

# A Provably-Efficient Model-Free Algorithm for Infinite-Horizon Average-Reward Constrained Markov Decision Processes

Honghao Wei,<sup>1</sup> Xin Liu,<sup>2</sup> Lei Ying<sup>1</sup>

<sup>1</sup> University of Michigan, Ann Arbor

<sup>2</sup> ShanghaiTech University

honghaow@umich.edu, liuxin7@shanghaitech.edu.cn, leiying@umich.edu

## Abstract

This paper presents a model-free reinforcement learning (RL) algorithm for infinite-horizon average-reward Constrained Markov Decision Processes (CMDPs). Considering a learning horizon  $K$ , which is sufficiently large, the proposed algorithm achieves  $\tilde{O}\left(\frac{\sqrt{SA}\kappa}{\delta}K^{\frac{5}{6}}\right)$  regret and zero constraint violation, where  $S$  is the number of states,  $A$  is the number of actions, and  $\kappa$  and  $\delta$  are two constants independent of the learning horizon  $K$ .

## Introduction

Reinforcement Learning has gained significant attention because of its successes in board games and video games such as Go (Silver et al. 2017) and StarCraft (Vinyals et al. 2019), and in highly-complex robotics systems (Andrychowicz et al. 2020). An agent’s objective in a typical RL problem is to maximize the cumulative reward through interacting with an unknown environment. In board games or video games, the outcomes of a random action are not consequential to the users (e.g. not life-threatening). However, a careless action in an engineering system might have catastrophic outcomes such as collisions and fatalities in robotics and autonomous driving (Ono et al. 2015; Garcia and Fernández 2012; Fisac et al. 2018) or surgical robotics (Richter, Orosco, and Yip 2019). Therefore, it is critical to strike a balance between reward maximization and safety in real-world applications. A standard formulation for RL with constraints is the Constrained Markov Decision Processes framework (Altman 1999), in which the agent aims at learning a policy that maximizes the expected cumulative reward under safety constraints during and after learning.

A CMDP problem can be solved using linear programming (LP) or the dual approach (Altman 1999) when the model (the transition kernel) is given. For example, (Zheng and Ratliff 2020) proposed a LP-based algorithm which learns the optimal policy while satisfying the constraints for a CMDP with a known model. Recent model-based algorithms (Singh, Gupta, and Shroff 2020; Brantley et al. 2020; Kalagarla, Jain, and Nuzzo 2021; Efroni, Mannor, and Pirotta 2020; Qiu et al. 2020) follow a similar approach but

learn the models from the data samples collected. This approach has also been utilized for CMDPs with linear function approximation (Ding et al. 2021) under the assumption that the transition kernel is linear. Leveraging the estimated model, the CMDPs can be solved approximately as long as the estimate becomes more and more accurately. The works mentioned above are proved to achieve sublinear constraint violation. A detailed discussion on the sample complexity of CMDPs of model-based approaches can be found in (HasanzadeZonuzi, Kalathil, and Shakkottai 2021). While model-based RL algorithms are sample efficient, they need to continuously solve LPs when the estimated models are updated, so these algorithms are often computationally inefficient and require a large memory to maintain a large number of model parameters.

Model-free algorithms, on the other hand, learn state or action value functions, instead of transition kernels, so require significantly less memory space and have lower computational complexity. In (Borkar 2005), the author proposes an actor-critic RL algorithm and shows its asymptotic global convergence using multi-timescale stochastic approximation theory for infinite-horizon average-reward CMDPs when the model is unknown. Policy gradient approaches have also been developed (Tessler, Mankowitz, and Mannor 2018; Stooke, Achiam, and Abbeel 2020; Yang et al. 2020) and seen successes in practice for solving constrained RL problems, though they lack regret and constraint violation analysis. (Ding et al. 2020; Xu, Liang, and Lan 2020; Chen, Dong, and Wang 2021) show that sublinear regrets and constraint violations are achievable when policy “simulators” (or generative models) are given. Two very recent works (Liu et al. 2021a; Wei, Liu, and Ying 2021) show that sublinear regret bound and zero violation are possible for episodic CMDPs without simulators. In particular, (Liu et al. 2021a) proposes a model-based algorithm and (Wei, Liu, and Ying 2022) presents a model-free algorithm, and (Bura et al. 2021) proves that it is possible to achieve zero violation during training given a safe baseline policy based on a model-based approach. Despite these significant developments, the following question is still open:

*Can we design efficient RL algorithms for infinite-horizon, average-reward CMDPs with provably regret and constraint violation guarantees?*

We answer this question affirmatively and present a

	Algorithm	Regret	Constraint Violations
Known Model	C-UCRL (Zheng and Ratliff 2020)	$\tilde{\mathcal{O}}(SA\sqrt{K}^{1.5})$	0
Model-based	UCRL-CMDP (Singh, Gupta, and Shroff 2020)	$\tilde{\mathcal{O}}(S\sqrt{AK}^{\frac{2}{3}})$	$\tilde{\mathcal{O}}(S\sqrt{AK}^{\frac{2}{3}})$
Known Model	CMDP-PSRL (Agarwal, Bai, and Aggarwal 2021)	$\tilde{\mathcal{O}}(\text{poly}(SAD)\sqrt{K})$	$\tilde{\mathcal{O}}(\text{poly}(SA)\sqrt{K})$
Model-based*	OptPess-LP (Liu et al. 2021a)	$\tilde{\mathcal{O}}(H^3\sqrt{S^3AK})$	0
Model-based*	OptPess-PrimalDual (Liu et al. 2021a)	$\tilde{\mathcal{O}}(H^3\sqrt{S^3AK})$	$\mathcal{O}(1)$
Model-based*	OPSRL(Bura et al. 2021)	$\tilde{\mathcal{O}}(\sqrt{S^4H^7AK})$	0
Model-free	<b>This Paper</b>	$\tilde{\mathcal{O}}\left(\frac{\sqrt{SA}}{\delta}K^{\frac{5}{6}}\right)$	0

Table 1: Regrets and constraint violations of RL algorithms for infinite-horizon average-reward CMDPs (\* These algorithms are designed for episodic CMDPs).  $S$  is the number of states,  $A$  is the number of actions,  $K$  is the number of steps,  $D$  is the diameter of the CMDP whose definition can be found in the supplementary material,  $\delta$  is the slackness that will be defined later (Eq. (10)), and  $\text{poly}(X)$  denotes a polynomial function of  $x$ . Throughout the paper, we use the notation  $\tilde{\mathcal{O}}$  to suppress log terms.  $\tilde{\mathcal{O}}(f(K))$  denotes  $\mathcal{O}(f(K)\log^n K)$  with  $n > 0$ .

model-free RL algorithm that achieves sub-linear regret and zero constraint violation. Table 1 compares the results in this paper with those in the literature. We remark that the proposed algorithm synthesizes the Triple-Q algorithm in (Wei, Liu, and Ying 2022) for episodic CMDPs and Optimistic Q-Learning (Wei et al. 2020) that reduces the average-reward problem to a discounted reward problem.

### Preliminaries

An infinite-horizon average-reward CMDP can be defined as  $(\mathcal{S}, \mathcal{A}, r, g, p)$ , where  $\mathcal{S}$  is the finite state space,  $\mathcal{A}$  is the finite action space,  $r(g) : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$  is the unknown reward (utility) function, and  $p : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$  is the transition kernel such that  $p(s'|s, a) := \mathbb{P}(s_{k+1} = s' | s_k = s, a_k = a)$  for  $s_k \in \mathcal{S}, a_k \in \mathcal{A}$ . A stationary policy is a mapping  $\pi : \mathcal{S} \rightarrow \mathcal{A}$ , the long-term average reward (reward rate) of a stationary policy  $\pi$  with initial state  $s \in \mathcal{S}$  is defined as

$$J_r^\pi(s) := \lim_{K \rightarrow \infty} \frac{1}{K} \mathbb{E} \left[ \sum_{k=1}^K r(s_k, \pi(s_k)) \middle| s_1 = s \right],$$

and the long-term average utility (utility rate) is defined as

$$J_g^\pi(s) := \lim_{K \rightarrow \infty} \frac{1}{K} \mathbb{E} \left[ \sum_{k=1}^K g(s_k, \pi(s_k)) \middle| s_1 = s \right].$$

We assume that under any stationary policy,  $s_k$  is an irreducible an aperiodic Markov chain, so it has a unique stationary distribution and the limits above are well defined. Letting  $s_\infty^\pi$  denote the Markov chain at steady-state under policy  $\pi$ , we have

$$J_r^\pi = \mathbb{E} [r(s_\infty, \pi(s_\infty))] \text{ and } J_g^\pi = \mathbb{E} [g(s_\infty, \pi(s_\infty))],$$

where we removed the dependence on the initial condition  $s$  because the stationary distribution is independent of the initial condition for a finite-state, irreducible and aperiodic Markov chain.

An optimal stationary policy  $\pi^*$  is defined to be the solution of the following problem:

$$\max_{\pi} J_r^\pi \quad \text{s.t.} \quad J_g^\pi \geq \rho. \quad (1)$$

This paper considers a constrained RL problem with  $K$  steps. At each step  $k$ , the agent observes state  $s_k$ , takes an action  $a_k$ , and receives reward  $r(s_k, a_k)$  and utility  $g(s_k, a_k)$ . The next state  $s_{k+1}$  is sampled according to the probability distribution  $p(\cdot | s_k, a_k)$ . Our goal is to develop an online RL algorithm, which may be nonstationary, that minimizes both the regret and the constraint violation defined below.

$$\text{Regret}(K) = \mathbb{E} \left[ \sum_{k=1}^K \left( J_r^{\pi^*} - r(s_k, a_k) \right) \right], \quad (2)$$

$$\text{Violation}(K) = \mathbb{E} \left[ \sum_{k=1}^K (\rho - g(s_k, a_k)) \right]. \quad (3)$$

When the transition kernel  $p(s' | s, a)$  is known, the optimal stationary policy that solves problem (1) can be obtained by solving the following LP problem (Altman 1999):

$$\max_{\{q(s,a):(s,a) \in \mathcal{S} \times \mathcal{A}\}} \sum_{s,a} q(s,a)r(s,a) \quad (4)$$

$$\text{s.t.} \quad \sum_{s,a} q(s,a)g(s,a) \geq \rho, \forall s \in \mathcal{S}, \forall a \in \mathcal{A} \quad (5)$$

$$q(s,a) \geq 0, \forall s \in \mathcal{S}, \forall a \in \mathcal{A} \quad (6)$$

$$\sum_{s,a} q(s,a) = 1 \quad (7)$$

$$\sum_a q(s,a) = \sum_{s',a'} p(s|s',a')q(s',a'), \quad (8)$$

where the  $q(s, a)$  is called the occupancy measure, which is defined as the set of distributions generated by executing the associated induced policy  $\pi$  in the infinite-horizon CMDP.  $\sum_a q(s, a)$  represents the probability the system is in state  $s$ , and  $\frac{q(s,a)}{\sum_{a'} q(s,a')}$  is the probability of taking action  $a$  in state  $s$ . The utility constraint is represented in (5). More details can be found in (Altman 1999).

To analyze the performance of our algorithm, we need to consider a tightened version of the above LP problem later,

which is defined below:

$$\begin{aligned} & \max_{\{q(s,a):(s,a) \in \mathcal{S} \times \mathcal{A}\}} \sum_{s,a} q(s,a)r(s,a) \quad (9) \\ \text{s.t. } & \sum_{s,a} q(s,a)g(s,a) \geq \rho + \epsilon, \forall s \in \mathcal{S}, \forall a \in \mathcal{A} \\ & (6) - (8), \end{aligned}$$

where  $\epsilon > 0$  is called a tightness constant. As in previous works (Ding et al. 2021, 2020; Efroni, Mannor, and Pirota 2020; Paternain et al. 2019), we make the following standard assumption of Slater's condition.

**Assumption 1.** (Slater's Condition). *There exist  $\delta > 0$  and a feasible solution  $q(s, a)$  to the LP such that*

$$\sum_{s,a} q(s, a)g(s, a) - \rho \geq \delta. \quad (10)$$

It is obvious that when  $\epsilon < \delta$  the problem (9) has a feasible solution due to Slater's condition. The Slater's condition is commonly assumed in previous works to ensure the LP problem has strong duality, see proofs in (Paternain et al. 2019, 2022). Unlike (Ding et al. 2021; Efroni, Mannor, and Pirota 2020), which assume  $\delta$  is known, and (Achiam et al. 2017; Liu et al. 2021a; Bura et al. 2021), where a strictly feasible policy is given, our assumption is less restrictive. Let

$$J_r^* = \sum_{s,a} q^*(s, a)r(s, a), \quad (11)$$

$$J_g^* = \sum_{s,a} q^*(s, a)g(s, a). \quad (12)$$

be the optimal reward rate and utility rate, where  $q^*(s, a)$  is the optimal solution obtained by solving the LP problem (4). Moreover it is obvious that  $J_r^*$  and  $J_g^*$  are independent of the initial state and we have  $J_r^* = J_r^{\pi^*}$  and  $J_g^* = J_g^{\pi^*}$ .

In the following, we use superscript  $*$  to denote the optimal policy achieved by solving the LP (4) of the original CMDP, and superscript  $\epsilon, *$  to denote the optimal policy related to the  $\epsilon$ -tightened version of LP (9).

## Primal-Dual-Based, Two-Time-Scale Optimistic SARSA (Triple-QA)

In this section, we introduce our algorithm Triple-QA (see Algorithm 1 for pseudo-code) which achieves sub-linear regret and zero constraint violation. The algorithm is inspired by Triple-Q, an algorithm for episodic CMDPs in (Wei, Liu, and Ying 2021). Triple-QA is for the infinite-horizon average-reward CMDPs with a different update rule. The algorithm further solves the discounted CMDPs with the discount factor  $\gamma$  close to one, an idea used in (Wei et al. 2020). The discounted CMDP is defined on the same state space, action space, reward/utility functions, the transition kernel. The intuition is that the reward of the discounted problem (scaled by  $1 - \gamma$ ) approaches to that of the average reward problem as  $\gamma$  goes to 1.

---

### Algorithm 1: Triple-QA

---

- 1: Initialize  $Q_1(s, a) = \hat{Q}_1(s, a) \leftarrow H = K^{\frac{1}{6}}$  and  $n_1(s, a) \leftarrow 0, \forall (s, a) \in \mathcal{S} \times \mathcal{A}, \gamma = 1 - \frac{1}{H}, \hat{V}_1(s) = H, \forall s \in \mathcal{S}$
  - 2: Choose  $\chi = K^{\frac{1}{3}}, \eta = K^{\frac{1}{6}}, \iota = 8 \log(\sqrt{2}K), \beta = \frac{2}{3}$ .
  - 3: Choose  $\epsilon = \frac{9\kappa\sqrt{SA\iota}}{K^{\frac{1}{6}}}, \kappa(Eq.(17))$ .
  - 4: Initialize  $\bar{C} \leftarrow 0, Z_1 \leftarrow 0$ .
  - 5: Define  $\alpha_\tau = \frac{\chi+1}{\chi+\tau}, b_\tau = \kappa\sqrt{\frac{(\chi+1)\iota}{\chi+\tau}}$ .
  - 6: **for** episode  $k = 1, \dots, K$  **do**
  - 7:   Take  $a_k = \arg \max_a \left( \hat{Q}_k(s_k, a) + \frac{Z}{\eta} \hat{C}_k(s_k, a) \right)$ .
  - 8:   Observe  $s_{k+1}$ .
  - 9:    $n_{k+1}(s_k, a_k) \leftarrow n_k(s_k, a_k) + 1, \tau \leftarrow n_{k+1}(s_k, a_k)$ .
  - 10:   Update  $Q_{k+1}(s_k, a_k) \leftarrow (1 - \alpha_\tau)Q_k(s_k, a_k)$
  - 11:      $+ \alpha_\tau[r(s_k, a_k) + \gamma \hat{V}_k(s_{k+1}) + b_\tau]$ ,
  - 12:   Update  $C_{k+1}(s_k, a_k) \leftarrow (1 - \alpha_\tau)C_k(s_k, a_k)$
  - 13:      $+ \alpha_\tau[g(s_k, a_k) + \gamma \hat{W}_k(s_{k+1}) + b_\tau]$ .
  - 14:   **if**  $Q_{k+1}(s_k, a_k) \leq \hat{Q}_k(s_k, a_k)$  and  $C_{k+1}(s_k, a_k) \leq \hat{C}_k(s_k, a_k)$  **then**
  - 15:      $\hat{Q}_{k+1}(s_k, a_k) \leftarrow Q_{k+1}(s_k, a_k)$
  - 16:      $\hat{C}_{k+1}(s_k, a_k) \leftarrow C_{k+1}(s_k, a_k)$
  - 17:   **else**
  - 18:      $\hat{Q}_{k+1}(s_k, a_k) \leftarrow \hat{Q}_k(s_k, a_k)$
  - 19:      $\hat{C}_{k+1}(s_k, a_k) \leftarrow \hat{C}_k(s_k, a_k)$
  - 20:    $\bar{C} \leftarrow \bar{C} + (1 - \gamma)\hat{C}_k(s_k, a_k)$
  - 21:    $a' = \arg \max_a \left( \hat{Q}_{k+1}(s_k, a) + \frac{Z}{\eta} \hat{C}_{k+1}(s_k, a) \right)$
  - 22:    $\hat{V}_{k+1}(s_k) \leftarrow \hat{Q}_{k+1}(s_k, a')$
  - 23:    $\hat{W}_{k+1}(s_k) \leftarrow \hat{C}_{k+1}(s_k, a')$
  - 24:   **if**  $k \bmod K^\beta = 0$  **then**
  - 25:      $Z \leftarrow \left( Z + \rho + \epsilon - \frac{\bar{C}}{K^\beta} \right)$
  - 26:     Reset  $\bar{C} \leftarrow 0, n_t(s, a) \leftarrow 0$ .
  - 27:     Reset  $\hat{Q}_{k+1}(s, a), Q_{k+1}(s, a), V_{k+1}(s)$  to  $H$
  - 28:     Reset  $\hat{C}_{k+1}(s, a), C_{k+1}(s, a), W_{k+1}(s)$  to  $H$
- 

Under the discounted CMDP setting, given a policy  $\pi$ , the reward value function  $V_k^\pi$  at step  $k$  is the expected cumulative rewards from step  $k$  under policy  $\pi$  :

$$V_k^\pi(s) = \mathbb{E} \left[ \sum_{i=k}^{\infty} \gamma^{i-k} r(s_i, \pi(s_i)) \middle| s_k = s \right].$$

The reward  $Q$ -function  $Q_k^\pi(s, a)$  at step  $k$  is the expected cumulative rewards when agent starts from a state-action pair  $(s, a)$  at step  $k$  and then follows policy  $\pi$  :

$$Q_k^\pi(s, a) = r(s, a) + \mathbb{E} \left[ \sum_{i=k}^{\infty} \gamma^{i-k} r(s_i, \pi(s_i)) \middle| \begin{matrix} s_k = s \\ a_k = a \end{matrix} \right].$$

Similarly, we use  $W_k^\pi(s) : \mathcal{S} \rightarrow \mathbb{R}^+$  and  $C_k^\pi(s, a) : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^+$  to denote the utility value function and utility  $Q$ -function at step  $k$ :

$$W_k^\pi(x) = \mathbb{E} \left[ \sum_{i=k}^{\infty} \gamma^{i-k} g(s_i, \pi(s_i)) \middle| s_k = s \right],$$

$$C_k^\pi(s, a) = g(s, a) + \mathbb{E} \left[ \sum_{i=k}^{\infty} \gamma^{i-k} g(s_i, \pi(s_i)) \mid \begin{array}{l} s_k = s \\ a_k = a \end{array} \right].$$

It is obvious that all the reward and utility value (Q-value) functions are bounded by  $\frac{1}{1-\gamma}$  because the reward and utility are bounded by 1. We define  $H = \frac{1}{1-\gamma}$ . Then given a state-action pair  $(s, a)$  at step  $k$ , our algorithm updates the estimate of reward (utility) Q-value functions of the discounted CMDP setting instead.

The design of the algorithm is based on the primal-dual approach for constrained optimization problems. Suppose that  $V^\pi(s)$  ( $W^\pi(s)$ ) is an accurate estimate of  $\frac{J_r^\pi}{1-\gamma}$  ( $\frac{J_g^\pi}{1-\gamma}$ ). The formal proof is deferred to the next section. Given Lagrangian multiplier  $\mu$ , we consider the following problem:

$$\begin{aligned} & \max_{\pi} J_r^\pi(s) + \mu(J_g^\pi(s) - \rho) \\ & \approx \max_{\pi} (1-\gamma)(V^\pi(s) + \mu W^\pi(s)) - \mu\rho \end{aligned}$$

which can be interpreted as an unconstrained MDP with a modified reward function  $(1-\gamma)(r + \mu g)$ .

The algorithm is an extension of Triple-Q (Wei, Liu, and Ying 2021) for episodic CMDPs by including the discount factor and replacing episode-by-episode updates with step-by-step updates. We adopt the same notations used in (Wei et al. 2020). Same as Triple-Q, the algorithm maintains an estimate  $\hat{V}_k(s)$  ( $\hat{W}_k(s)$ ) for the optimal value function  $V^*(s)$  ( $W^*(s)$ ) and  $\hat{Q}_k(s, a)$  ( $\hat{C}_k(s, a)$ ) for the optimal Q-function  $Q^*(s, a)$  ( $C^*(s, a)$ ). At each step  $k$ , after observing state  $s$ , the agent selects action  $a_k^*$  based on the combined Q-value:

$$a_k^* \in \arg \max_a \hat{Q}_k(s, a) + \frac{Z}{\eta} \hat{C}_k(s, a), \quad (13)$$

where  $\frac{Z}{\eta}$  can be treated as an estimate of the Lagrange multiplier  $\mu$ . Similar to (Wei, Liu, and Ying 2021), we need to carefully tune the frequency of updating the Lagrange multiplier to balance the convergence and optimality. Updating it too frequent would lead to divergence and too infrequent would result in a large regret and large constraint violation. The algorithm tackles this difficulty by updating  $Z$  at a slow time-scale, i.e., every  $K^\beta$  steps in line 25 – 26 in Algorithm 1, with the following update

$$Z \leftarrow \left( Z + \rho + \epsilon - \frac{\bar{C}}{K^\beta} \right)^+, \quad (14)$$

where  $(x)^+ = \max\{x, 0\}$ , and  $\bar{C}$  is the summation of all  $(1-\gamma)\hat{C}_k(s_k, a_k)$  of the steps in the previous frame, where each frame consists of  $K^\beta$  consecutive steps.

During each frame, the algorithm learns the combined Q functions for fixed  $Z$  at a fast time scale. The estimates of reward and utility value functions are updated after observing a new state-action pair.

It is important to note that for a CMDP,

$$V^*(s) \neq \max_a Q^*(s, a). \quad (15)$$

This means optimistic Q-learning algorithms for unconstrained MDPs (e.g. (Jin et al. 2018; Wei et al. 2020; Dong

et al. 2019)) cannot be used for estimating the optimal value functions of CMDPs. Instead, Triple-Q (Wei, Liu, and Ying 2021) and this algorithms use a SARSA-type updating rule, as shown in line 11 – 14.

We note that the optimal policy for a CMDP is stochastic in general. The policy under our algorithm is a stochastic policy because the virtual queue  $Z$  varies during and after the learning process, which results in a stochastic policy.

We further introduce additional notations before presenting our main theorem. Let  $v^\pi(s)$  and  $w^\pi(s)$  denote the reward and utility relative value functions for state  $s$  under average-reward setting, and  $q^\pi(s, a)$ ,  $c^\pi(s, a)$  be the reward and utility Q value functions for any state-action pair  $(s, a)$ . Based on the Bellman equation, we have

$$\begin{aligned} J_r^\pi + q^\pi(s, a) &= r(s, a) + \mathbb{E}_{s' \sim p(\cdot | s, a)} [v^\pi(s')] \\ v^\pi(s) &= \sum_a q^\pi(s, a) \mathbb{P}(\pi(s) = a) \\ J_g^\pi + c^\pi(s, a) &= g(s, a) + \mathbb{E}_{s' \sim p(\cdot | s, a)} [w^\pi(s')] \\ w^\pi(s) &= \sum_a c^\pi(s, a) \mathbb{P}(\pi(s) = a) \end{aligned}$$

Define

$$sp(f) = \max_{s \in \mathcal{S}} f(s) - \min_{s \in \mathcal{S}} f(s) \quad (16)$$

to be the span of the function  $f$ . It is well known that the span of the optimal reward relative value function  $sp(v^*)$  and utility relative value function  $sp(w^*)$  are bounded for weakly communication or ergodic MDPs. In particular, they are bounded by the diameter of the MDP (Lattimore and Szepesvári 2020).

Let

$$\kappa = \max_{0 \leq \epsilon \leq \rho/2} (\max\{sp(v^{\epsilon,*}), sp(w^{\epsilon,*}), 1\}) \quad (17)$$

and assume that  $\kappa$  which is used in the algorithm is known beforehand as in (Wei et al. 2020, 2021). We can always substitute them with any upper bound (e.g. the diameter) when it is unknown. We now state the main results in the following theorem.

**Theorem 1.** Assume  $K \geq \left(\frac{18\kappa\sqrt{SA\iota}}{\delta}\right)^6$  and let  $\epsilon = \frac{9\kappa\sqrt{SA\iota}}{K^{\frac{1}{6}}}$  such that  $\epsilon \leq \frac{\delta}{2}$ . By choosing  $m = K^{\frac{1}{6}} \log K$ ,  $H = K^{\frac{1}{6}}$ ,  $\eta = K^{\frac{1}{6}}$ ,  $\chi = K^{\frac{1}{3}}$ , and  $\beta = \frac{2}{3}$ , Algorithm 1 guarantees

$$\text{Regret}(K) \leq \tilde{O} \left( \frac{\sqrt{SA\kappa}}{\delta} K^{\frac{5}{6}} \right)$$

$$\text{Violation}(K) \leq \frac{92K^{\frac{2}{3}}}{\delta} \log \left( \frac{24}{\delta} \right) - \sqrt{SA\iota} K^{\frac{5}{6}} = 0,$$

where  $\iota = 32 \log(\sqrt{2}K)$ .  $\square$

## Proof of the Main Theorem

### Notations

Throughout the paper, we use shorthand notation

$$\{f - g\}(x) = f(x) - g(x),$$

where  $f(\cdot)$  and  $g(\cdot)$  the the same argument value. Similarly,

$$\{(f - g)q\}(x) = (f(x) - g(x))q(x).$$

Due to the page limit, we will only present several key lemmas and the key intuitions in this section. The complete proof can be found in the appendix.

### Regret Analysis

We start the proof by adding and subtracting the corresponding terms to the regret defined in (2), we obtain

$$\begin{aligned} \text{Regret}(K) &= \mathbb{E} \left[ \sum_{k=1}^K (J_r^* - r(s_k, a_k)) \right] \\ &= \mathbb{E} \left[ \sum_{k=1}^K (J_r^* - J_r^{\epsilon, *}) \right] \end{aligned} \quad (18)$$

$$+ \mathbb{E} \left[ \sum_{k=1}^K (J_r^{\epsilon, *} - (1 - \gamma)V^{\epsilon, *}(s_k)) \right] \quad (19)$$

$$+ \mathbb{E} \left[ \sum_{k=1}^K (1 - \gamma) (V^{\epsilon, *}(s_k) - \hat{Q}_k(s_k, a_k)) \right] \quad (20)$$

$$+ \mathbb{E} \left[ \sum_{k=1}^K ((1 - \gamma)\hat{Q}_k(s_k, a_k) - r(s_k, a_k)) \right]. \quad (21)$$

We will bound each of the four terms above in the following sequence of lemmas.

Term (18) is the difference between the original CMDP and its corresponding  $\epsilon$ -tighten version which is a perturbation of the original problem. We establish a bound by using the following lemma. The proof is deferred to supplementary material.

**Lemma 2.** *Under assumption 1, given  $\epsilon \leq \delta$ , we have*

$$\sum_{t=1}^K (J_r^* - J_r^{\epsilon, *}) \leq \frac{\epsilon K}{\delta} \quad (22)$$

For the second term (19), we establish a bound by using Lemma 3, which shows the difference between value functions of the average-reward problem and the value functions of the discounted setting problem is small. The proof is based on the Bellman equations under teh two settings. The proof follows Lemma 2 in (Wei et al. 2020) closely.

**Lemma 3.** *For an arbitrary policy  $\pi$ , we have*

$$J_r^\pi - (1 - \gamma)V^\pi(s) \leq (1 - \gamma)sp(v^\pi(s)), \quad (23)$$

$$|V^\pi(s_1) - V^\pi(s_2)| \leq 2sp(v^\pi(s)); \quad (24)$$

$$J_g^\pi - (1 - \gamma)W^\pi(s) \leq (1 - \gamma)sp(w^\pi(s)), \quad (25)$$

$$|W^\pi(s_1) - W^\pi(s_2)| \leq 2sp(w^\pi(s)), \quad (26)$$

where  $V^\pi(s)$  is the value function for the discounted setting under policy  $\pi$ , and  $J_r^\pi(J_g^\pi)$  is the reward (utility) rate under policy  $\pi$ .

Then it is easy to obtain

$$J_r^{\epsilon, *} - (1 - \gamma)V^{\epsilon, *}(s) \leq (1 - \gamma)\kappa, \quad (27)$$

Next we establish a bound on term (20) by using the Lyapunov-drift analysis. In unconstrained MDPs, the bound is established by showing that optimistic Q-learning guarantees that  $\hat{Q}_k(s, a)$  is an overestimate of  $Q^*(s, a)$ . However this does not hold in CMDPs because the algorithm needs to consider reward and utility simultaneously so  $\hat{Q}_k(s, a)$  is not necessarily an overestimate of  $Q^*(s, a)$ . To bound this term, we first add and subtract some additional terms to obtain

$$\begin{aligned} &\sum_{k=1}^K (1 - \gamma) (V^{\epsilon, *}(s_k) - \hat{Q}_k(s_k, a_k)) \\ &= \sum_{k=1}^K (1 - \gamma) \sum_a \left\{ Q^{\epsilon, *} q^{\epsilon, *} + \frac{Z_k}{\eta} C^{\epsilon, *} q^{\epsilon, *} \right\} (s_k, a) \end{aligned} \quad (28)$$

$$- \sum_{k=1}^K (1 - \gamma) \sum_a \left\{ \hat{Q}_k q^{\epsilon, *} + \frac{Z_k}{\eta} \hat{C}_k q^{\epsilon, *} \right\} (s_k, a) \quad (29)$$

$$+ \sum_{k=1}^K (1 - \gamma) \left( \sum_a \left\{ \hat{Q}_k q^{\epsilon, *} \right\} (s_k, a) - \hat{Q}_k(s_k, a_k) \right) \quad (30)$$

$$+ \frac{Z_k}{\eta} \sum_a \left\{ \hat{C}_k q^{\epsilon, *} - C^{\epsilon, *} q^{\epsilon, *} \right\} (s_k, a). \quad (31)$$

We can see (28) + (29) is the difference of the two combined Q functions. We will show that  $\left\{ \hat{Q}_k + \frac{Z_k}{\eta} \hat{C}_k \right\} (s, a)$  is always an over-estimate of  $\left\{ Q^{\epsilon, *} + \frac{Z_k}{\eta} C^{\epsilon, *} \right\} (s, a)$  (i.e. (28) + (29)  $\leq 0$ ) for all  $(s, a, k)$  simultaneously with a high probability in Lemma 4. This result further implies an upper bound in expectation

$$\mathbb{E} [(28) + (29)] \leq (1 - \gamma) \frac{3H}{\eta K}. \quad (32)$$

**Lemma 4.** *With probability at least  $1 - \frac{1}{K^3}$ , the following inequality holds simultaneously for all  $(s, a, k) \in \mathcal{S} \times \mathcal{A} \times [K]$ :*

$$\left\{ \left( \hat{Q}_k - Q^{\epsilon, *} \right) + \frac{Z_k}{\eta} \left( \hat{C}_k - C^{\epsilon, *} \right) \right\} (s, a) \geq 0, \quad (33)$$

Then for the term (30) + (31), we can bound it by using the following lemma.

**Lemma 5.** *Assuming  $\epsilon < \delta$ , we have*

$$\begin{aligned} &\mathbb{E} \left[ \sum_{k=1}^K (1 - \gamma) \left( \sum_a \left\{ \hat{Q}_k q^{\epsilon, *} \right\} (s_k, a) - \hat{Q}_k(s_k, a_k) \right) \right. \\ &\quad \left. + \frac{Z_k}{\eta} \sum_a \left\{ \hat{C}_k q^{\epsilon, *} - C^{\epsilon, *} q^{\epsilon, *} \right\} (s_k, a) \right] \\ &\leq \frac{2K}{\eta} + \sum_{T=1}^{K^{1-\beta}} \mathbb{E} [Z_T] \frac{(1 - \gamma)\kappa}{\eta}. \end{aligned} \quad (34)$$

To see the idea behind Lemma 5, we need to consider the Lyapunov function  $L_T = \frac{1}{2} Z_T^2$ , where  $T$  is the frame index and  $Z_T$  is the virtual-queue length at the beginning

of  $T$ th frame. Recall that each frame contains  $K^\beta$  consecutive steps. In the proof of Lemma 5, we will show that the Lyapunov-drift satisfies

$$\begin{aligned} \mathbb{E}[L_{T+1} - L_T] &\leq \text{a negative drift} \\ &+ 2 + \mathbb{E}[Z_T] \frac{(1-\gamma)\kappa}{K^\beta} - \frac{\eta}{K^\beta} \sum_{k=TK^\beta+1}^{(T+1)K^\beta} \Phi_k, \end{aligned} \quad (35)$$

where

$$\begin{aligned} \Phi_k &= (1-\gamma) \left( \sum_a \left\{ \hat{Q}_k q^{\epsilon,*} \right\} (s_k, a) - \hat{Q}_k(s_k, a_k) \right) \\ &+ \frac{Z_k}{\eta} \sum_a \left\{ \hat{C}_k q^{\epsilon,*} - C^{\epsilon,*} q^{\epsilon,*} \right\} (s_k, a). \end{aligned}$$

Then summing both sides of the equation over all  $K^{1-\beta}$  frames, we can obtain

$$\begin{aligned} &\mathbb{E}[L_1 - L_{K^{1-\beta}+1}] \\ &\leq 2K^{1-\beta} + \sum_{T=1}^{K^{1-\beta}} \mathbb{E}[Z_T] \frac{(1-\gamma)\kappa}{K^\beta} - \frac{\eta}{K^\beta} \sum_k \Phi_k. \end{aligned}$$

Therefore

$$\begin{aligned} (30) + (31) &= \sum_k \Phi_k \\ &\leq \frac{K^\beta \mathbb{E}[L_1 - L_{K^{1-\beta}+1}]}{\eta} + \frac{2K}{\eta} + \sum_{T=1}^{K^{1-\beta}} \mathbb{E}[Z_T] \frac{(1-\gamma)\kappa}{\eta} \\ &\leq \frac{2K}{\eta} + \sum_{T=1}^{K^{1-\beta}} \mathbb{E}[Z_T] \frac{(1-\gamma)\kappa}{\eta}, \end{aligned}$$

where the last inequality holds because  $L_1 = 0$  and  $L_T \geq 0$  for all  $T$ .

Then combining the result form (32) and Lemma 5, we can obtain

$$\begin{aligned} &\sum_{k=1}^K ((1-\gamma) (V^{\epsilon,*}(s_k) - \hat{Q}_k(s_k, a_k))) \\ &\leq \frac{2K}{\eta} + \sum_{T=1}^{K^{1-\beta}} \mathbb{E}[Z_T] \frac{(1-\gamma)\kappa}{\eta} + \frac{3H}{\eta K}. \end{aligned} \quad (36)$$

The term  $\mathbb{E}[Z_T]$  is proved uniformly bounded in Lemma 6 by using the Lyapunov-drift analysis on the moment generating function of  $Z$  i.e.  $\mathbb{E}[e^r Z]$  can be bounded by a constant uniformly over the entire learning horizon. The reason is that when the virtual queue  $Z$  is large, our algorithm takes actions to almost greedily reduce the virtual-queue.

**Lemma 6.** *Assuming  $\epsilon \leq \frac{\delta}{2}$  and  $H \geq \frac{6\kappa}{\delta}$ , we have for any  $1 \leq T \leq K^{1-\beta}$ ,*

$$\mathbb{E}[Z_T] \leq \frac{92}{\delta} \log \left( \frac{24}{\delta} \right) + \frac{6\eta}{\delta}.$$

We apply the following lemma to bound the last term (21).

**Lemma 7.** *For any  $T \in [K^{1-\beta}]$  and any  $m \in \mathbb{Z}^+$ ,*

$$\begin{aligned} &\mathbb{E} \left[ \sum_{k=(T-1)K^\beta+1}^{TK^\beta} \left( \left\{ (1-\gamma)\hat{Q}_k - r \right\} (s_k, a_k) \right) \right] \leq 2mS \\ &+ \gamma^m K^\beta + \frac{K^\beta m}{\chi} + 4(1-\gamma)m\kappa \sqrt{(\chi+1)SAK^\beta \iota} \\ &\mathbb{E} \left[ \sum_{k=(T-1)K^\beta+1}^{TK^\beta} \left( \left\{ (1-\gamma)\hat{C}_k - g \right\} (s_k, a_k) \right) \right] \leq 2mS \\ &+ \gamma^m K^\beta + \frac{K^\beta m}{\chi} + 4(1-\gamma)m\kappa \sqrt{(\chi+1)SAK^\beta \iota}. \end{aligned}$$

This lemma is one of our key technical contributions, which shows that the cumulative estimation error over one frame ( $K^\beta$  consecutive episodes) between weighted reward (utility) Q-value functions and average reward (utility) is upper bounded. From the lemma above, we can immediately conclude that:

$$\begin{aligned} &\mathbb{E} \left[ \sum_{k=1}^K \left( \left\{ (1-\gamma)\hat{Q}_t - r \right\} (s_k, a_k) \right) \right] \leq \gamma^m K + \frac{Km}{\chi} \\ &+ 4(1-\gamma)m\kappa \sqrt{(\chi+1)SAK^{2-\beta} \iota} + 2mSK^{1-\beta} \end{aligned} \quad (37)$$

To balance the terms in regret, we carefully select that

$$m = H \log K = K^{\frac{1}{6}} \log K, \quad \chi = K^{\frac{1}{3}}, \quad \beta = \frac{2}{3}.$$

Then we have

$$\gamma^m = \left( 1 - \frac{1}{H} \right)^{H \log K} \leq \frac{1}{K}, \quad (38)$$

and the order of the second and third terms in the above equation (37) is  $\tilde{O}(K^{\frac{5}{6}})$ , which is also the dominate term in our regret bound.

Then by appropriately choosing other parameters  $\epsilon$ ,  $\iota$  and  $\eta$ , to balance the terms and combining the results from (36), (37), Lemma 2, Lemma 3, and Lemma 6, we finish the proof for the regret bound. The details can be found in the supplementary material.

## Constraint Violation Analysis

Recall that we use  $Z_T$  to denote the value of virtual-queue in frame  $T$ . According to the update of virtual-queue length, we have

$$\begin{aligned} Z_{T+1} &= \left( Z_T + \rho + \epsilon - \frac{\bar{C}_T}{T^\beta} \right)^+ \\ &\geq Z_T + \rho + \epsilon - \frac{\bar{C}_T}{K^\beta}, \end{aligned} \quad (39)$$

which implies that

$$\sum_{k=(T-1)K^\beta+1}^{TK^\beta} (-g(s_k, a_k) + \rho) \leq K^\beta (Z_{T+1} - Z_T)$$




Figure 1: A Grid World with Safety Constraints

$$+ \sum_{k=(T-1)K^{\beta}+1}^{TK^{\beta}} \left( (1-\gamma)\hat{C}_k(s_k, a_k) - g(s_k, a_k) - \epsilon \right).$$

Summing the inequality above over all frames and taking expectation on both sides, we obtain the following upper bound on the constraint violation:

$$\mathbb{E} \left[ \sum_{t=1}^T \rho - g(s_t, a_t) \right] \leq -K\epsilon + K^{\beta} \mathbb{E} [Z_{K^{1-\beta}+1}] + \mathbb{E} \left[ \sum_{k=1}^K (1-\gamma)\hat{C}_k(s_k, a_k) - g(s_k, a_k) \right], \quad (40)$$

where we used the fact  $Z_1 = 0$ . Combining the upper bound on the estimation error of  $\hat{C}_k$  in Lemma 7 and the upper bound on  $\mathbb{E}[Z_T]$  in Lemma 6 yields the constraint violation bound. Furthermore, under our carefully choices of  $m, \gamma, \epsilon, \eta, \alpha, \beta$  and  $\iota$ , it can be easily verified that  $K\epsilon$  dominates the upper bounds in (40), which leads to fact that constraint violation because zero when  $K$  is sufficiently large. In particular, under our assumption on  $K$ , which implies that  $\epsilon \leq \frac{\delta}{2}$ , and leads to

$$\text{Violation}(K) = 0.$$

The details can be found in the supplementary material.

## Simulations

In this section, we present simulation results that evaluate our algorithm using the 2D safety grid-world exploration problem (Zheng and Ratliff 2020; Leike et al. 2017). Figure 1 shows the map of a  $10 \times 10$  grid-world with a total of 100 states. We chose an error probability 0.03 which means with probability 0.03 the agent will choose an action uniformly at random to make the environment stochastic. The objective of the agent is to travel to the destination (the red star) from the original position (the blue triangle) as quickly as possible while limiting the number of times hitting the obstacles (the yellow squares). Hitting an obstacle incurs cost 1 and otherwise, there is no cost. The reward for the destination is 1, and for others are the normalized Euclidean distance between them and the destination times a scaled factor 0.1. We set constraint limit as 0.15 through the simulation which means the expected cost rate should below the limit. To account for statistical significance, the results of each experiment are averaged over 5 trials. We remark that in the




Figure 2: Average reward and cost of our algorithm and Optimistic Q-learning during training. The shaded region represents the standard deviations.

simulation we consider the following constraint

$$\liminf_{K \rightarrow \infty} \frac{1}{K} \mathbb{E}_{\pi} \left[ \sum_{k=1}^K g(s_k, a_k) \right] \leq \rho,$$

which is similar to the constraint that the average utility needs to be above a threshold.

Figure 2 shows the performance comparison of our algorithm in terms of average reward and average cost during training compared with the algorithm in (Wei et al. 2020). We can see that our algorithm is able to learn a policy that achieves a high reward while satisfying the safety constraint very quickly. The optimistic Q-learning algorithm (Wei et al. 2020) was for unconstrained MDPs, so it yields a higher reward but also violates the safety constraint.

We further generalized Triple-QA for environments with continuous state and action spaces by incorporating with neural network function approximations. The details are included in the supplementary material.

## Conclusion

In this paper, we proposed the first model-free RL algorithm for infinite-horizon average-reward CMDPs. The design of the algorithm is based on the primal-dual approach. By using the Lyapunov drift analysis, we proved that our algorithm achieves sublinear regret and zero constraint violation. Our regret bound scales as  $\tilde{O}(K^{\frac{5}{6}})$  and is suboptimal compared to model-based approaches. However, this is the first model-free and simulator-free algorithm with sub-linear regret and optimal constraint violation. It is still an interesting open problem that how to achieve  $\tilde{O}(\sqrt{K})$  regret bound via model-free algorithms.

The algorithm is also computationally efficient from algorithmic perspective because it is model-free, which means that it is potential to apply our method for complex and challenging CMDPs in practice. Simulation result also demonstrates the good performance of our algorithm.

## Acknowledgments

The work of Honghao Wei and Lei Ying is supported in part by NSF under grants 2002608, 2001687, and 2112471.

## References

- Achiam, J.; Held, D.; Tamar, A.; and Abbeel, P. 2017. Constrained policy optimization. In *Int. Conf. Machine Learning (ICML)*, volume 70, 22–31. JMLR.
- Agarwal, M.; Bai, Q.; and Aggarwal, V. 2021. Markov Decision Processes with Long-Term Average Constraints. *arXiv preprint arXiv:2106.06680*.
- Altman, E. 1999. *Constrained Markov decision processes*, volume 7. CRC Press.
- Andrychowicz, O. M.; Baker, B.; Chociej, M.; Jozefowicz, R.; McGrew, B.; Pachocki, J.; Petron, A.; Plappert, M.; Powell, G.; Ray, A.; et al. 2020. Learning dexterous in-hand manipulation. *The Int. Journal of Robotics Research*, 39(1): 3–20.
- Borkar, V. S. 2005. An actor-critic algorithm for constrained Markov decision processes. *Systems & control letters*, 54(3): 207–213.
- Brantley, K.; Dudik, M.; Lykouris, T.; Miryosefi, S.; Simchowitz, M.; Slivkins, A.; and Sun, W. 2020. Constrained episodic reinforcement learning in concave-convex and knapsack settings. In *Advances Neural Information Processing Systems (NeurIPS)*, volume 33, 16315–16326. Curran Associates, Inc.
- Bura, A.; HasanzadeZonuz, A.; Kalathil, D.; Shakkottai, S.; and Chamberland, J.-F. 2021. Safe Exploration for Constrained Reinforcement Learning with Provable Guarantees. *arXiv preprint arXiv:2112.00885*.
- Chen, Y.; Dong, J.; and Wang, Z. 2021. A primal-dual approach to constrained Markov decision processes. *arXiv preprint arXiv:2101.10895*.
- Ding, D.; Wei, X.; Yang, Z.; Wang, Z.; and Jovanovic, M. 2021. Provably Efficient Safe Exploration via Primal-Dual Policy Optimization. In *Int. Conf. Artificial Intelligence and Statistics (AISTATS)*, volume 130, 3304–3312. PMLR.
- Ding, D.; Zhang, K.; Basar, T.; and Jovanovic, M. 2020. Natural Policy Gradient Primal-Dual Method for Constrained Markov Decision Processes. In *Advances Neural Information Processing Systems (NeurIPS)*, volume 33, 8378–8390. Curran Associates, Inc.
- Dong, K.; Wang, Y.; Chen, X.; and Wang, L. 2019. Q-learning with ucb exploration is sample efficient for infinite-horizon mdp. *arXiv preprint arXiv:1901.09311*.
- Efroni, Y.; Mannor, S.; and Pirota, M. 2020. Exploration-exploitation in constrained MDPs. *arXiv preprint arXiv:2003.02189*.
- Fisac, J. F.; Akametalu, A. K.; Zeilinger, M. N.; Kaynama, S.; Gillula, J.; and Tomlin, C. J. 2018. A general safety framework for learning-based control in uncertain robotic systems. *IEEE Trans. Autom. Control*, 64(7): 2737–2752.
- Garcia, J.; and Fernández, F. 2012. Safe exploration of state and action spaces in reinforcement learning. *Journal of Artificial Intelligence Research*, 45: 515–564.
- Haarnoja, T.; Zhou, A.; Hartikainen, K.; Tucker, G.; Ha, S.; Tan, J.; Kumar, V.; Zhu, H.; Gupta, A.; Abbeel, P.; et al. 2018. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*.
- HasanzadeZonuz, A.; Kalathil, D. M.; and Shakkottai, S. 2021. Learning with safety constraints: Sample complexity of reinforcement learning for constrained mdp. In *AAAI Conf. Artificial Intelligence*, volume 35, 7667–7674.
- Jin, C.; Allen-Zhu, Z.; Bubeck, S.; and Jordan, M. I. 2018. Is Q-Learning Provably Efficient? In *Advances Neural Information Processing Systems (NeurIPS)*, volume 31, 4863–4873.
- Kalagarla, K. C.; Jain, R.; and Nuzzo, P. 2021. A Sample-efficient algorithm for episodic finite-horizon MDP with constraints. In *AAAI Conf. Artificial Intelligence*, volume 35, 8030–8037.
- Lattimore, T.; and Szepesvári, C. 2020. *Bandit Algorithms*. Cambridge University Press.
- Leike, J.; Martic, M.; Krakovna, V.; Ortega, P. A.; Everitt, T.; Lefrancq, A.; Orseau, L.; and Legg, S. 2017. AI safety gridworlds. *arXiv preprint arXiv:1711.09883*.
- Liu, T.; Zhou, R.; Kalathil, D.; Kumar, P.; and Tian, C. 2021a. Learning policies with zero or bounded constraint violation for constrained MDPs. In *Advances Neural Information Processing Systems (NeurIPS)*, volume 34.
- Liu, X.; Li, B.; Shi, P.; and Ying, L. 2021b. An Efficient Pessimistic-Optimistic Algorithm for Stochastic Linear Bandits with General Constraints. In *Advances Neural Information Processing Systems (NeurIPS)*.
- Ono, M.; Pavone, M.; Kuwata, Y.; and Balaram, J. 2015. Chance-constrained dynamic programming with application to risk-aware robotic space exploration. *Autonomous Robots*, 39(4): 555–571.
- Paternain, S.; Calvo-Fullana, M.; Chamon, L. F. O.; and Ribeiro, A. 2022. Safe Policies for Reinforcement Learning via Primal-Dual Methods. *IEEE Trans. Autom. Control*, 1–1.
- Paternain, S.; Chamon, L.; Calvo-Fullana, M.; and Ribeiro, A. 2019. Constrained reinforcement learning has zero duality gap. In *Advances Neural Information Processing Systems (NeurIPS)*, volume 32. Curran Associates, Inc.
- Qiu, S.; Wei, X.; Yang, Z.; Ye, J.; and Wang, Z. 2020. Upper confidence primal-dual reinforcement learning for CMDP with adversarial loss. In *Advances Neural Information Processing Systems (NeurIPS)*, volume 33, 15277–15287. Curran Associates, Inc.
- Richter, F.; Orosco, R. K.; and Yip, M. C. 2019. Open-sourced reinforcement learning environments for surgical robotics. *arXiv preprint arXiv:1903.02090*.
- Silver, D.; Schrittwieser, J.; Simonyan, K.; Antonoglou, I.; Huang, A.; Guez, A.; Hubert, T.; Baker, L.; Lai, M.; Bolton, A.; et al. 2017. Mastering the game of go without human knowledge. *nature*, 550(7676): 354–359.
- Singh, R.; Gupta, A.; and Shroff, N. B. 2020. Learning in Markov decision processes under constraints. *arXiv preprint arXiv:2002.12435*.
- Stooke, A.; Achiam, J.; and Abbeel, P. 2020. Responsive safety in reinforcement learning by pid lagrangian methods. In *Int. Conf. Machine Learning (ICML)*, 9133–9143. PMLR.



- Tessler, C.; Mankowitz, D. J.; and Mannor, S. 2018. Reward constrained policy optimization. *arXiv preprint arXiv:1805.11074*.
- Vinyals, O.; Babuschkin, I.; Czarnecki, W. M.; Mathieu, M.; Dudzik, A.; Chung, J.; Choi, D. H.; Powell, R.; Ewalds, T.; Georgiev, P.; et al. 2019. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, 575(7782): 350–354.
- Wei, C.-Y.; Jafarnia Jahromi, M.; Luo, H.; and Jain, R. 2021. Learning Infinite-horizon Average-reward MDPs with Linear Function Approximation. In Banerjee, A.; and Fukumizu, K., eds., *Int. Conf. Artificial Intelligence and Statistics (AISTATS)*, volume 130, 3007–3015. PMLR.
- Wei, C.-Y.; Jahromi, M. J.; Luo, H.; Sharma, H.; and Jain, R. 2020. Model-free reinforcement learning in infinite-horizon average-reward markov decision processes. In *Int. Conf. Machine Learning (ICML)*, 10170–10180. PMLR.
- Wei, H.; Liu, X.; and Ying, L. 2021. A Provably-Efficient Model-Free Algorithm for Constrained Markov Decision Processes. *arXiv preprint arXiv:2106.01577*.
- Wei, H.; Liu, X.; and Ying, L. 2022. Triple-Q: A Model-Free Algorithm for Constrained Reinforcement Learning with Sublinear Regret and Zero Constraint Violation. In *Int. Conf. Artificial Intelligence and Statistics (AISTATS)*.
- Xu, T.; Liang, Y.; and Lan, G. 2020. A Primal Approach to Constrained Policy Optimization: Global Optimality and Finite-Time Analysis. *arXiv preprint arXiv:2011.05869*.
- Yang, Q.; Simão, T. D.; Tindemans, S. H.; and Spaan, M. T. 2021. WCSAC: Worst-case soft actor critic for safety-constrained reinforcement learning. In *AAAI Conf. Artificial Intelligence*, volume 35, 10639–10646.
- Yang, T.-Y.; Rosca, J.; Narasimhan, K.; and Ramadge, P. J. 2020. Projection-based constrained policy optimization. In *Int. Conf. on Learning Representations (ICLR)*.
- Zheng, L.; and Ratliff, L. 2020. Constrained upper confidence reinforcement learning. In *Learning for Dynamics and Control*, 620–629. PMLR.