu} J &\approx \mathbb{E}_{s_t \sim \rho^\beta} \left[\nabla_{\theta^\mu} Q(s, a | \theta^Q) \Big|_{s=s_t, a=\mu(s_t | \theta^\mu)} \right] \\ &= \mathbb{E}_{s_t \sim \rho^\beta} \left[\nabla_a Q(s, a | \theta^Q) \Big|_{s=s_t, a=\mu(s_t)} \right. \\ &\quad \left. \times \nabla_{\theta^\mu} \mu(s | \theta^\mu) \Big|_{s=s_t} \right] \end{aligned} \quad (29)$$

To stabilize the training process, DDPG also requires the target networks to compute temporal differences. Therefore, y_t is obtained from the target critic-network $\theta^{Q'}$ and actor-network $\theta^{\mu'}$. The final loss function can be given

$$L(\theta^Q) = \mathbb{E}_{s_t \sim \rho^\beta, a_t \sim \beta, r_t \sim E} \left[(Q(s_t, a_t | \theta^Q) - y_t)^2 \right] \quad (30)$$

$y_t = r(s_t, a_t) + \gamma Q(s_{t+1}, \mu'(s_{t+1}) | \theta^{Q'}),$

where y_t is the Q-value computed by the target critic-network.

Algorithm 1: D3PG for Task Partitioning and Offloading and Frequency Control

Input: *epoch_no*, $\mathcal{M}$

Output: *loss*, *gradients*

//1. Initialization:

Initialize replay memory $\mathcal{D}$;

Randomly initialize θ^Q and θ^μ ;

Initialize actor-network $\mu(s | \theta^Q)$ with weights θ^μ ;

Initialize critic-network $Q(s, a | \theta^Q)$

Initialize target networks Q' and μ' weights:

$\theta^{Q'} \leftarrow \theta^Q, \theta^{\mu'} \leftarrow \theta^\mu$;

for *episode* $\leftarrow 1$ to $\mathcal{M}$ **do**

Initialize a random process $\mathcal{N}$ for action exploration;

Preprocess initial state: $S \leftarrow \psi(< x_1 >)$;

for *time step*: $\tau \leftarrow 1$ to T_{max} **do**

// 2. Generate training data:

Select action according to the current policy and exploration noise:

$a_t = \text{Dir}(\theta_t^\phi) \oplus \left[\mu(s_t | \theta_t^f) + \mathcal{N}_t \right]$;

Execute action a_t and observe reward r_t and next state S' ;

Store experience (S, A, R, S') in $\mathcal{D}$;

// 3. Learning:

Obtain random mini-batch of (s_i, a_i, r_i, s_{i+1}) from $\mathcal{D}$;

$y_i = r_i + \gamma Q'_i(s_{i+1}, \mu'(s_{i+1} | \theta^{\mu'}))$;

Update critic by minimizing the loss:

$L(\theta) = \frac{1}{N} \sum_i (y_i - Q(s, a | \theta_i^Q))^2$;

Update critics: $\theta_i \leftarrow \min_{\theta_i} L(\theta)$;

Update the actor policy using the sampled policy gradient:

$\nabla_{\theta^\mu} J \approx \frac{1}{N} \nabla_a Q(s, a | \theta^Q) \nabla_{\theta^\mu} \mu(s | \theta^\mu)$

Update the target critic network:

$\theta^{Q'} \leftarrow \tau \theta^Q + (1 - \tau) \theta^{Q'}$;

Update the target actor network:

$\theta^{\mu'} \leftarrow \tau \theta^\mu + (1 - \tau) \theta^{\mu'}$;

D. Training process

The training process is shown in Alg. 1. The algorithm has three blocks, including initialization, data generation and collection, and model training.

The first block is to initialize variables and the networks with random weights, creating an experience reply buffer, copying the networks to the target network. As mentioned before, we have two networks, the actor-network and the critic-network, and each network has a target network to stabilize the training. The experience reply buffer maintains the training data collected from the interaction with the MEC network environment.

The second block of the algorithm is to collect data by interacting with the environment. As mentioned earlier, the action consists of a Dirichlet distribution and Ornstein-Uhlenbeck process. Every interaction with the MEC generates a training sample, and each sample includes the current observation

state, reward(feedback), next state, and a termination flag. The collected datasets are stored in the experience reply buffer, which is a queue-like container. The experience reply buffer has a fixed size, and it discards the oldest data when it receives new data.

The third block is for training the networks in the model. During the training, the target policy actor has added a smooth factor ϵ with clipped range. Again, the noise is only added to the frequency control sub-actions because the rest of the sub-actions are sampled from the Dirichlet distribution. As the standard actor-critic setting, the policy is optimized with respect to the Q-values defined by the critics. Further, the target networks are updated with soft-update and delay update methods. The target network updates are delayed to reduce variance. This method is similar to the fixation method introduced in DQN; the only difference is that it updates the network more frequently than the fixation method. The soft-updates keep a significant amount of the original weights instead of completely overwriting the networks so that the model does not have to wait for a long time to update the networks to avoid high variance. The portion of weights updated to the target networks can control the factor τ .

V. SIMULATION RESULTS

In this section, we present the details of the simulation and results analysis. We adopt Numpy [46] as a tool for data preprocessing and Pytorch¹ to build the DRL models. Again, we consider the simulation has two parts, including the MEC network environment and the DRL model. The MEC network has various network entities such as edge servers and IoTs users. As we consider the MEC networks are heterogeneous, and the edge servers are configured with different computational resources; and the IoTs users frequently generate various tasks to offload. In addition, the MEC network also maintains the network properties such as the channel gain and transmit speed matrices. The entities of the MEC networks are simulated with processors. To verify our model, we compared our model (D3PG) with the existing methods, including DDPG, DDPG with *softmax* (DDPG-softmax), Twin Delayed Deep Deterministic Policy Gradient (TD3) [47] and greedy algorithm. The DRL models are implemented with Pytorch. The key parameter settings are summarized in Table. I.

TABLE I: Parameter Settings

Parameter	Value
Signal to Noise Ratio (dB)	100
Task Data Size (bits)	$[2 \times 10^5, 2 \times 10^7]$
Task Computing Size (cycles)	$[8 \times 10^6, 1 \times 10^7]$
Server Max Frequency (Hz)	$[2 \times 10^9, 8 \times 10^9]$
Number of Online Users	[10, 1000]
Number of Edge Server s	[5, 50]
Batch Size	256
Learning Rate α	5×10^{-4}
Discount Factor γ	0.9

As mentioned in previous sections, both the DDPG and TD3 models have two types of neural networks, one for taking the

Fig. 4: Reward

actions called the actor-network and the other for evaluating the actor-network called critic-networks. For both models, the number of layers and the number of neurons in the hidden layer are the same. Specifically, the actor-networks have five layers, and the number of neurons is the size of the state space, 256, 512, 256, and the size of the action space, respectively. Similarly, the critic-networks have five layers, and the number of neurons in the state space plus the size of the action space, 256, 512, 256, and 1, respectively. Note that the TD3 model has two critic-networks while the DDPG has a single critic-network only. Therefore, TD3 consumes more computational power than DDPG models.

Considering the randomness in the MEC network and DRL models, we have run five experiments and average the results. Fig. 4 shows the rewards with respect to the episodes, and the D3PG are converged to the optimal policies around 1,500 episodes. Each reward has three components: the completed number of tasks, energy consumption, and time cost; the weights of components allow the network providers to configure based on their applications and business purposes. As we can see from Fig. 4, the D3PG model can achieve better results than other models because the Dirichlet distribution captures partitioning actions to improve the policies. DDPG-softmax outperforms the original version of DDPG and TD3 because *softmax* can capture partitioning actions; however, it is highly likely to converge local optima because *softmax* does not have an exploration mechanism to explore the optimal policies. In fact, this DDPG-softmax has relatively good results because we add noise to actions as in TD3 to help *softmax* explore the partition actions. Both original DDPG and TD3 perform poorly because we have to force the partition actions to satisfy the action space constraint $\sum_j^n \phi_j = 1$, which can degrade the overall performance. The greedy algorithm does not require a learning process and can collect more rewards at early episodes; it outperforms the standard TD3 and DDPG. The D3PG and DDPG-softmax can adequately address the action space to maximize the accumulation of long-term rewards and accumulate more rewards at later episodes.

¹<https://pytorch.org/>

Fig. 5: Completed Tasks

Fig. 6: Task Completed Ratio

Note that Fig. 4 shows the TD3 model accumulates negative rewards and performs poorer than other models. However, it learns to maximize the completed number of tasks before expiring, as shown in Fig. 5. Although the DRL model is learning to maximize the expected long-term rewards, we can decompose the rewards to verify that the model can address joint optimization in an end-to-end manner. Fig. 5 shows the completed number of tasks before expiring. The edge servers can only process a minimal number of tasks at each episode at the beginning because the models take random actions that fail to allocate the resources properly. As the models interact with the MEC environment to learn and improve the policies, they can allocate the resource optimally and serve a maximum number of offloaded tasks. Fig. 6 shows the ratio of the number of tasks processed before expiring to the total tasks. Note that we set the computational cost and data size of tasks considerably large for the edge servers. Therefore, some of the tasks are even impossible to be completed before their corresponding deadlines, and the completion ratio is only for comparison purposes. The completion ratio is much higher if we reduce the size of the tasks and the computational demand. Although the greedy algorithm can collect more rewards than DDPG and TD3, the completed number of the tasks is less than the learning methods. The learning agents need to stabilize the environment and process tasks as many as possible to collect more rewards in an episode. The D3PG outperforms other methods in terms of completed tasks and task completed ratio because the constraints do not cripple the model. Moreover, the Dirichlet distribution can capture the uncertainty of the environment and explore optimal policies.

Similarly, the models can reduce the energy consumption while maintaining the number of completed tasks, as shown in Fig. 7. The models can save energy through frequency control because the energy consumption is proportional to the square of the CPU spinning frequencies. Therefore, the DRL models can find optimal frequencies to process the offloaded tasks to balance completed tasks and energy consumption. The TD3 consumes more energy than the rest of the models, and

Fig. 7: Energy Consumption

that explains why TD3 collects less rewards shown in Fig. 4. Fig. 7 shows the energy cost per task, and the get expired tasks also consume energy. We can compute wasted energy by multiple the energy consumption with the task completed ratio shown in Fig. 6. Note that the weights of the energy and other terms affect how the model optimizes the energy cost, and the network provider can adjust the weights based on their demands. Although D3PG consumes more energy than DDPG-softmax, they consume almost the same amount of energy to process each task as shown in Fig. 8, and the D3PG model can save more energy than DDPG and TD models. Interestingly, the greedy algorithm achieves the best results for tasks completion ratio, but it cannot stabilize the edge servers and plan the resource for long-term as shown in Fig. 10.

The models also can reduce the time cost while maximizing the number of completed tasks and save energy, as shown in Fig. 9. Although the reward function has contained the number of completed tasks before expiring, reduce the time cost can improve the user experience. As we can see from the Fig. 9, the average time cost decrease as the models converge to the

Fig. 8: Energy to Tasks

Fig. 9: Average Time Cost

near-optimal policies. Fig. 9 shows that the D3PG model saves more time than other models, which improves the quality of service. The DDPG and DDPG-softmax models do not learn to reduce the time cost because the weights of time consumption are relatively small. Note that both energy consumption and average time cost are for the total cost of the completed and expired tasks.

Fig. 10 shows the stability of models, which is measured by the number of time steps of each epoch; the stability is measured by the number of steps that the MEC servers can persist in each episode. The task partitioning and offloading for the MEC can be considered continuous tasks from a reinforcement learning perspective; there is no endpoint of the offloading unless the servers are overloaded or crashed. However, for the simplicity of training, we formulate the task partitioning and offloading as episodic reinforcement learning task; therefore, we set the episodes to end when one of the edge servers is overloaded. In this simulation, we also add the external trigger to terminate the episodes, and the episode is forced to terminate when the number of time steps greater

Fig. 10: Stability

than a threshold value, set as 1,000. Again, we run training the models for five times and average the results. As we can see from Fig. 10 TD3 and D3PG models can reach nearly 1,000 epochs, which indicates most of the epochs are stopped by the simulator, and they can maintain a stable MEC network. The greedy algorithm only chooses the action that maximizes the current reward for each time step but does not plan resources for the long term; therefore, the edge servers are easily overloaded under the greedy algorithm control.

VI. CONCLUSIONS

In this work, we have studied task partitioning and computation offloading in a dynamic environment with multiple IoT devices and multiple edge servers. We developed an end-to-end DRL method to partition and offload tasks and allocate the edge servers' computing power to achieve joint optimization of expected long-term rewards. The model is optimized to maximize the completed tasks before the deadlines, minimizing energy consumption and simultaneously minimizing the time cost. In order to deal with the constrained hybrid action space, we proposed a novel DRL model, namely D3GP, by integrating the Dirichlet distribution into DDGP to make decisions for task partitioning and the Ornstein-Uhlenbeck process for frequency control. The developed model has been verified with extensive simulations and comparisons with existing methods, and the results show that our model outperforms state-of-the-art DRL models. For future work, we will study the optimization for task partitioning and offloading based on the sub-tasks that have dependency relations and task priorities.

REFERENCES

- [1] X. Li and L. Da Xu, "A review of internet of things—resource allocation," *IEEE Internet of Things Journal*, pp. 1–1, 2020.
- [2] N. Zhang, X. Fang, Y. Wang, S. Wu, H. Wu, D. Kar, and H. Zhang, "Physical-layer authentication for internet of things via wfrt-based gaussian tag embedding," *IEEE Internet of Things Journal*, vol. 7, no. 9, pp. 9001–9010, 2020.
- [3] Y. Chen, N. Zhang, Y. Zhang, X. Chen, W. Wu, and X. Shen, "Energy efficient dynamic offloading in mobile edge computing for internet of things," *IEEE Transactions on Cloud Computing*, 2019.

- [4] M. Li, J. Gao, L. Zhao, and X. Shen, "Deep reinforcement learning for collaborative edge computing in vehicular networks," *IEEE Transactions on Cognitive Communications and Networking*, vol. 6, no. 4, pp. 1122–1135, 2020.
- [5] S. Wang, M. Chen, X. Liu, and C. Yin, "Task and resource allocation in mobile edge computing: An improved reinforcement learning approach," in *2019 IEEE Globecom Workshops (GC Wkshps)*, 2019, pp. 1–6.
- [6] H. Wu, Z. Zhang, C. Guan, K. Wolter, and M. Xu, "Collaborate edge and cloud computing with distributed deep learning for smart city internet of things," *IEEE Internet of Things Journal*, vol. 7, no. 9, pp. 8099–8110, 2020.
- [7] S. Yu, X. Wang, and R. Langar, "Computation offloading for mobile edge computing: A deep learning approach," in *2017 IEEE 28th Annual International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC)*, 2017, pp. 1–6.
- [8] Y. Lecun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [9] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. Cambridge, MA, USA: A Bradford Book, 2018.
- [10] V. Mnih, K. Kavukcuoglu, D. Silver, and et al, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, Feb. 2015.
- [11] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning," in *4th International Conference on Learning Representations, ICLR 2016*, 2016.
- [12] L. Ale, N. Zhang, X. Fang, X. Chen, S. Wu, and L. Li, "Delay-aware and energy-efficient computation offloading in mobile-edge computing using deep reinforcement learning," *IEEE Transactions on Cognitive Communications and Networking*, vol. 7, no. 3, pp. 881–892, 2021.
- [13] M. Hausknecht and P. Stone, "Deep reinforcement learning in parameterized action space," in *4th International Conference on Learning Representations, ICLR 2016*, 2016, pp. 1–12.
- [14] W. Wu, N. Chen, C. Zhou, M. Li, X. Shen, W. Zhuang, and X. Li, "Dynamic ran slicing for service-oriented vehicular networks via constrained learning," *IEEE Journal on Selected Areas in Communications*, vol. 39, no. 7, pp. 2076–2089, 2021.
- [15] S. Bi, L. Huang, and Y. A. Zhang, "Joint optimization of service caching placement and computation offloading in mobile edge computing systems," *IEEE Transactions on Wireless Communications*, vol. 19, no. 7, pp. 4947–4963, 2020.
- [16] Z. Liang, Y. Liu, T. Lok, and K. Huang, "Multiuser computation offloading and downloading for edge computing with virtualization," *IEEE Transactions on Wireless Communications*, vol. 18, no. 9, pp. 4298–4311, 2019.
- [17] J. Liu and Q. Zhang, "Adaptive task partitioning at local device or remote edge server for offloading in mec," in *2020 IEEE Wireless Communications and Networking Conference (WCNC)*, 2020, pp. 1–6.
- [18] C. Fiandrino, N. Allio, D. Kliazovich, P. Giaccone, and P. Bouvry, "Profiling performance of application partitioning for wearable devices in mobile cloud and fog computing," *IEEE Access*, vol. 7, pp. 12 156–12 166, 2019.
- [19] J. Cao, L. Yang, and J. Cao, "Revisiting computation partitioning in future 5g-based edge computing environments," *IEEE Internet of Things Journal*, vol. 6, no. 2, pp. 2427–2438, 2019.
- [20] W. He, S. Guo, S. Guo, X. Qiu, and F. Qi, "Joint dnn partition deployment and resource allocation for delay-sensitive deep learning inference in iot," *IEEE Internet of Things Journal*, vol. 7, no. 10, pp. 9241–9254, 2020.
- [21] X. Shen, J. Gao, W. Wu, K. Lyu, M. Li, W. Zhuang, X. Li, and J. Rao, "Ai-assisted network-slicing based next-generation wireless networks," *IEEE Open Journal of Vehicular Technology*, vol. 1, pp. 45–66, 2020.
- [22] X. Lyu, C. Ren, W. Ni, H. Tian, R. P. Liu, and E. Dutkiewicz, "Optimal online data partitioning for geo-distributed machine learning in edge of wireless networks," *IEEE Journal on Selected Areas in Communications*, vol. 37, no. 10, pp. 2393–2406, 2019.
- [23] K. Yang, Y. Shi, W. Yu, and Z. Ding, "Energy-efficient processing and robust wireless cooperative transmission for edge inference," *IEEE Internet of Things Journal*, vol. 7, no. 10, pp. 9456–9470, 2020.
- [24] X. Xu, Y. Ding, S. X. Hu, M. Niemier, J. Cong, Y. Hu, and Y. Shi, "Scaling for edge inference of deep neural networks," *Nature Electronics*, vol. 1, no. 4, pp. 216–222, 2018.
- [25] C. Xu, J. Ren, L. She, Y. Zhang, Z. Qin, and K. Ren, "Edgesanitizer: Locally differentially private deep inference at the edge for mobile data analytics," *IEEE Internet of Things Journal*, vol. 6, no. 3, pp. 5140–5151, 2019.
- [26] L. Ale, N. Zhang, H. Wu, D. Chen, and T. Han, "Online proactive caching in mobile edge computing using bidirectional deep recurrent neural network," *IEEE Internet of Things Journal*, vol. 6, no. 3, pp. 5520–5530, 2019.
- [27] L. Ale, N. Zhang, X. Fang, X. Chen, S. Wu, and L. Li, "Delay-aware and energy-efficient computation offloading in mobile edge computing using deep reinforcement learning," *IEEE Transactions on Cognitive Communications and Networking*, pp. 1–1, 2021.
- [28] K. Zhang, Y. Zhu, S. Leng, Y. He, S. Maharjan, and Y. Zhang, "Deep learning empowered task offloading for mobile edge computing in urban informatics," *IEEE Internet of Things Journal*, vol. 6, no. 5, pp. 7635–7647, 2019.
- [29] N. Cheng, F. Lyu, W. Quan, C. Zhou, H. He, W. Shi, and X. Shen, "Space/aerial-assisted computing offloading for iot applications: A learning-based approach," *IEEE Journal on Selected Areas in Communications*, vol. 37, no. 5, pp. 1117–1129, 2019.
- [30] J. Baek and G. Kaddoum, "Heterogeneous task offloading and resource allocations via deep recurrent reinforcement learning in partial observable multi-fog networks," *IEEE Internet of Things Journal*, pp. 1–1, 2020.
- [31] Z. Yang, Y. Liu, Y. Chen, and N. Al-Dhahir, "Cache-aided noma mobile edge computing: A reinforcement learning approach," *IEEE Transactions on Wireless Communications*, vol. 19, no. 10, pp. 6899–6915, 2020.
- [32] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [33] S. Yu, X. Chen, Z. Zhou, X. Gong, and D. Wu, "When deep reinforcement learning meets federated learning: Intelligent multi-timescale resource management for multi-access edge computing in 5g ultra dense network," *IEEE Internet of Things Journal*, pp. 1–1, 2020.
- [34] W. Masson, P. Ranchod, and G. Konidaris, "Reinforcement learning with parameterized actions," *30th AAAI Conference on Artificial Intelligence, AAAI 2016*, pp. 1934–1940, 2016.
- [35] M. Khamassi, G. Velentzas, T. Tsitsimis, and C. Tzafestas, "Active exploration and parameterized reinforcement learning applied to a simulated human-robot interaction task," in *2017 First IEEE International Conference on Robotic Computing (IRC)*, 2017, pp. 28–35.
- [36] J. Xiong, Q. Wang, Z. Yang, P. Sun, L. Han, Y. Zheng, H. Fu, T. Zhang, J. Liu, and H. Liu, "Parameterized deep q-networks learning: Reinforcement learning with discrete-continuous hybrid action space," *CoRR*, vol. abs/1810.06394, 2018.
- [37] H. Fu, H. Tang, J. Hao, Z. Lei, Y. Chen, and C. Fan, "Deep multi-agent reinforcement learning with discrete-continuous hybrid action spaces," *CoRR*, vol. abs/1903.04959, 2019.
- [38] S. J. A. Raza and M. Lin, "Constructive policy: Reinforcement learning approach for connected multi-agent systems," in *2019 IEEE 15th International Conference on Automation Science and Engineering (CASE)*, 2019, pp. 257–262.
- [39] M. Neunert, A. Abdolmaleki, M. Wulfmeier, T. Lampe, J. T. Springenberg, R. Hafner, F. Romano, J. Buchli, N. Heess, and M. A. Riedmiller, "Continuous-discrete reinforcement learning for hybrid control in robotics," *CoRR*, vol. abs/2001.00449, 2020.
- [40] Y. Wang, M. Sheng, X. Wang, L. Wang, and J. Li, "Mobile-edge computing: Partial computation offloading using dynamic voltage scaling," *IEEE Transactions on Communications*, vol. 64, no. 10, pp. 4268–4282, 2016.
- [41] S. Guo, J. Liu, Y. Yang, B. Xiao, and Z. Li, "Energy-efficient dynamic computation offloading and cooperative task scheduling in mobile cloud computing," *IEEE Transactions on Mobile Computing*, vol. 18, no. 2, pp. 319–333, 2019.
- [42] L.-J. Lin, "Self-improving reactive agents based on reinforcement learning, planning and teaching," *Machine Learning*, vol. 8, no. 3–4, p. 293–321, May 1992.
- [43] R. S. Sutton, "Learning to predict by the methods of temporal differences," *Machine Learning*, vol. 3, no. 1, p. 9–44, Aug. 1988.
- [44] D. Silver, G. Lever, N. Heess, T. Degris, D. Wierstra, and M. Riedmiller, "Deterministic policy gradient algorithms," in *Proceedings of the 31st International Conference on Machine Learning*, vol. 32, no. 1, Beijing, China, 22–24 Jun 2014, pp. 387–395.
- [45] C. Watkins and P. Dayan, "Q-learning," *Machine Learning*, vol. 8, no. 3, pp. 279–292, May 1992.
- [46] C. R. Harris, K. J. Millman, and et al., "Array programming with numpy," *Nature*, vol. 521, no. 585, p. 357–362, 2020.
- [47] S. Fujimoto, H. Van Hoof, and D. Meger, "Addressing Function Approximation Error in Actor-Critic Methods," in *35th International Conference on Machine Learning, ICML 2018*, vol. 4, 2018, pp. 2587–2601.

- [48] K. Wang, Z. Hu, Q. Ai, Y. Zhong, J. Yu, P. Zhou, L. Chen, and H. Shin, "Joint offloading and charge cost minimization in mobile edge computing," *IEEE Open Journal of the Communications Society*, vol. 1, pp. 205–216, 2020.