

Received 10 November 2022; revised 14 March 2023 and 6 June 2023; accepted 7 June 2023.
 Date of publication 14 June 2023; date of current version 12 July 2023.

The associate editor coordinating the review of this article and approving it for publication was N. H. Tran.

Digital Object Identifier 10.1109/TMLCN.2023.3285543

Reinforcement Learning With Non-Cumulative Objective

WEI CUI¹ (Graduate Student Member, IEEE), AND WEI YU¹ (Fellow, IEEE)

The Edward S. Rogers Sr. Department of Electrical and Computer Engineering, University of Toronto, Toronto, ON M5S 3G4, Canada

CORRESPONDING AUTHOR: W. CUI (cuiwei2@ece.utoronto.ca)

This work was supported by the Natural Sciences and Engineering Research Council (NSERC) of Canada through the Canada Research Chairs Program.

ABSTRACT In reinforcement learning, the objective is almost always defined as a *cumulative* function over the rewards along the process. However, there are many optimal control and reinforcement learning problems in various application fields, especially in communications and networking, where the objectives are not naturally expressed as summations of the rewards. In this paper, we recognize the prevalence of non-cumulative objectives in various problems, and propose a modification to existing algorithms for optimizing such objectives. Specifically, we dive into the fundamental building block for many optimal control and reinforcement learning algorithms: the Bellman optimality equation. To optimize a non-cumulative objective, we replace the original *summation operation* in the Bellman update rule with a generalized operation corresponding to the objective. Furthermore, we provide sufficient conditions on the form of the generalized operation as well as assumptions on the Markov decision process under which the globally optimal convergence of the generalized Bellman updates can be guaranteed. We demonstrate the idea experimentally with the *bottleneck* objective, i.e., the objectives determined by the *minimum* reward along the process, on classical optimal control and reinforcement learning tasks, as well as on two network routing problems on maximizing the flow rates.

INDEX TERMS Reinforcement learning, optimal control, Markov decision process, wireless network, routing.

I. INTRODUCTION

IN REINFORCEMENT learning (RL), an agent performs a sequence of actions to optimize a certain objective, over an environment modeled as a Markov decision process (MDP) [1]. The objective value is determined by the collection of intermediate rewards the agent receives until the MDP is terminated (or an absorbing state is reached). In most of the literature, the objective is defined as the summation of these intermediate rewards, which corresponds to the *summation operation* in the Bellman optimality equation [2] when computing the value function. Such *cumulative* objectives indeed capture the ultimate goals for many problems, such as Atari games [3], stock trading [4], advertisement placements [5], and so on. Nonetheless, there are many problems with objectives that do not translate to summations of rewards.

Specifically, in the field of wireless communications, there are many system optimization problems that can be formulated and decomposed into sequences of optimization decisions, whose global objectives cannot be readily expressed as summations of rewards from individual

optimization decisions. Examples of such problems include but are not limited to *max-min* optimizations in routing and resources allocation [6], [7], [8], [9], *harmonic mean* maximization for traffic engineering [10] and for transmission system optimization [11], [12], the *proportional fairness* optimizations for wireless communications [13], [14], [15], and so on. In this paper, we recognize the prevalence of problems with non-cumulative objectives, and propose modifications to many existing optimal control and RL algorithms for optimizing such objectives¹.

In the optimal control or reinforcement learning literature, one class of problems with non-cumulative objectives are the problems where only terminal states matter, such as the game of *Go* [16], [17] or *Chess* [18]. Researchers managed to cast the objectives into summations of rewards, by assigning every reward a zero value except for the terminal reward. Problems seeking fast task completions form

¹The code for this paper is available at: https://github.com/willtop/Reinforcement_Learning_With_Non-Cumulative_Objective

another class of examples, such as maze navigation or the mountain-car control task [19]. Researchers cast the objectives as cumulative rewards by assigning a penalty for each action the agent takes before reaching the destination [20]. There are also researches on objectives that are not easily cast into summations, such as the objectives as the *average reward* [21], [22], [23], [24], [25], [26], [27]. To optimize the average reward, besides computing the summation of rewards, the number of steps is either tracked explicitly [22], or taken to the limit at infinity (for cyclic non-terminating MDPs) [23], [24]. Regardless, the summation operation in the Bellman optimality equation remains in these proposed algorithms. There have been two works [28], [29] exploring maximum-reward objectives, with applications on financial derivatives and medicine design. These works recognize the possibility of modifying the Bellman optimality equation, however their scopes are restricted to the maximum-reward objective formulation, instead of generalizing to a larger class of objective functions or proposing universal conditions for convergence. Furthermore, for MDPs whose state transition is a stochastic function of the input, the convergence to the global optimal policy cannot be guaranteed for the approach in [28] and [29].

In this paper, we generalize the optimal control and reinforcement learning objectives to a variety of functions over the intermediate rewards. To optimize the generalized objectives, we exploit the flexibility in the Bellman optimality equation and modify it accordingly to the generalized objective functions. Specifically, we replace the summation operation in the Bellman optimality equation by new operations catering to the non-cumulative objective functions. Through this approach, we can readily adapt the existing optimal control or reinforcement learning algorithms to optimizing non-cumulative objectives, without needing to re-engineer a new set of artificial rewards just to cast the objectives into a summation of rewards. Furthermore, we provide the theoretical analysis on the generalized Bellman updates, and propose sufficient conditions on the form of the new operation as well as the assumptions on the MDP under which the global optimality of the converged value function and the corresponding greedy policy can be guaranteed.

By expanding the possibilities of the objective functions, we are now able to solve problems with objectives that are intrinsically non-cumulative. For experiments, we focus on the *bottleneck* objective: the objective as the minimum reward of all intermediate rewards. To optimize bottleneck objectives, we replace the summation operation in the Bellman optimality equation by the minimization operation, and apply the generalized Bellman update rule to learn the value function. In numerical simulations, we first re-formulate two classical reinforcement learning problems: the *CartPole* problem [30] and the Atari game *Breakout*, with bottleneck objectives. Through optimizing these problems with the proposed generalized Bellman updates, we obtain competitive

performances by policies with different strategies from the classical solutions.

We further experiment on two network communication applications with bottleneck objectives: the problem of finding the single-path maximum flow on a directed graph as an optimal control task, as well as joint routing and spectrum access over a wireless ad hoc network as a reinforcement learning problem. The proposed approach achieves excellent performances on both problems that are otherwise difficult to solve using the conventional formulation and learning algorithms. Specifically, for the wireless ad hoc network problem, a prior work [6] has explored the Monte-Carlo estimation approach for learning the value function. In contrast, the proposed generalized update rule allows for the adaptation of the highly efficient temporal difference learning technique [31] to the generalized objective formulation, which results in noticeably faster and more stable learning progress. Furthermore, as the wireless ad hoc network problem is essentially a multi-agent reinforcement learning (MARL) problem, the results obtained also suggest that the proposed approach is readily compatible and effective under the multi-agent reinforcement learning setting.

The rest of the paper is organized as follows. In Section II, we introduce the general problem description on optimizing non-cumulative objectives, as well as several examples where non-cumulative objectives are applicable. In Section III, we formally propose the method of the generalized Bellman update rules, and provide theoretical convergence and optimality analysis. We provide the detailed problem formulations on several example applications, and elaborate on how the proposed generalizations can be applied to optimizing such specific applications in Section IV, followed by the numerical simulations and analysis of the results in Section V. Lastly, we draw conclusions in Section VI.

II. GENERALIZED OPTIMAL CONTROL & REINFORCEMENT LEARNING FORMULATION

A. CONVENTIONAL FORMULATION

Let \mathcal{S} and \mathcal{A} denote the state space and the action space of an MDP. At time step t , the agent observes a state $s_t \in \mathcal{S}$, executes an action $a_t \in \mathcal{A}$, and receives a reward $r_t \in \mathcal{R}$ while transiting to the next state $s_{t+1} \in \mathcal{S}$. We use $\{p_{R_t|S_t, A_t}(r_t|s_t, a_t)\}_{t=1,2,\dots}$ and $\{p_{S_{t+1}|S_t, A_t}(s_{t+1}|s_t, a_t)\}_{t=1,2,\dots}$ to denote the reward distribution and the state transition distribution of the MDP, but often omit the subscripts for notational simplicity, e.g., as in $\{p(r_t|s_t, a_t)\}_{t=1,2,\dots}$ and $\{p(s_{t+1}|s_t, a_t)\}_{t=1,2,\dots}$. In most of the literature, the objective is defined as the summation of all intermediate rewards the agent received along the process:

$$u = r_1 + \gamma r_2 + \gamma^2 r_3 + \dots, \quad (1)$$

where $\gamma \in (0, 1)$ is the *discount factor* to encourage the agent focuses more on rewards closer in time. The study of control (when both the reward distribution $p(r_t|s_t, a_t)$ and the state transition distribution $p(s_{t+1}|s_t, a_t)$ are known)

or reinforcement learning (when neither $p(r_t|s_t, a_t)$ nor $p(s_{t+1}|s_t, a_t)$ is known) is to find a policy π for the agent to select actions based on states as $a_t \sim \pi(s_t)$, $\forall t$, such that u in Eq. (1) is optimized.

Corresponding to Eq. (1), the *value function* is defined as the future cumulative rewards the agent expect to receive under a specific policy. Let $\mathcal{V} = \{(s, a) | s \in \mathcal{S}, a \in \mathcal{A}\}$ denote the set of all possible state-action pairs. The value function $Q^\pi \in \mathcal{R}^{|\mathcal{V}|}$ is a vector containing the future cumulative rewards expected starting from each (s, a) tuple, with:

$$Q^\pi(s_t, a_t) = \mathbb{E}_{\substack{\{p(r_{t'}|s_{t'}, a_{t'})\}_{t'=t, t+1, \dots} \\ \{p(s_{t'+1}|s_{t'}, a_{t'})\}_{t'=t, t+1, \dots} \\ \{a_{t'+1} \sim \pi(s_{t'+1})\}_{t'=t, t+1, \dots}}} [r_t + \gamma r_{t+1} + \gamma^2 r_{t+2} + \dots | s_t, a_t] \quad (2)$$

$$= \mathbb{E}_{\substack{p(r_t|s_t, a_t) \\ p(s_{t+1}|s_t, a_t) \\ a_{t+1} \sim \pi(s_{t+1})}} [r_t + \gamma Q^\pi(s_{t+1}, a_{t+1}) | s_t, a_t]. \quad (3)$$

As shown in [32], for stationary single-agent fully-observable MDPs, there exists a deterministic global-optimal policy π^* , with its value function denoted as Q^* , with the following relationship:

$$\pi^*(s_t) = \operatorname{argmax}_a Q(s_t, a) \quad (4)$$

Essentially, $\pi^*(s_t)$ is a *deterministic distribution* with all its probability density on the single action that maximizes $Q(s_t, a_t)$. Therefore, π^* is commonly referred to as a *greedy policy*.

For optimal control, Q^* can be computed by Eq. (3) with π being the global optimal greedy policy π^* , leading to the *Bellman optimality equation*:

$$Q^*(s_t, a_t) = \mathbb{E}_{\substack{p(r_t|s_t, a_t) \\ p(s_{t+1}|s_t, a_t)}} [r_t + \gamma \max_{a_{t+1}} Q^*(s_{t+1}, a_{t+1}) | s_t, a_t]. \quad (5)$$

Meanwhile, for reinforcement learning, Q^* is learned through iterative updates of sample-based approximations to Eq. (5), known as the *Bellman update*:

$$Q(s_t, a_t) \leftarrow r_t + \gamma \max_{a_{t+1}} Q(s_{t+1}, a_{t+1}), \quad (6)$$

where the superscript on Q is dropped, since during these updates, Q does not necessarily correspond to the value function of any policy. We note that in Bellman updates, as shown by Eq. (6), the updated estimations for Q are obtained through bootstrapping from the current estimations. This learning technique is commonly known as *temporal difference learning* [31], which enjoys low estimation variance and high learning efficiency. Eq. (6) is used directly in the value-based algorithms such as SARSA [33], Q-learning [34], with the process commonly referred to as *value iteration*; and policy-based algorithms such as the class of Actor-Critic methods [35].

B. GENERALIZED NON-CUMULATIVE OBJECTIVES

While it is proper to express the objective as Eq. (1) in many scenarios, there exist applications where the objective u is intrinsically some other function over the intermediate rewards. In this paper, to generalize the class of objectives that can be optimized, we formulate the objectives as *general functions* over intermediate rewards:

$$u = f(r_1, r_2, r_3, \dots). \quad (7)$$

Examples for such objectives can be seen from a wide variety of problems, which include, but are not limited to, the following classes of problems:

- The *bottleneck* of the intermediate rewards along the process, which fits into the large class of *max-min optimization problems* [6], [7], [8], [9]. Among these max-min optimizations, the *network routing* problems are perhaps the most standout examples.
- The *largest* reward among the intermediate rewards along the process [28], [29].
- The *harmonic mean* of the intermediate rewards along the process, such as the average traveling velocity, electrical resistance in circuits, density of mixture. It has also been used in wireless communications as a measure of fairness among users [10].

Among various non-cumulative objectives, the objective of the *bottleneck reward* is particularly prevalent. An important class of problems with bottleneck objectives are the network routing problems. Consider a data flow in a communication network consisting of multiple links, the highest rate the flow supports is the rate of the bottleneck link (i.e. the link with the lowest rate). Correspondingly, network routing problems are best formulated by the bottleneck objective. We describe such problems in detail in Section IV.

III. LEARNING ALGORITHMS WITH GENERALIZED BELLMAN UPDATES

This section aims to generalize optimal control and reinforcement learning to MDPs with non-cumulative objectives as Eq. (7) by modifying the operation within the Bellman updates in Eq. (6). We present sufficient conditions on the modified operation as well as assumptions on the underlying MDPs such that the Bellman updates still maintain the global optimal convergence property. Furthermore, we provide examples of frequently-encountered non-cumulative objectives with corresponding operations that satisfy the conditions for convergence.

A. BELLMAN UPDATE WITH GENERALIZED OPERATIONS

Observing Eq. (6), the update target of the new iteration consists of three fundamental elements:

- Intermediate reward r_t ,
- Value function at next state-action pair $Q(s_{t+1}, a_{t+1})$,
- Summation operation to combine a and b.

In this paper, we explore substitutions of c in the Bellman optimality equation and its update rule by an alternative

computational operation, which we refer to as the *generalized Bellman update operation*, denote by $g(\cdot, \cdot)$. The operation takes \mathbf{a} and \mathbf{b} as the two arguments. As the result, we generalize Eq. (6) to the following form:

$$Q(s_t, a_t) \leftarrow g\left(r_t, \gamma \max_{a_{t+1}} Q(s_{t+1}, a_{t+1})\right). \quad (8)$$

Through this generalized Bellman update operation, we are able to adapt the highly efficient temporal difference learning technique, as well as many popular reinforcement learning algorithms that based on it (e.g. SARSA, Q-learning, Actor-Critic), to optimizing the non-cumulative objectives, with minimal changes to these algorithms.

To determine which generalized objective functions $f(\dots)$ as per Eq. (7) can be optimized with such generalized Bellman updates, the first criterion is that the objective function needs to have *optimal substructure* [36], as a fundamental requirement of dynamic programming. Furthermore, for learning based algorithms with value function approximators (such as neural networks), it is desirable to have the value function Q with fixed-dimension outputs from state to state (under most scenarios, the value function is a scalar function). This corresponds to the requirement that the objective function should be computable by iteratively updating a fixed number of statistics over its arguments (i.e. the intermediate rewards). When the two requirements are satisfied, we can deduce the proper operation $g(\cdot, \cdot)$ from the objective function $f(\dots)$ on a case-by-case basis.

B. CONDITIONS FOR CONVERGENCE

To facilitate the theoretical analysis, we denote each step of the value iteration by the function mapping $F^\pi : \mathcal{R}^{|\mathcal{V}|} \rightarrow \mathcal{R}^{|\mathcal{V}|}$. The superscript π indicates that the policy π is used for action selection in the one-step look ahead target computation. Correspondingly, we have:

$$(F^\pi Q)(s_t, a_t) = \mathbb{E}_{\substack{p(r_t|s_t, a_t) \\ p(s_{t+1}|s_t, a_t) \\ a_{t+1} \sim \pi(s_{t+1})}} g\left(r_t, \gamma Q(s_{t+1}, a_{t+1})\right), \quad (9)$$

where the expectation is understood as conditioned under (s_t, a_t) . When the deterministic greedy policy as derived from the current Q is used, the value iteration is denoted by F^* as follows:

$$(F^*Q)(s_t, a_t) = \mathbb{E}_{\substack{p(r_t|s_t, a_t) \\ p(s_{t+1}|s_t, a_t)}} g\left(r_t, \gamma \max_{a_{t+1}} Q(s_{t+1}, a_{t+1})\right). \quad (10)$$

The original Bellman updates in Eq. (6) enjoy convergence to the global optimal value function, as shown in [37] and [38]. To generalize this convergence property for the generalized updates as Eq. (10), we present a sufficient condition on $g(\cdot, \cdot)$ for ensuring convergence to a unique value function in the following theorem.

*Theorem 1: On a single-agent fully observable MDP, the series of value functions obtained from iteratively applying the generalized Bellman update rule $Q \leftarrow F^*Q$ as in Eq. (10)*

is guaranteed to converge to a unique convergence point in $\mathcal{R}^{|\mathcal{V}|}$ from any arbitrary starting point, if $g(\cdot, \cdot) : \mathcal{R} \times \mathcal{R} \rightarrow \mathcal{R}$ satisfies the following condition:

$$|g(a, b) - g(a, c)| \leq |b - c| \quad \forall a, b, c \in \mathcal{R} \quad (11)$$

The mathematical proof of this theorem is presented in Appendix A.

Note that in Theorem 1, we do not claim that the greedy policy resulted from the converged value function is the global optimal policy. Besides an additional condition we need on the operation $g(\cdot, \cdot)$ (to be introduced in the next subsection), the main reason is that after generalizing the Bellman update operation to $g(\cdot, \cdot)$, the value function learned through the value iteration process is no longer guaranteed to be the true expectation of the objective value as defined in Eq. (7), when the state transition functions and reward functions are stochastic. We elaborate on this observation in the following subsection.

C. SUBOPTIMALITY WITH STOCHASTIC TRANSITIONS AND REWARDS

With the generalized operation $g(\cdot, \cdot)$, we express the objective in (7) with $g(\cdot, \cdot)$ as:

$$u = f(r_1, r_2, r_3, \dots) = g\left(r_1, \gamma g(r_2, \gamma g(r_3, \dots))\right). \quad (12)$$

We have shown a condition on $g(\cdot, \cdot)$ for convergence in Theorem 1 for obtaining Q^* . Nonetheless, we observe that Q^* does not necessarily recover the true expectation of u when stochastic state transitions and rewards are considered. To illustrate this, consider an episode starting from the state s_1 . Under the greedy policy π^* derived from Q^* , we take the expectation over $p(r_t|s_t, a_t)$ and $p(s_{t+1}|s_t, a_t)$ on Eq. (12), which leads to:

$$\begin{aligned} & \mathbb{E}_{\substack{\{a_t = \pi^*(s_t)\}_{t=1,2,\dots} \\ \{p(r_t|s_t, a_t)\}_{t=1,2,\dots} \\ \{p(s_{t+1}|s_t, a_t)\}_{t=1,2,\dots}}} [u(r_1, r_2, r_3, \dots)] \\ &= \mathbb{E}_{\substack{\{a_t = \pi^*(s_t)\}_{t=1,2,\dots} \\ \{p(r_t|s_t, a_t)\}_{t=1,2,\dots} \\ \{p(s_{t+1}|s_t, a_t)\}_{t=1,2,\dots}}} \left[g\left(r_1, \gamma g(r_2, \gamma g(r_3, \dots))\right) \right]. \quad (13) \end{aligned}$$

Meanwhile, starting from $t = 1$, with the converged Q^* obtained from the generalized Bellman updates, we have:

$$\begin{aligned} & \mathbb{E}_{a_1 = \pi^*(s_1)} [Q^*(s_1, a_1)] \\ &= \mathbb{E}_{\substack{a_1 = \pi^*(s_1) \\ p(r_1|s_1, a_1) \\ p(s_2|s_1, a_1) \\ a_2 = \pi^*(s_2)}} \left[g\left(r_1, \gamma Q^*(s_2, a_2)\right) \right] \\ &= \mathbb{E}_{\substack{a_1 = \pi^*(s_1) \\ p(r_1|s_1, a_1) \\ p(s_2|s_1, a_1) \\ a_2 = \pi^*(s_2)}} \left[g\left(r_1, \gamma \mathbb{E}_{\substack{p(r_2|s_2, a_2) \\ p(s_3|s_2, a_2) \\ a_3 = \pi^*(s_3)}} [g(r_2, \gamma Q^*(s_3, a_3))]\right) \right]. \quad (15) \end{aligned}$$

Comparing Eq. (13) and Eq. (15), for $\mathbb{E}_{a_1 = \pi^*(s_1)} [Q^*(s_1, a_1)]$ to be equal to the expectation of u under π^* , $p(r_t|s_t, a_t)$, and $p(s_{t+1}|s_t, a_t)$, we require $g(\cdot, \cdot)$ to be exchangeable

with $\mathbb{E}_{\pi^*}[\cdot]$, $\mathbb{E}_{p(r_t|s_t, a_t)}[\cdot]$, and $\mathbb{E}_{p(s_{t+1}|s_t, a_t)}[\cdot]$. With π^* being the deterministic greedy policy as in Eq. (4), the operation $\mathbb{E}_{\pi^*}[\cdot]$ can always be exchanged with $g(\cdot, \cdot)$. However, if $p(r_t|s_t, a_t)$ or $p(s_{t+1}|s_t, a_t)$ is stochastic, $\mathbb{E}_{p(r_t|s_t, a_t)}[\cdot]$ or $\mathbb{E}_{p(s_{t+1}|s_t, a_t)}[\cdot]$ is not necessarily exchangeable with $g(\cdot, \cdot)$. In this case, Eq. (13) and Eq. (15) can potentially evaluate to different values, and therefore π^* derived from Q^* may be suboptimal.

Under this observation, in order to obtain a global optimality guarantee on the greedy policy π^* , we constrain the scope to deterministic MDPs. Furthermore, we introduce an additional condition on the generalized operation $g(\cdot, \cdot)$ in order to establish global optimality, as formally stated in the following theorem:

Theorem 2: Given a non-cumulative objective function u and its corresponding generalized Bellman update operation $g(\cdot, \cdot) : \mathcal{R} \times \mathcal{R} \rightarrow \mathcal{R}$ satisfying the condition Eq. (11) from Theorem 1, let Q^ denote the convergence point of the value iteration (from iteratively applying the generalized Bellman update rule as in Eq. (10)). For an MDP with deterministic $p(r_t|s_t, a_t)$ and $p(s_{t+1}|s_t, a_t)$, the greedy policy π^* derived from Q^* is guaranteed to be the global optimal policy, if $g(\cdot, \cdot)$ satisfies the following additional condition:*

$$b \geq c \text{ implies } g(a, b) \geq g(a, c) \quad \forall a, b, c \in \mathcal{R} \quad (16)$$

The mathematical proof of this theorem is provided in Appendix B.

We note that the assumptions on MDPs in Theorem 2 are satisfied by a large class of optimal control and reinforcement learning problems: e.g., board games including Go and Chess, a subset of Atari games, the class of network routing problems (such as the problems to be studied in Section IV), and so on.

To summarize, given any general MDP, we may generalize its objective function and apply the generalized Bellman update as in Eq. (10) to try to learn its value function. If the generalized update operation satisfies the condition as in Eq. (11) in Theorem 1, the value iteration is guaranteed to converge to a unique convergence point. Furthermore, if the underlying MPD satisfies the assumptions in Theorem 2, and the update operation satisfies the condition Eq. (16), the convergence point is the optimal value function and the greedy policy π^* derived from the value function is guaranteed to be the global optimal policy.

D. EXAMPLES OF GENERALIZED OBJECTIVES AND BELLMAN UPDATE OPERATIONS

We introduce several widely applicable non-cumulative objectives, and present the corresponding modified Bellman update operations. In Appendix B, we provide the proofs that these operations satisfy the conditions in Theorem 1 and Theorem 2.

1) BOTTLENECK REWARD OBJECTIVE

The objective u is the minimum (i.e. *bottleneck*) intermediate reward in the process:

$$u(r_1, r_2, r_3, \dots) = \min(r_1, r_2, r_3, \dots). \quad (17)$$

The corresponding modified Bellman update operation is:

$$\begin{aligned} g(r_t, \gamma \max_{a_{t+1}} Q(s_{t+1}, a_{t+1})) \\ = \min(r_t, \gamma \max_{a_{t+1}} Q(s_{t+1}, a_{t+1})), \end{aligned} \quad (18)$$

where the discount factor γ is useful for encouraging the agent to postpone the occurrences of negative rewards that often correspond to undesired or failure outcomes.

The proof that the bottleneck update operation satisfies both conditions in Theorem 1 and Theorem 2 is presented in Appendix C.

2) MAXIMUM REWARD OBJECTIVE

The objective u is the maximum intermediate rewards within the process:

$$u(r_1, r_2, r_3, \dots) = \max(r_1, r_2, r_3, \dots). \quad (19)$$

The corresponding modified Bellman update operation is:

$$\begin{aligned} g(r_t, \gamma \max_{a_{t+1}} Q(s_{t+1}, a_{t+1})) \\ = \max(r_t, \gamma \max_{a_{t+1}} Q(s_{t+1}, a_{t+1})). \end{aligned} \quad (20)$$

The proof that the maximum update operation Eq. (20) satisfies both conditions in Theorem 1 and Theorem 2 follows the same logic as the proof for the bottleneck update operation shown in Appendix C.

3) HARMONIC MEAN REWARD OBJECTIVE

Assuming $r_t > 0, \forall t$, and the process is always terminated after a fixed number of steps, the objective u is the harmonic mean of all intermediate rewards within the process:

$$U(r_1, r_2, \dots, r_T) = \frac{1}{\frac{1}{r_1} + \frac{1}{r_2} + \frac{1}{r_3} + \dots + \frac{1}{r_T}}, \quad (21)$$

where we omit the constant reward count. Examples of such applications with harmonic mean objectives include:

- Optimize average traveling speed over a trip consisting of a fixed number of intervals.
- Minimize resistance in a circuit with a fixed number of resistors in parallel connection.
- Optimize mixture density (e.g. alloys) with a fixed number of selections on equal-weight components.

Although technically maximizing Eq. (21) is equally valid as minimizing the summation of inverse of rewards, we present it as an example of a non-cumulative objective function with a modified Bellman update operation (as shown below) that satisfies the proposed convergence conditions.

The corresponding modified Bellman update operation is:

$$g(r_t, \gamma \max_{a_{t+1}} Q(s_{t+1}, a_{t+1})) = \frac{1}{\frac{1}{r_t} + \frac{1}{\gamma \max_{a_{t+1}} Q(s_{t+1}, a_{t+1})}}. \quad (22)$$

The proof that the harmonic mean update operation satisfies both conditions in Theorem 1 and in Theorem 2 is presented in Appendix D.

IV. APPLICATIONS OF GENERALIZED REINFORCEMENT LEARNING

A. CLASSICAL REINFORCEMENT LEARNING PROBLEMS WITH BOTTLENECK OBJECTIVES

We first re-examine classical reinforcement learning problems, formulated with the bottleneck objectives as introduced in Section III-D.1. In many classical optimal control and reinforcement learning applications, the agent's success is largely based on its ability to *avoid failure or defeat*. This is particularly the case when the MDPs lack significant intermediate milestones or checkpoints, such as the *CartPole* problem and the Atari game *Breakout*. Instead of regarding such tasks as collecting as many rewards as possible, the agent can interpret the tasks with the equally valid strategy of *avoiding the worst outcome* (corresponding to the lowest reward) as much as possible.

Conventionally, both tasks are formulated with the cumulative objective, each with an incremental rewarding scheme. In the *CartPole* task, a positive reward is assigned to the agent for every timestep it maintains the pole in the upright position; while in Atari, a positive reward is assigned each time the agent breaks a brick with the bouncing ball.

To formulate the task with the bottleneck objective for such classical tasks, we assign a negative reward to the agent when an undesired or failure event occurs after executing a certain action. For the other actions that do not directly lead to the failure events, we simply assign a zero intermediate reward. In the *CartPole* task, the agent aims to control the cart to vertically balance the pole. When the pole falls outside a pre-defined angle range, a negative reward is assigned to the agent. Similarly, for the Atari game *Breakout*, the agent controls the movement of a paddle to catch and reflect a bouncing ball upwards to destroy layers of bricks located above. Each time the agent fails to catch the falling ball with the paddle, it is assigned a negative reward. With the discount factor γ applied on rewards over time steps, the later the negative rewards occur, the higher the bottleneck objective is.

By optimizing the bottleneck objective, the agent is able to learn *alternative strategies* to these classical problems: For *CartPole*, the strategy is to prevent the pole from falling for as long as possible. For *Breakout*, the strategy is to keep the ball in play for a maximized duration through controlling the paddle to constantly catch and reflect the ball, which translates to, although not always most efficiently, maximizing the bricks destroyed and thus achieving competitive game scores.

B. SINGLE-PATH MAXIMUM-FLOW ROUTING WITH BOTTLENECK OBJECTIVE ON A GRAPH

1) PROBLEM SETUP

Consider a communication network modeled as a directed graph $G = (\mathcal{N}, \mathcal{E})$, where the set of nodes \mathcal{N} corresponds to the transmission nodes, and the set of edges \mathcal{E} corresponds to the communication links between the nodes.

A single-path data flow is routed through the network, from a fixed source node $n_s \in \mathcal{N}$ towards a fixed destination node $n_t \in \mathcal{N}$. Each directed edge $e^{n_i \rightarrow n_j} \in \mathcal{E}$ from $n_i \in \mathcal{N}$ to $n_j \in \mathcal{N}$ represents the transmission link from n_i to n_j , and is assigned with a link rate capacity $r(e^{n_i \rightarrow n_j}) = r_{n_i \rightarrow n_j}$. We set $r_{n_i \rightarrow n_j} = 0$ when there is no link from n_i to n_j in the network. The optimal routing problem is that of finding an ordered sequence of relay nodes as transmission hops, to form the route such that the bottleneck rate is maximized:

$$\text{maximize}_{n_1, n_2, \dots, n_m} \min(r_{n_s \rightarrow n_1}, r_{n_1 \rightarrow n_2}, \dots, r_{n_{m-1} \rightarrow n_m}, r_{n_m \rightarrow n_t}), \quad (23)$$

where $\{n_i\}_{i \in \{1 \dots m\}}$ denote the m relay nodes (with the number m adjustable) forming the route of the flow.

2) GENERALIZED OPTIMAL CONTROL SOLUTION

To find the single-path maximum flow within a given network represented by a directed graph as described above, we formulate the routing process as an MDP: the agent moves along the frontier node of the route, and makes sequential decisions on the selection of the node for each hop, until the destination node is reached. For the state space \mathcal{S} , each state is uniquely identified by the frontier node the agent resides on. Specifically, we use $\mathbf{s}^{n_i} \in \mathcal{S}$ to denote the state that the current frontier node of the partially established route is node n_i . For the action space \mathcal{A} , we use $\mathbf{a}^{n_i \rightarrow n_j} \in \mathcal{A}$ to denote the action to move from node n_i to node n_j . Lastly, as specified in the problem setup, $r_{n_i \rightarrow n_j}$ corresponds to the reward for the action $\mathbf{a}^{n_i \rightarrow n_j}$, which is the link rate capacity of the link from n_i to n_j .

To optimize the objective (23), for each state and action pair $(\mathbf{s}^{n_i}, \mathbf{a}^{n_i \rightarrow n_j})$, the generalized update is as follows:

$$Q(\mathbf{s}^{n_i}, \mathbf{a}^{n_i \rightarrow n_j}) \leftarrow \min(r_{n_i \rightarrow n_j}, \gamma \max_{n_k} Q(\mathbf{s}^{n_k}, \mathbf{a}^{n_j \rightarrow n_k})). \quad (24)$$

From the converged Q^* , we obtain the global optimal greedy policy π^* (guaranteed by the results in Section III-B and III-C), following which produces the flow route supporting the global maximal flow rate.

C. WIRELESS AD HOC NETWORK ROUTING AND SPECTRUM ACCESS WITH BOTTLENECK OBJECTIVE

1) PROBLEM SETUP

Consider the physical-layer routing problem as discussed in [6]. In a wireless ad hoc network with a set of transmission nodes \mathcal{N} , a set of data flows \mathcal{K} is to be established, each consisted of multiple hops with their own pairs of source and

destination nodes. A set of frequency bands \mathcal{B} is available for transmission, each with a bandwidth of W Hz. We focus on two optimization tasks for these data flows: routing and spectrum access. The task of routing is to select an ordered list of intermediary relay nodes from \mathcal{N} to form the route for each data flow. The task of spectrum access is to select a frequency band from \mathcal{B} for the transmission of each hop in the route of each flow. We represent the route for flow $k \in \mathcal{K}$ as an ordered list denoted by $\mathbf{n}^{(k)}$:

$$\mathbf{n}^{(k)} = (n_0^{(k)}, n_1^{(k)}, n_2^{(k)} \dots n_m^{(k)}, n_{m+1}^{(k)}), \quad (25)$$

where $n_0^{(k)}$ and $n_{m+1}^{(k)}$ represent the fixed source and destination node for flow k , and $\{n_i^{(k)}\}_{i \in \{1 \dots m\}}$ represent the m relay nodes (with the number m adjustable) forming the route of flow k . We represent the spectrum access solution for flow k as an ordered list denoted by $\mathbf{b}^{(k)}$, containing the selected frequency band of each hop:

$$\mathbf{b}^{(k)} = (b_1^{(k)}, b_2^{(k)}, b_3^{(k)} \dots b_{m+1}^{(k)}), \quad (26)$$

where $b_i^{(k)} \in \mathcal{B}$ denotes the frequency band selected for the i -th hop in the route of flow k , with $i \in \{1 \dots m + 1\}$. As the global topology of the ad hoc network is not available as inputs, the agents need to learn to infer the network topology during the routing process.

Consider a link from node n_i to node n_j over frequency band b . Let $h_{(n_i \rightarrow n_j, b)} \in \mathcal{C}$ denote its channel coefficient. The maximum transmission rate of this link is based on the *signal to interference plus noise ratio* (SINR) as follows:

$$\text{SINR}_{(n_i \rightarrow n_j, b)} = \frac{x_{n_i, b} |h_{(n_i \rightarrow n_j, b)}|^2 p}{\sum_{\substack{n_l \neq n_i, n_j \\ n_l \in \mathcal{N}}} x_{n_l, b} |h_{(n_l \rightarrow n_j, b)}|^2 p + \sigma^2}, \quad (27a)$$

$$r_{(n_i \rightarrow n_j, b)} = W \log(1 + \text{SINR}_{(n_i \rightarrow n_j, b)}). \quad (27b)$$

where p and σ^2 denote the transmit power of each node and the background noise power on each frequency band. The binary control variable $x_{n_i, b}$ indicates whether the node n_i is transmitting on the band b or idle. The objective for each flow $u^{(k)}$ is the transmission rate it supports, which is the bottleneck link rate:

$$u^{(k)} = r_{\min}^{(k)} = \min_{i=0, 1, 2, \dots, m} r(n_i^{(k)} \rightarrow n_{i+1}^{(k)}, b_{i+1}^{(k)}). \quad (28)$$

The global objective u over all data flows is then defined as the average of the bottleneck rates over all data flows:

$$u = \frac{\sum_{k \in \mathcal{K}} u^{(k)}}{|\mathcal{K}|} \quad (29)$$

2) GENERALIZED REINFORCEMENT LEARNING SOLUTION

For the physical layer routing and spectrum access problem, we assign one agent per data flow, with each agent moving along the frontier node of its flow and making hop-by-hop decisions. With multiple data flows to be jointly optimized, this problem is essentially a multi-agent reinforcement learning problem, with higher complexity than the maximum flow

routing problem on a graph as in Section IV-B. By optimizing this problem with the bottleneck objective formulation and the generalized Bellman updates, we demonstrate that the proposed approach is competitive and highly effective in the setting of multi-agent reinforcement learning.

For better parameter efficiency, we only train one set of parameters shared among all agents. We assume the wireless network is only partially observable to each agent, meaning Q^* is no longer guaranteed to be global optimal. Nonetheless, as shown in later simulations, the corresponding π^* is still competitive. We adopt the MDP formulation as in [6]. At each step, each agent gathers 4 pieces of information on frequency band b for each of its c closest neighboring nodes: the agent-to-neighbor distance; the neighbor-to-destination distance; the angle between agent-to-neighbor and agent-to-destination directions; and the signal interference on the neighbor on band b . With this information, the agent forms the state s on band b with $s \in \mathbb{R}^{4c}$, $\forall s \in \mathcal{S}$. For the action space \mathcal{A} , the agent has $c + 1$ actions on band b : one action for connecting with each of the c nodes via b , and one action for *reprobing* on b (if none of the c nodes is suitable). We use $\mathbf{s}^{(n_i, b)}$ to denote the state that the frontier node of the partially established flow is node n_i and that the transmission to the next hop uses the band b . We use $\mathbf{a}^{(n_i \rightarrow n_j, b)}$ to denote the agent's action to establish the link from node n_i to node n_j using band b , which is assigned the reward as the rate of this link $r_{(n_i \rightarrow n_j, b)}$. During training, these rewards as link rates are computed after the routes are formed.

As the bottleneck rate is not expressible as summations, [6] uses the Monte-Carlo method [39] for estimating the value function. The key improvement we propose over [6] is to utilize the modified Bellman update rule for training the agents in the off-policy fashion, providing higher data efficiency, faster convergence, and better performances. Using Eq. (18), the generalized updates for training each agent are:

$$Q(\mathbf{s}^{(n_i, b)}, \mathbf{a}^{(n_i \rightarrow n_j, b)}) \leftarrow \min(r_{(n_i \rightarrow n_j, b)}, \gamma \max_{n_k, b'} Q(\mathbf{s}^{(n_j, b')}, \mathbf{a}^{(n_j \rightarrow n_k, b')})). \quad (30)$$

After predicting Q values for all frequency bands, the agent selects the action with the single highest Q value among all bands to establish the new link, which specifies not only the optimal node as the next hop, but also the optimal frequency band for transmission to that node.

V. SIMULATIONS

We experiment on the optimal control and reinforcement learning problems and compare solutions from the conventional Bellman updates and from our proposed generalized Bellman updates. We use following terms to refer to each algorithm:

- *Q-Min*: Optimal control solution or RL policy based on the value function obtained from the generalized Bellman update rule as Eq. (8) and Eq. (18).

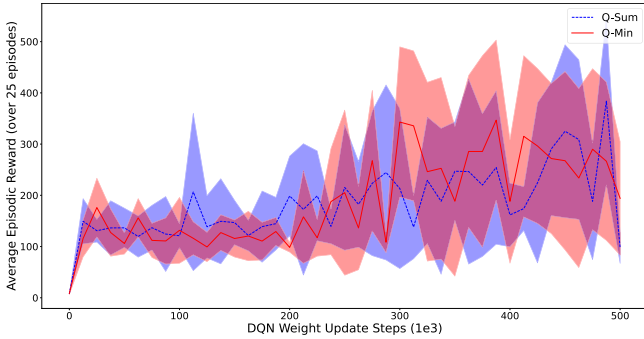


FIGURE 1. Learning progress on CartPole: From 6 full training trails under 6 random seeds, we plot the mean as the line and the standard deviation as the shadow.

- *Q-Sum*: Optimal control solution or RL policy based on the value function obtained from the conventional Bellman update rule as Eq. (6).

A. CLASSICAL REINFORCEMENT LEARNING PROBLEMS

We use the double-DQN architecture [40] to model the agents. During training, the ϵ -greedy policy with decaying ϵ is used for collecting experiences, along with prioritized experience replay [41] for sampling training batches in each update step.

1) CARTPOLE TASK

To solve the CartPole task with the Q-Min algorithm, when the pole falls outside of the pre-defined angle range ($\pm 12^\circ$ from the up-right position), we assign a negative reward of -1 to the agent. To encourage the agent to postpone negative reward occurrence, we use a discount factor $\gamma = 0.95$ in Eq. (18). For learning with the Q-Sum algorithm, we follow the conventional incremental rewarding scheme that has been long used in this task.

We illustrate the agent learning progress under both algorithms in Fig. 1, where we evaluate each agent's performance, averaged over 25 new episodes, after each 12500 update steps of training. We note that we stick with the conventional cumulative objective for CartPole as the performance metric when visualizing the learning progresses of both algorithms as shown in Fig. 1, which allows us to compare both algorithms directly. A competitive performance by the Q-Min algorithm on the cumulative objective would indicate that our alternative bottleneck objective is also a viable for formulating the task.

As shown by the numerical results, besides the oscillations in both learning curves (as DQN is known for unstable learning), the Q-Min agent and the Q-Sum agent learn to balance the pole at a similar pace throughout training. The close results between the two algorithms validate that the bottleneck objective is indeed a suitable alternative to the CartPole objective formulation.

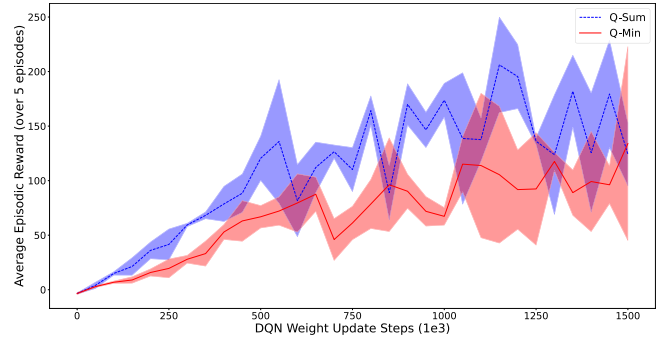


FIGURE 2. Learning progress on Breakout: From 3 full training trails under 3 random seeds, we plot the mean as the line and the standard deviation as the shadow.

2) ATARI BREAKOUT GAME

To solve Atari with the proposed Q-Min algorithm, we utilize a simple reward scheme under Q-Min: we assign a negative reward of -1 to the agent each time it fails to catch the ball with the paddle, and set $\gamma = 0.98$ in Eq. (18) to encourage the agent to postpone such failure events. For learning with the Q-Sum algorithm, we follow the conventional incremental rewarding scheme originally built into the Atari game engine.

We present the learning progress of Q-Min and Q-Sum in Fig. 2, with each agent's performance evaluated and averaged over 5 new game runs, after each 50 thousand update steps of training. Similar to the learning progress visualization for CartPole, we also use the conventional cumulative objective for the original Breakout game as the performance metric when plotting the learning curves of both algorithms.

Unlike in CartPole, the Q-Min agent shows a slightly slower learning progress and lower performance for Breakout. This is likely due to the Q-Min agent not learning the strictly optimal trajectories of redirecting the ball for hitting the most bricks, as its sole objective is to keep the ball in play. Nonetheless, even with simpler and more sparse rewards than the rewards used by Q-Sum, the Q-Min agent still manages to achieve relatively close performances to the Q-Sum agent, especially at the late training stage. The results illustrate the viability of interpreting Breakout with the bottleneck objective formulation.

We emphasize again that, specifically for these two classical problems, our goal is not to show that the proposed Q-Min algorithm is strictly superior to the conventional Q-Sum algorithm. After all, these two problems have long served as the canonical examples for the conventional reinforcement learning problems formulated with the cumulative objectives. Instead, we have shown that it is also valid to interpret and optimize these classical problems with the bottleneck objective formulation. Through learning with the proposed generalized Bellman update rule, the agent is capable of achieving the performances comparable with the results from the convention reinforcement learning approach as presented above. Essentially, When optimizing the agent under the




FIGURE 3. The maximum network flow problem as modeled by a directed graph.

bottleneck objective formulation, the agent learns an alternative game playing strategy for both CartPole and Breakout: to avoid or delay the failure event for as long as possible.

B. SINGLE-PATH MAXIMUM-FLOW ROUTING ON GRAPH

We consider the directed graph network shown in Fig. 3, and perform the Q-Min algorithm with Eq. (24) until convergence. With the MDP in this problem being finite, we set $\gamma = 1$, which simplifies the numerical results with Q values precisely equal to the future bottleneck rates.

In Table 1, we present the iterations of both the Q-Min algorithm and the Q-Sum algorithm. In the first row of the table, we adopt the simplified notations for Q values: we use $Q_{n_i \rightarrow n_j}$ to uniquely denote the state-action value function $Q(s_{n_i}, a_{(n_i, n_j)})$ in Eq. (24). We adopt synchronized iterations, where in each new iteration, the Q value in the right-hand-side of Eq. (24) comes from the previous iteration. All the iterations of value function updates are shown until convergence.

For the Q-Min algorithm, it takes 4 iterations of the generalized Bellman updates to converge. From the resulted Q^* , we deduce the optimal policy π^* producing the following optimal flow route:

$$s \rightarrow b \rightarrow a \rightarrow d \rightarrow t. \quad (31)$$

This route obtained supports a flow rate of 5, which is indeed the global optimal flow rate.

On the other hand, for the Q-Sum algorithm, the convergence speed is lower than Q-Min, as it takes 5 iterations of the regular Bellman updates to converge. Furthermore, the deduced optimal policy results in the following flow route:

$$s \rightarrow b \rightarrow a \rightarrow c \rightarrow d \rightarrow t. \quad (32)$$

This route supports a flow rate of 4, which is sub-optimal and inferior to the route obtained by the Q-Min algorithm.

C. WIRELESS AD HOC NETWORK ROUTING AND SPECTRUM ACCESS

1) EXPERIMENT SETTINGS

We simulate wireless ad hoc networks in a 1000m \times 1000m region with $|\mathcal{K}| = 3$ data flows and $|\mathcal{B}| = 8$ frequency bands. We adopt the same specifications as in [6] as we aim to compare results and illustrate the effectiveness of the proposed

Q-Min algorithm. Specifically, we consider the short-range outdoor model ITU-1411 with a distance-dependent path-loss to model all wireless channels, over all frequency bands at 2.4GHz carrier frequency. Shadowing and fast-fading are not considered in the simulation setting. This corresponds to an outdoor environment (e.g., a rural or remote area), where the strengths of the wireless links are mostly functions of the distances between the transmitters and the receivers. We assume each of the $|\mathcal{B}| = 8$ frequency bands has a 5MHz bandwidth for signal transmission. All antennas equipped at the transmission nodes have a height of 1.5m and 2.5dBi antenna gain. We assume a transmit power of 30dBm for all nodes and background noise at -130 dBm/Hz.

To generate realistic wireless network layouts, the node locations are randomly generated with varying node densities over the region. Specifically, we divide the 1000m \times 1000m network region into nine equal sub-regions, and randomly locate (6, 8, 7, 6, 5, 10, 8, 9, 6) nodes within each of the nine sub-regions correspondingly.

2) TRAINING CONVERGENCE SPEED COMPARISON

We train each set of $|\mathcal{K}| = 3$ agents with three algorithms: Q-Min, Q-Sum, and the algorithm by [6]:

- *Q-MC*: RL policy based on the value function obtained from the Monte-Carlo episodic estimations of future bottleneck rewards, computed at the end of episodes.

We generate 380,000 wireless ad hoc network layouts for training the agents under each algorithm, under the following training schedule:

- Initial 30,000 layouts are used for random routing on collecting initial experience.
- The middle 300,000 layouts are used for the ϵ -greedy policy based routing, with the ϵ value follows linear annealing from 1.0 to 0.1 throughout the training over these layouts.
- The final 50,000 layouts are used with $\epsilon = 0$ for the final convergence stage.

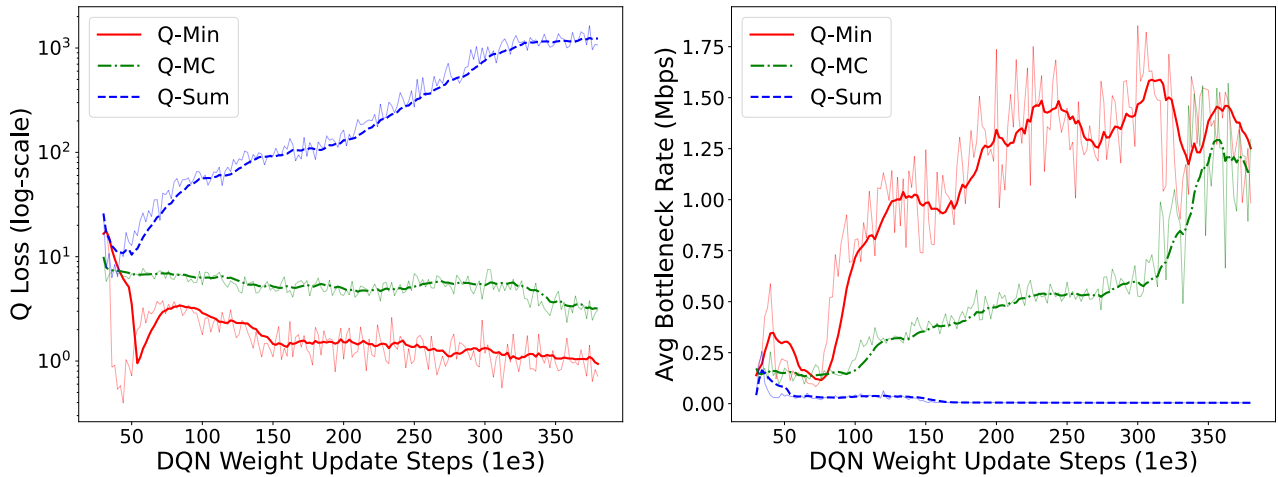
We use the Dueling-DQN architecture [42] to model all the agents, with the neural network specifications listed in Table 2, same as in [6]. Since the rewards as link rates are dense throughout the MDP, uniform sampling is sufficient for experience replay. We use $c = 10$ as the number of neighbors the agent explores each time. The state inputs to the DQNs are therefore 40-component vectors, i.e. $s \in \mathcal{R}^{40}, \forall s \in \mathcal{S}$.

For the same reasons as in Section V-B, we set $\gamma = 1$ in Eq. (30). During training, we track both the mean-squared-errors for predictions on Q , and the routing performances over 100 newly generated network layouts at each update step of training. The training curves for all three algorithms are displayed in Fig. 4.

Shown by the learning curves, the conventional Q-Sum agents collectively achieve the worst learning progress and simply fail to convergence on the Q value estimations. While both the Q-Min agents and the Q-MC agents converge to comparable performances, the Q-Min agents enjoy a much

TABLE 1. Q-min value iterations on graph network routing (shaded entries depict the routing selection under π^*).

Algorithm	Update Step	$Q_{d \rightarrow t}$	$Q_{c \rightarrow t}$	$Q_{c \rightarrow d}$	$Q_{a \rightarrow c}$	$Q_{a \rightarrow d}$	$Q_{b \rightarrow d}$	$Q_{b \rightarrow c}$	$Q_{b \rightarrow a}$	$Q_{s \rightarrow a}$	$Q_{s \rightarrow b}$
Q-Min	0	0	0	0	0	0	0	0	0	0	0
	1	5	3	0	0	0	0	0	0	0	0
	2	5	3	4	3	5	3	3	0	0	0
	3	5	3	4	4	5	3	4	5	4	3
	4 (Converged)	5	3	4	4	5	3	4	5	4	5
Q-Sum	0	0	0	0	0	0	0	0	0	0	0
	1	5	3	4	8	5	3	9	7	4	6
	2	5	3	9	12	10	8	13	15	12	13
	3	5	3	9	17	10	8	18	19	16	21
	4	5	3	9	17	10	8	18	24	21	25
5 (Converged)	5	3	9	17	10	8	18	24	21	30	

**FIGURE 4.** Learning progress of various algorithms on the physical layer routing problem: The dotted lines are the performance values of all update steps; the solid lines are the smoothed curves, with each point being the average over the last 10 update steps.**TABLE 2.** Dueling-DQN neural network specifications.

Model Module	Hidden Layer	Neurons
Feature Learning Module	1st	150
	2nd	150
State-Value Estimation Stream	1st	100
	2nd	1 (1 state value)
Action-Advantage Estimation Stream	1st	100
	2nd	11 (11 actions)

faster convergence speed. This illustrates the advantage of the temporal difference learning over the Monte-Carlo method, which is made possible for non-cumulative objectives with the proposed generalized update rules.

To better understand why Q-Min achieves noticeably faster convergence than Q-MC, we emphasize that the Monte-Carlo estimations used by Q-MC are highly affected by the random explorations especially at the early stage of training. Certain random explorations might lead to an extremely low bottleneck rate for the newly established route. This bottleneck rate is then used as the Monte-Carlo estimation on the value function during training the Q-MC agent. Thus, the value

estimations learned by the Q-MC agents suffer from low qualities significantly at the beginning of training.

On the other hand, with the proposed generalized update operation, the bottleneck objective can be estimated by the temporal difference learning technique as in the Q-Min algorithm. As an *off-policy*¹ learning algorithm, the temporal difference learning estimations are much more resilient to the random explorations, since the estimation target is obtained through one-step bootstrapping on the already learned value function. Therefore, it is not a surprise to see the significant improvements on the training efficiency and convergence speed by the Q-Min agents.

3) PERFORMANCES ON BOTTLENECK RATES

We present in Table 3 test results on the number of links established on each flow, as well as the achieved bottleneck flow rates for these data flows, over 1000 newly generated testing wireless ad hoc networks. Furthermore, for each method, we collect the bottleneck rates of all data flows over all testing wireless networks, and present the cumulative distribution

¹An off-policy algorithm separates the policy that the value estimation is based on from the sampling policy, which is desired when the sampling policy is highly noisy (e.g. with many random explorations).

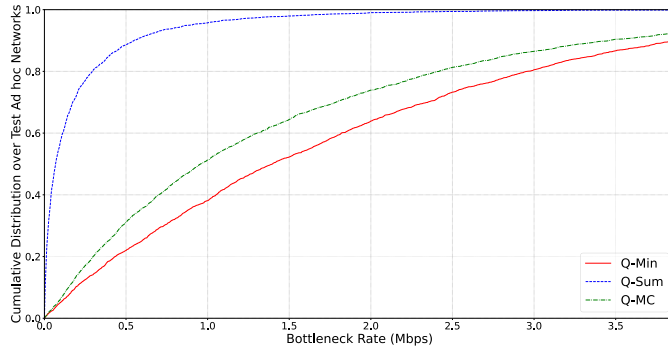


FIGURE 5. Cumulative distribution function of flow bottleneck rates in 1000 testing wireless ad hoc networks.

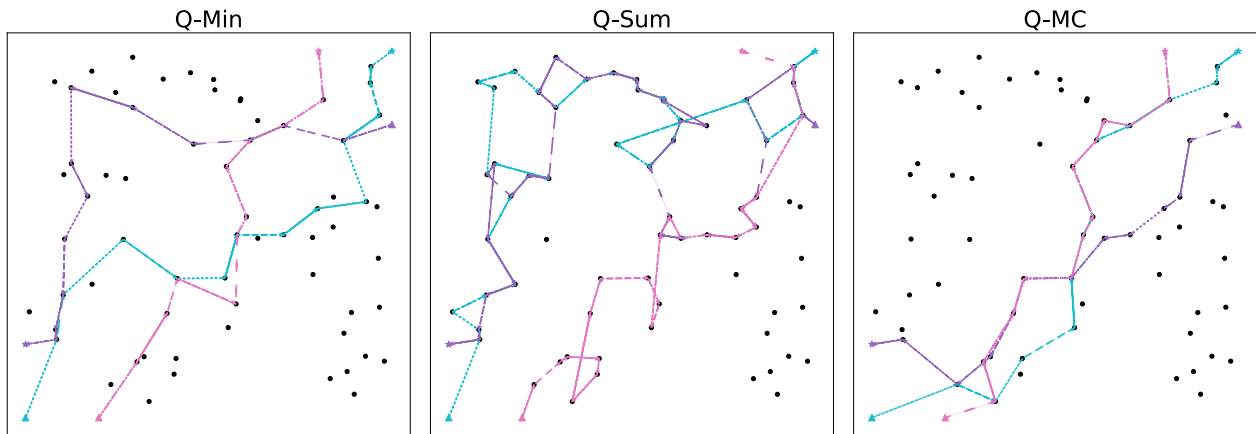


FIGURE 6. Routes for data flows (data flows differentiated by colors; frequency bands differentiated by linestyles).

TABLE 3. Average bottleneck flow rate performances.

Methods	Average # Links per Flow	Average Bottleneck Rate
Q-Min	13.0	1.76 Mbps
Q-MC	13.8	1.42 Mbps
Q-Sum	31.2	0.21 Mbps

function (CDF) of these bottleneck rates in Fig. 5. As shown by both the statistics and the distributions of the flow rates, the Q-Min agents achieve the best routing results, whereas the Q-Sum agents perform the worst by a large margin, while having much higher numbers of links over the established data flows.

We visualize the optimized routes by each RL algorithm over a random wireless ad hoc network in Fig. 6. The Q-Min agents learn to establish links with medium lengths. This policy ensures a certain level of channel strength for the bottleneck links, without constructing too many links to avoid excessive interference which is detrimental to the bottleneck link rates. Furthermore, the Q-Min agents also learn to spatially spread out data flows as well as the frequency bands used among links for effective interference mitigation.

On the other hand, the Q-Sum agents learn the policy that connects unnecessarily many short links to form routes, neglecting the importance of the bottleneck link within each flow. Evidently, the conventional reinforcement learning formulation is unsuitable for solving such routing problems. For this reason, in application fields such as network communications, generalizing the objective function and its learning rule through our proposed approach is an essential optimization technique.

VI. CONCLUSION

This paper recognizes the possibilities of formulating optimal control or reinforcement learning objectives as non-cumulative functions over rewards, and generalizes existing algorithms to optimizing such objectives. Specifically, we explore the generalized operations in the Bellman update rule, for which we provide the global convergence conditions with mathematical proofs. We also recognize the assumptions required on the MDP state transitions and reward functions for ensuring the global optimality on the obtained policies. With the generalized objectives and learning algorithms, we are able to unveil alternative strategies to classical optimal control or reinforcement learning problems, and more importantly, realize the possibilities for solving new

problems with intrinsically non-cumulative objectives, which are frequently encountered in the fields such as network communications. This opens up directions for a broader range of applications for optimal control and reinforcement learning techniques.

APPENDIX A PROOF OF THEOREM 1

Lemma 1: If $g(\cdot, \cdot)$ satisfies the condition Eq. (11) in Theorem 1, then the generalized value function update F^* as in Eq. (10) is a contraction mapping.

Proof: In the following mathematical expressions, for the simplicity of notations, we use $p(r_t, s_{t+1}|s_t, a_t)$ as the shorthand notation for the joint distribution of $p_{R_t|S_t, A_t}(r_t|s_t, a_t)$ and $p_{S_{t+1}|S_t, A_t}(s_{t+1}|s_t, a_t)$. We also assume $p(r_t, s_{t+1}|s_t, a_t)$ is a discrete distribution. For continuous distribution, the proof still holds with summations substituted by integrations when computing the expectations.

For any pair of value functions $\forall Q^1, Q^2 \in \mathcal{R}^{|\mathcal{V}|}$, we have:

$$\begin{aligned} & \|F^*Q^1 - F^*Q^2\|_\infty \\ &= \max_{s_t, a_t} \left| (F^*Q^1)(s_t, a_t) - (F^*Q^2)(s_t, a_t) \right| \end{aligned} \quad (33)$$

$$\begin{aligned} &= \max_{s_t, a_t} \left| \sum_{r_t, s_{t+1}} p(r_t, s_{t+1}|s_t, a_t) g(r_t, \gamma \max_{a_{t+1}} Q^1(s_{t+1}, a_{t+1})) \right. \\ &\quad \left. - \sum_{r_t, s_{t+1}} p(r_t, s_{t+1}|s_t, a_t) g(r_t, \gamma \max_{a_{t+1}} Q^2(s_{t+1}, a_{t+1})) \right| \end{aligned} \quad (34)$$

$$\begin{aligned} &= \max_{s_t, a_t} \left| \sum_{r_t, s_{t+1}} p(r_t, s_{t+1}|s_t, a_t) [g(r_t, \gamma \max_{a_{t+1}} Q^1(s_{t+1}, a_{t+1})) \right. \\ &\quad \left. - g(r_t, \gamma \max_{a_{t+1}} Q^2(s_{t+1}, a_{t+1}))] \right| \end{aligned} \quad (35)$$

$$\begin{aligned} &\leq \max_{s_t, a_t} \sum_{r_t, s_{t+1}} p(r_t, s_{t+1}|s_t, a_t) \left| g(r_t, \gamma \max_{a_{t+1}} Q^1(s_{t+1}, a_{t+1})) \right. \\ &\quad \left. - g(r_t, \gamma \max_{a_{t+1}} Q^2(s_{t+1}, a_{t+1})) \right| \end{aligned} \quad (36)$$

$$\begin{aligned} &\leq \max_{s_t, a_t} \sum_{r_t, s_{t+1}} p(r_t, s_{t+1}|s_t, a_t) \left| \gamma \max_{a_{t+1}} Q^1(s_{t+1}, a_{t+1}) \right. \\ &\quad \left. - \gamma \max_{a_{t+1}} Q^2(s_{t+1}, a_{t+1}) \right| \end{aligned} \quad (37)$$

$$\begin{aligned} &\leq \gamma \max_{s_t, a_t} \sum_{r_t, s_{t+1}} p(r_t, s_{t+1}|s_t, a_t) \max_{a_{t+1}} |Q^1(s_{t+1}, a_{t+1}) \\ &\quad - Q^2(s_{t+1}, a_{t+1})| \end{aligned} \quad (38)$$

$$\leq \gamma \max_{s_t, a_t} \sum_{r_t, s_{t+1}} p(r_t, s_{t+1}|s_t, a_t) \|Q^1 - Q^2\|_\infty \quad (39)$$

$$\leq \gamma \|Q^1 - Q^2\|_\infty, \quad (40)$$

where (37) follows from the condition Eq. (11) in Theorem 1; (38) follows from the fact that for any two functions f_1 and f_2 , we have $|\sup_x f_1(x) - \sup_x f_2(x)| \leq \sup_x |f_1(x) - f_2(x)|$; and lastly, (40) follows from the normalization of the probability distribution $p(r_t, s_{t+1}|s_t, a_t)$. \square

With Lemma 1 established, we can readily prove the main theorem of the global convergence of the value function updates.

Starting from any arbitrary value function initialization point $Q^0 \in \mathcal{R}^{|\mathcal{V}|}$, consider the value iteration process of iteratively applying the mapping F^* . With $\mathcal{R}^{|\mathcal{V}|}$ being a Banach space and F^* being a contraction mapping (by Lemma 1), according to the Banach's fixed-point theorem [43], the process is guaranteed to converge to a unique convergence point Q^* .

APPENDIX B PROOF OF THEOREM 2

Lemma 2: If $g(\cdot, \cdot)$ satisfies the condition Eq. (16) in Theorem 2, then the generalized value function update F^* as in Eq. (10) is monotonic, i.e., $\forall Q^1, Q^2 \in \mathcal{R}^{|\mathcal{V}|}$, if $Q^1 \geq Q^2$, then $F^*Q^1 \geq F^*Q^2$ always holds.²

Proof:

$$\begin{aligned} & Q^1 \geq Q^2 \\ & \implies Q^1(s, a) \geq Q^2(s, a), \quad \forall s, a \end{aligned} \quad (41)$$

$$\implies \max_a Q^1(s, a) \geq \max_a Q^2(s, a), \quad \forall s \quad (42)$$

$$\implies g(r, \gamma \max_a Q^1(s, a)) \geq g(r, \gamma \max_a Q^2(s, a)), \quad \forall s, r \quad (43)$$

where Eq. (43) follows from the condition Eq. (16) in Theorem 2.

Now we introduce the time step into the equations, and consider any given state and action pair at time t : s_t and a_t . Let $s = s_{t+1}$ in Eq. (43) be the state the agent is in after executing a_t on s_t ; and let $r = r_t$ be the reward from executing a_t on s_t . We then have:

$$\begin{aligned} & Q^1 \geq Q^2 \\ & \implies g(r_t, \gamma \max_a Q^1(s_{t+1}, a)) \geq g(r_t, \gamma \max_a Q^2(s_{t+1}, a)), \\ & \quad \forall r_t, s_{t+1} \end{aligned} \quad (44)$$

$$\implies (F^*Q^1)(s_t, a_t) \geq (F^*Q^2)(s_t, a_t), \quad \forall s_t, a_t \quad (45)$$

$$\implies F^*Q^1 \geq F^*Q^2 \quad (46)$$

\square

To show that Q^* is the optimal point, let π_0 be an arbitrary initial policy and $Q^0 \in \mathcal{R}^{|\mathcal{V}|}$ be the corresponding value function (therefore $Q^0 = F^{\pi_0}Q^0$), we have:

$$Q^0 = F^{\pi_0}Q^0 \leq F^*Q^0 \quad (47)$$

where the inequality follows from the maximization over the actions by the greedy policy mapping F^* . Applying F^* again on both sides of the inequality, and by the monotonicity from Lemma 2 we have:

$$Q^0 \leq F^*Q^0 \leq F^{*2}Q^0 \quad (48)$$

where F^2 denotes iteratively applying the mapping F twice. After iteratively applying F^* until convergence, we arrive at

²The notation $Q^1 \geq Q^2$ implies $Q^1(s, a) \geq Q^2(s, a), \forall s, a$.

a chain of inequalities ending with the unique convergence point:

$$Q^0 \leq F^*Q^0 \leq F^{*2}Q^0 \leq F^{*3}Q^0 \leq \dots \leq \lim_{n \rightarrow \infty} F^{*n}Q^0 = Q^* \quad (49)$$

Since Q^0 is an arbitrary value function, we have shown that the unique point of convergence Q^* is indeed the global maximum value function.

Furthermore, given the assumptions that $p(r_t|s_t, a_t)$ and $p(s_{t+1}|s_t, a_t)$ are deterministic as required by Theorem 2, we have $g(\cdot, \cdot)$ being exchangeable with $\mathbb{E}_{p(r_t|s_t, a_t)}[\cdot]$ and $\mathbb{E}_{p(s_{t+1}|s_t, a_t)}[\cdot]$. By bringing the expectations inside the operation $g(\cdot, \cdot)$, we can see that Eq. (13) and Eq. (15) are equivalent. Correspondingly, the converged point of the value iteration is truly the expectation value of the objective function as defined in Eq. (7). Therefore, the greedy policy π^* derived from the value function Q^* is truly the global optimal policy.

APPENDIX C PROOF OF CONVERGENCE PROPERTIES ON THE BOTTLENECK UPDATE OPERATION

To show that the bottleneck update operation Eq. (18) satisfies the condition by Eq. (11) in Theorem 1:

Proof: Without loss of generality, we assume $b \geq c$, then we have:

- If $c \leq b \leq a$, then:

$$|g(a, b) - g(a, c)| = |\min(a, b) - \min(a, c)| = |b - c|. \quad (50)$$

- If $c \leq a < b$, then:

$$|g(a, b) - g(a, c)| = |\min(a, b) - \min(a, c)| \quad (51)$$

$$= a - c \quad (52)$$

$$< |b - c|. \quad (53)$$

- If $a < c \leq b$, then:

$$|g(a, b) - g(a, c)| = |\min(a, b) - \min(a, c)| \quad (54)$$

$$= 0 \quad (55)$$

$$< |b - c|. \quad (56)$$

□

To show that the bottleneck operation Eq. (18) satisfies the condition by Eq. (16) in Theorem 2:

Proof: Given $b \geq c$, we have:

- If $a \leq c$, then:

$$g(a, b) = \min(a, b) = a = \min(a, c) = g(a, c). \quad (57)$$

- If $c < a \leq b$, then:

$$g(a, b) = \min(a, b) = a > c = \min(a, c) = g(a, c). \quad (58)$$

- If $a > b$, then:

$$g(a, b) = \min(a, b) = b \geq c = \min(a, c) = g(a, c). \quad (59)$$

□

APPENDIX D PROOF OF CONVERGENCE PROPERTIES ON THE HARMONIC MEAN UPDATE OPERATION

To show that the harmonic mean operation Eq. (22) satisfies the condition by Eq. (11) in Theorem 1:

Proof: With our assumption of positive rewards, we have:

$$|g(a, b) - g(a, c)| = \left| \frac{1}{\frac{1}{a} + \frac{1}{b}} - \frac{1}{\frac{1}{a} + \frac{1}{c}} \right| \quad (60)$$

$$= \frac{|b - c|}{\left(\frac{1}{a} + \frac{1}{b}\right)\left(\frac{1}{a} + \frac{1}{c}\right)bc} \quad (61)$$

$$\leq \frac{|b - c|}{\left(\frac{1}{b}\right)\left(\frac{1}{c}\right)bc} \quad (62)$$

$$= |b - c| \quad (63)$$

□

To show that the harmonic mean operation Eq. (22) satisfies the condition by Eq. (16) in Theorem 2:

Proof: Given $b \geq c$. With our assumption of positive rewards, we have:

$$g(a, b) - g(a, c) = \frac{1}{\frac{1}{a} + \frac{1}{b}} - \frac{1}{\frac{1}{a} + \frac{1}{c}} \quad (64)$$

$$= \frac{b - c}{\left(\frac{1}{a} + \frac{1}{b}\right)\left(\frac{1}{a} + \frac{1}{c}\right)bc} \quad (65)$$

$$\geq 0 \quad (66)$$

□

REFERENCES

- [1] R. Bellman, "A Markovian decision process," *J. Math. Mech.*, vol. 6, no. 5, pp. 679–684, 1957.
- [2] R. Bellman, "On the theory of dynamic programming," *Proc. Nat. Acad. Sci. USA*, vol. 38, no. 8, pp. 716–719, 1952.
- [3] V. Mnih et al., "Human-level control through deep reinforcement learning," *Nature*, vol. 518, pp. 529–533, Feb. 2015.
- [4] H. Bühler, L. Gonon, J. Teichmann, and B. Wood, "Deep hedging," 2018, *arXiv:1802.03042*.
- [5] H. Cai et al., "Real-time bidding by reinforcement learning in display advertising," in *Proc. 10th ACM Int. Conf. Web Search Data Mining*, Feb. 2017, pp. 661–670.
- [6] W. Cui and W. Yu, "Scalable deep reinforcement learning for routing and spectrum access in physical layer," *IEEE Trans. Commun.*, vol. 69, no. 12, pp. 8200–8213, Dec. 2021.
- [7] H. Yang, Y. Ye, and X. Chu, "Max-min energy-efficient resource allocation for wireless powered backscatter networks," *IEEE Wireless Commun. Lett.*, vol. 9, no. 5, pp. 688–692, May 2020.
- [8] M. Bashar, K. Cumanan, A. G. Burr, M. Debbah, and H. Q. Ngo, "On the uplink max-min SINR of cell-free massive MIMO systems," *IEEE Trans. Wireless Commun.*, vol. 18, no. 4, pp. 2021–2036, Apr. 2019.
- [9] H. Xie, J. Xu, and Y. Liu, "Max-min fairness in IRS-aided multi-cell MISO systems with joint transmit and reflective beamforming," *IEEE Trans. Wireless Commun.*, vol. 20, no. 2, pp. 1379–1393, Feb. 2021.

- [10] M. Dashti, P. Azmi, and K. Navaie, "Harmonic mean rate fairness for cognitive radio networks with heterogeneous traffic," *Trans. Emerg. Telecommun. Technol.*, vol. 24, no. 2, pp. 185–195, Mar. 2013.
- [11] M. O. Hasna and M.-S. Alouini, "Harmonic mean and end-to-end performance of transmission systems with relays," *IEEE Trans. Commun.*, vol. 52, no. 1, pp. 130–135, Jan. 2004.
- [12] M. O. Hasna and M. S. Alouini, "End-to-end performance of transmission systems with relays over Rayleigh-fading channels," *IEEE Trans. Wireless Commun.*, vol. 2, no. 6, pp. 1126–1131, Nov. 2003.
- [13] G. Caire, R. R. Müller, and R. Knopp, "Hard fairness versus proportional fairness in wireless communications: The single-cell case," *IEEE Trans. Inf. Theory*, vol. 53, no. 4, pp. 1366–1385, Apr. 2007.
- [14] K. Shen and W. Yu, "Distributed pricing-based user association for downlink heterogeneous cellular networks," *IEEE J. Sel. Areas Commun.*, vol. 32, no. 6, pp. 1100–1113, Jun. 2014.
- [15] W. Cui and W. Yu, "A clustering approach to wireless scheduling," in *Proc. IEEE 21st Int. Workshop Signal Process. Adv. Wireless Commun. (SPAWC)*, May 2020, pp. 1–5.
- [16] D. Silver et al., "Mastering the game of go with deep neural networks and tree search," *Nature*, vol. 529, no. 7587, pp. 484–489, Jan. 2016.
- [17] D. Silver et al., "Mastering the game of go without human knowledge," *Nature*, vol. 550, no. 7676, pp. 354–359, Oct. 2017.
- [18] J. Baxter, A. Tridgell, and L. Weaver, "Learning to play chess using temporal differences," *Mach. Learn.*, vol. 40, no. 3, pp. 243–263, 2000.
- [19] S. P. Singh and R. S. Sutton, "Reinforcement learning with replacing eligibility traces," *Mach. Learn.*, vol. 22, nos. 1–3, pp. 123–158, 1996.
- [20] L. Kaelbling, "Hierarchical learning in stochastic domains: Preliminary results," in *Int. Conf. Mach. Learn. (ICML)*, 1993, pp. 167–173.
- [21] D. Blackwell, "Discrete dynamic programming," *Ann. Math. Statist.*, vol. 33, no. 2, pp. 719–726, Jun. 1962.
- [22] A. Jalali and M. Ferguson, "Computationally efficient adaptive control algorithms for Markov chains," in *Proc. 28th IEEE Conf. Decis. Control*, Dec. 1989, pp. 1283–1288.
- [23] A. Schwartz, "A reinforcement learning method for maximizing undiscounted rewards," in *Int. Conf. Mach. Learn. (ICML)*, Jul. 1993, pp. 298–305.
- [24] S. Mahadevan, "Average reward reinforcement learning: Foundations, algorithms, and empirical results," *Mach. Learn.*, vol. 22, nos. 1–3, pp. 159–195, Mar. 1996.
- [25] P. Tadepalli and D. Ok, "Model-based average reward reinforcement learning," *Artif. Intell.*, vol. 100, nos. 1–2, pp. 177–224, Apr. 1998.
- [26] S. Proper and P. Tadepalli, "Scaling model-based average-reward reinforcement learning for product delivery," in *Proc. Eur. Conf. Mach. Learn. (ECML)*, Sep. 2006, pp. 735–742.
- [27] C.-Y. Wei, M. J. Jahromi, H. Luo, H. Sharma, and R. Jain, "Model-free reinforcement learning in infinite-horizon average-reward Markov decision processes," in *Proc. Int. Conf. Mach. Learn. (ICML)*, 2020, pp. 10170–10180.
- [28] K. H. Quah and C. Quek, "Maximum reward reinforcement learning: A non-cumulative reward criterion," *Expert Syst. Appl.*, vol. 31, no. 2, pp. 351–359, Aug. 2006.
- [29] S. K. Gottipati et al., "Maximum reward formulation in reinforcement learning," 2020, *arXiv:2010.03744*.
- [30] K. H. Lundberg and T. W. Barton, "History of inverted-pendulum systems," *IFAC Proc. Volumes*, vol. 42, no. 24, pp. 131–135, 2010.
- [31] R. S. Sutton, "Learning to predict by the methods of temporal differences," *Mach. Learn.*, vol. 3, no. 1, pp. 9–44, Aug. 1988.
- [32] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods*. Upper Saddle River, NJ, USA: Prentice-Hall, 1989.
- [33] G. A. Rummery and M. Niranjan, "On-line Q -learning using connectionist systems," Dept. Eng., Cambridge Univ., Cambridge, U.K., Tech. Rep. CUED/F-INFENG/TR 166, Sep. 1994.
- [34] C. J. C. H. Watkins and P. Dayan, " Q -learning," *Mach. Learn.*, vol. 8, nos. 3–4, pp. 279–292, May 1992.
- [35] V. R. Konda and J. N. Tsitsiklis, "Actor-critic algorithms," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, 2000, pp. 1–7.
- [36] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*, 3rd ed. Cambridge, MA, USA: MIT Press, Jun. 2009.
- [37] T. Jaakkola, M. I. Jordan, and S. P. Singh, "On the convergence of stochastic iterative dynamic programming algorithms," *Neural Comput.*, vol. 6, no. 6, pp. 1185–1201, Nov. 1994.
- [38] C. Szepesvári, *Algorithms for Reinforcement Learning*. San Rafael, CA, USA: Morgan and Claypool, Jun. 2010.
- [39] N. Metropolis and S. Ulam, "The Monte Carlo method," *J. Amer. Statist. Assoc.*, vol. 44, no. 247, pp. 335–341, Sep. 1949.
- [40] H. van Hasselt, A. Guez, and D. Silver, "Deep reinforcement learning with double Q -learning," in *Proc. Assoc. Adv. Artif. Intell. (AAAI)*, Feb. 2016, pp. 1–7.
- [41] T. Schaul, J. Quan, I. Antonoglou, and D. Silver, "Prioritized experience replay," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, May 2016, pp. 1–21.
- [42] Z. Wang, T. Schaul, M. Hessel, H. van Hasselt, M. Lanctot, and N. de Freitas, "Dueling network architectures for deep reinforcement learning," *Proc. Mach. Learn. Res.*, vol. 48, pp. 1995–2013, Jun. 2016.
- [43] S. Banach, "Sur les opérations dans les ensembles abstraits et leur application aux équations intégrales," *Fundamenta Mathematicae*, vol. 3, no. 1, pp. 133–181, 1922.



WEI CUI (Graduate Student Member, IEEE) received the B.A.Sc. degree in Engineering Science, the M.A.Sc. degree in Electrical and Computer Engineering, and the Ph.D. degree in Electrical and Computer Engineering from University of Toronto, Toronto, ON, Canada, in 2017, 2019, and 2023, respectively.

His research interests include machine learning, optimization, and wireless communication.



WEI YU (Fellow, IEEE) received the B.A.Sc. degree in computer engineering and mathematics from the University of Waterloo, Waterloo, ON, Canada, and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, USA. He is currently a Professor and Canada Research Chair in Information Theory and Wireless Communications with the Department of Electrical and Computer Engineering, University of Toronto, Toronto, ON, Canada. He is

a Fellow of the Canadian Academy of Engineering and a member of the College of New Scholars, Artists, and Scientists, The Royal Society of Canada. He was a recipient of the Steacie Memorial Fellowship in 2015, the IEEE Marconi Prize Paper Award in Wireless Communications in 2019, the IEEE Communications Society Award for Advances in Communication in 2019, the IEEE Signal Processing Society Best Paper Award in 2008, 2017, and 2021, the *Journal of Communications and Networks* Best Paper Award in 2017, and the IEEE Communications Society Best Tutorial Paper Award in 2015. He served as the President of the IEEE Information Theory Society in 2021. He served as the Chair of the Signal Processing for Communications and Networking Technical Committee, IEEE Signal Processing Society, from 2017 to 2018. He has served as an Area Editor for the IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS, an Associate Editor for IEEE TRANSACTIONS ON INFORMATION THEORY, and an Editor for the IEEE TRANSACTIONS ON COMMUNICATIONS and IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS. He was an IEEE Communications Society Distinguished Lecturer from 2015 to 2016.