

Adaptive Primal-Dual Method for Safe Reinforcement Learning

Wei Qin Chen
Rensselaer Polytechnic Institute
Troy, NY, USA
chenw18@rpi.edu

James Onyejizu
Rensselaer Polytechnic Institute
Troy, NY, USA
onyejj@rpi.edu

Long Vu
IBM T.J. Watson Research Center
Yorktown Heights, NY, USA
lhvu@us.ibm.com

Lan Hoang
IBM Research
Daresbury, Warrington, UK
lan.hoang@ibm.com

Dharmashankar Subramanian
IBM T.J. Watson Research Center
Yorktown Heights, NY, USA
dharmash@us.ibm.com

Koushik Kar
Rensselaer Polytechnic Institute
Troy, NY, USA
kark@rpi.edu

Sandipan Mishra
Rensselaer Polytechnic Institute
Troy, NY, USA
mishrs2@rpi.edu

Santiago Paternain
Rensselaer Polytechnic Institute
Troy, NY, USA
paters@rpi.edu

ABSTRACT

Primal-dual methods have a natural application in Safe Reinforcement Learning (SRL), posed as a constrained policy optimization problem. In practice however, applying primal-dual methods to SRL is challenging, due to the inter-dependency of the learning rate (LR) and Lagrangian multipliers (dual variables) each time an embedded unconstrained RL problem is solved. In this paper, we propose, analyze and evaluate adaptive primal-dual (APD) methods for SRL, where two adaptive LRs are adjusted to the Lagrangian multipliers so as to optimize the policy in each iteration. We theoretically establish the convergence, optimality and feasibility of the APD algorithm. Finally, we conduct numerical evaluation of the practical APD algorithm with four well-known environments in Bullet-Safey-Gym employing two state-of-the-art SRL algorithms: PPO-Lagrangian and DDPG-Lagrangian. All experiments show that the practical APD algorithm outperforms (or achieves comparable performance) and attains more stable training than the constant LR cases. Additionally, we substantiate the robustness of selecting the two adaptive LRs by empirical evidence.

KEYWORDS

Safe Reinforcement Learning; Adaptive Primal-Dual; Adaptive Learning Rates

ACM Reference Format:

Wei Qin Chen, James Onyejizu, Long Vu, Lan Hoang, Dharmashankar Subramanian, Koushik Kar, Sandipan Mishra, and Santiago Paternain. 2024. Adaptive Primal-Dual Method for Safe Reinforcement Learning. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), Auckland, New Zealand, May 6 – 10, 2024*, IFAAMAS, 9 pages.



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand. © 2024 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

1 INTRODUCTION

Reinforcement learning (RL) has a rich history of solving a wide range of decision-making problems. Recently, RL has succeeded in training large language models such as ChatGPT [1], playing video games at superhuman level [2–4], mastering Go [5, 6], and manipulating robotics [7, 8]. RL problems, in general, are formulated as Markov Decision Processes (MDPs). In this work, we are interested in conditions where the underlying dynamics are unknown, the optimal policy thus needs to be learned from data (samples). The goal for an agent is to explore the environment so that it is able to maximize the expected cumulative reward.

Nevertheless, classical RL techniques might lead to risky/unsafe actions [9–11], if they are only concerned with the reward. Therefore, safety constitutes a foundational aspect in realistic domains or physical entities. Specifically, in the realm of robot navigation, ensuring collision avoidance [12, 13] is essential for their proper functioning, and to ensure the preservation of human safety in the vicinity. Taking into account the safety requirements motivates the development of policy optimization under safety guarantees [14–16].

A common approach is to employ the framework of Constrained MDPs (CMDPs) [17] where auxiliary costs (analogous to reward) are considered in the constraints. This framework has gained widespread adoption for inducing safe behaviors [16, 18–22]. We briefly introduce these work below.

1.1 Related Work

The state-of-the-art algorithms for solving CMDPs commonly include two types of methods: primal methods and primal-dual methods.

1.1.1 Primal Methods. [18] develops a constrained policy optimization (CPO) algorithm for SRL that searches the feasible policy within the confines of the trust region while guaranteeing a monotonic performance improvement as well as constraint satisfaction. Projection-based constrained policy optimization (PCPO) [19] employs TRPO [23], the cutting-edge unconstrained RL algorithm first, and then projects the policy back into the feasible set. Nevertheless,

both CPO and PCPO suffer from the approximation error and expensive computation of the inversion of high-dimensional Hessian matrices. To address these issues, [20] proposes first-order constrained optimization in policy space (FOCOPS), which optimizes the constrained policy in the non-parametric space, and subsequently derives the first-order gradients of the ℓ_2 loss function in the parameterization space. Another alternative is the penalized proximal policy optimization (P3O) [24], which solves the constrained policy iteration based on the minimization of an equivalent unconstrained problem. However, both FOCOPS and P3O introduce more auxiliary parameters that need to be learned.

1.1.2 Primal-Dual Methods. Lagrangian-based primal-dual algorithms like primal-dual optimization (PDO) [16] and reward constrained policy optimization (RCPO) [21] have succeeded in solving CMDP optimization problems. Nevertheless, the convergence guarantees of these algorithms are limited to either local (locally optimal policies or stationary-point) [16, 21, 25] or asymptotic scenarios [26]. [22] establishes the analysis on the duality gap for CMDPs within the policy space and provides a provably dual descent algorithm under the assumption of access to a non-convex optimization oracle. However, obtaining the solution to a primal non-convex problem remains an issue, thus lacking global convergence guarantees. To tackle these challenges, [27] proposes a natural policy gradient primal-dual (NPG-PD) method that achieves non-asymptotic global convergence with sublinear rates regarding both the optimality gap and the constraint violation. [28] provides the first provably efficient online policy optimization algorithm—optimistic primal-dual proximal policy optimization (OPDOP) and establishes the bounds on the regret and constraint violation. [29] proposes a safe primal-dual (SPG) algorithm that is used to solve a CMDP problem with probabilistic safety constraints. Nonetheless, achieving satisfactory performance with the existing primal-dual methods remains challenging, mainly due to their sensitivity to hyper-parameters such as learning rates.

1.2 Main Contribution

This paper addresses a core challenge in applying primal-dual methods to SRL problems, the inter-dependence of the primal *Learning Rate (LR)* and *Lagrangian Multiplier (LM)* (dual variable) parameters in the primal-dual method. We provide analytical expressions (bounds) of the amount of progress made in each step of the primal-dual algorithm, based on which we develop two adaptive LR choices that optimize these bounds. The two LR choices have an inverse-linear and inverse-quadratic dependence on the LM, and are incorporated into the proposed adaptive primal-dual (APD) algorithm. We provide theoretical analyses of the convergence, return optimality, and constraint feasibility of the APD algorithm. Finally, we numerically evaluate the practical version of APD (PAPD) algorithm using four environments in the Bullet-Safety-Gym[30], and compare the performance of PAPD with constant-LR solutions using two state-of-the-art constrained RL methods: PPO-Lagrangian (PPOL) and DDPG-Lagrangian (DDPGL). We also numerically demonstrate the robustness of PAPD algorithm with respect to certain key parameter choices.

2 SAFE REINFORCEMENT LEARNING

MDPs are defined by a tuple $(S, A, R, \mathbb{P}, \mathcal{U}, \gamma)$ [31], where S is the state space, A is the action space, $R : S \times A \times S \rightarrow \mathbb{R}$ is the reward function describing the quality of the decision. For any $\hat{S} \subset S$, $s_t \in S$, $a_t \in A$, $t \in \{0, 1, \dots\}$, $\mathbb{P}(s_{t+1} \in \hat{S} | s_t, a_t)$ (i.e., the probability of s_{t+1} being in \hat{S} given s_t and a_t) is the transition probability describing the dynamics of the system, $\mathcal{U}(\hat{S}) := \mathbb{P}(s \in \hat{S})$ is the starting state distribution, and γ is the discount factor. Note that a table of notations is provided in the supplementary material of [32], which serves to facilitate tracking and enhance comprehension for the readers.

Consider a parameterization space Θ and a probability density function $P(\cdot)$. Given $\theta \in \Theta$, a parameterized stationary *policy* $\pi_\theta : S \rightarrow P(A)$ maps states to probability distributions over the set of actions, and $\pi_\theta(a|s)$ indicates the probability density of drawing action $a \in A$ in the corresponding state $s \in S$. Common parameterizations include neural networks (NN) and Radial Basis Functions (RBFs). In this work, we are particularly interested in situations where the state transition distributions are unknown, and thus the policies need to be computed through interacting with the environment.

In the context of MDPs, the objective is to find an optimal parameter that maximizes the expected discounted return

$$J_R(\pi_\theta) = \mathbb{E}_{\tau \sim \pi_\theta} \left[\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t, s_{t+1}) \right], \quad (1)$$

where $\tau = (s_0, a_0, s_1, a_1, \dots)$ denotes a sample trajectory. Given fixed initial state distribution $s_0 \sim \mathcal{U}$ and transition distribution $s_{t+1} \sim \mathbb{P}(\cdot | s_t, a_t)$, let us define a shorthand $\tau \sim \pi_\theta$ indicating that the distribution over trajectories depends on the policy through $a_t \sim \pi_\theta(\cdot | s_t)$.

CMDPs [17] impose additional constraints on the allowable policies. More concretely, auxiliary cost functions $C_i : S \times A \times S \rightarrow \mathbb{R}$, $i = 1, 2, \dots, m$ are introduced. We make the following assumption about the costs.

ASSUMPTION 1. Consider $B > 0$ and $C = [C_1, \dots, C_m]^T$, assume we have bounded costs such that $\|C\| \leq B$.

Analogous to the expected discounted return, define the expected discounted cost as

$$J_{C_i}(\pi_\theta) = \mathbb{E}_{\tau \sim \pi_\theta} \left[\sum_{t=0}^{\infty} \gamma^t C_i(s_t, a_t, s_{t+1}) \right]. \quad (2)$$

Denote by d_i , $i = 1, 2, \dots, m$ the desired threshold for the expected discounted cost. Accordingly, SRL problems can be formulated as

$$\begin{aligned} \theta^* \in \arg \max_{\theta \in \Theta} J_R(\pi_\theta) \\ \text{s.t. } J_{C_i}(\pi_\theta) \leq d_i, \quad i = 1, 2, \dots, m. \end{aligned} \quad (3)$$

The SRL problem posed in (3) can be solved by applying gradient-based methods on a regularized objective function [33] or using primal-dual methods [16, 29] to achieve local optimal solutions. The performance of the first type of method can heavily depend on the choice of the regularization parameter, and this choice is problem dependent for which there is no easy method. In general, for a given choice of regularization parameter, the first type of method may




Figure 1: Learning curves of PPOL over five independent runs with fixed LM values of 1 and 5. The horizontal axis represents time steps. Cost limit $\mathbf{d} = 10$ (black dashed line) in all experiments. LR = 0.0006 outperforms LR = 0.0003 at LM = 1 (red curve is infeasible), while the opposite holds when LM = 5.

not produce optimal results. In this work, therefore, we focus on the primal-dual method, where the regularization parameter is the LM (dual variable) that is progressively updated. We introduce this method in the next section and motivate the need for an adaptive LR in the primal update step.

3 ADAPTIVE PRIMAL-DUAL ALGORITHM

3.1 Motivation

Primal-dual methods rely on iterative training of the policy parameters θ to minimize the Lagrangian of (3):

$$\mathcal{L}(\theta, \lambda) \doteq -J_R(\pi_\theta) + \lambda^T (J_C(\pi_\theta) - \mathbf{d}), \quad (4)$$

where $J_C(\cdot) = [J_{C_1}(\cdot), J_{C_2}(\cdot), \dots, J_{C_m}(\cdot)]^T$, $\mathbf{d} = [d_1, d_2, \dots, d_m]^T$, and $\lambda \in \mathbb{R}^m$ is the LM.

As stated in Section 1, despite the impressive capabilities demonstrated by state-of-the-art algorithms like PDO [16] and RCPO [21] in addressing a wide range of SRL problems, these primal-dual like algorithms still encounter challenges of selecting an appropriate LR for the embedded RL problem during the training process. Indeed, it is natural to posit that employing a constant LR throughout the training process might not be optimal, given that the LM undergoes continuous changes. We illustrate this through a numerical example in Figure 1. Figure 1 presents learning curves (both Return and Cost) for the PPOL method over five independent runs using fixed LM values of 1 and 5. We use the *SafetyCarRun-v0* environment from the *Bullet-Safety-Gym* [30] for this test. Notably, as LM transitions from 1 to 5, the LR needs adjustment from 0.0006 to 0.0003 for optimal results (in terms of return and cost). This means, maintaining a consistent LR across varying LMs might lead to sub-optimal results. Specifically, deploying an optimized LR tailored for an LM value of 5 (green curve) in a scenario meant for LM value of 1 (red curve) and vice versa leads to compromised performance. Therefore, accounting for the dependence between LM and LR is crucial for achieving good performance in SRL.

3.2 Adaptive Learning Rate

A naive approach to addressing the above challenge is to solve the primal problem in (4) for different values of λ (LM), each with its optimized LR. However, this approach does not scale when the LM is multi-dimensional (i.e., there are multiple safety constraints) or has a large range. Primal-dual algorithms offer a practical way to find the LM (dual variable) λ^* that maximizes the dual function

$$d(\lambda) = \min_{\theta \in \mathbb{R}^d} -J_R(\pi_\theta) + \lambda^T (J_C(\pi_\theta) - \mathbf{d}). \quad (5)$$

Consequently the optimal policy (primal variable) θ^* is computed through iterative coordinated updates of the primal and dual variables

$$\begin{cases} \theta_{k+1} = \theta_k - \eta_k \nabla_\theta \mathcal{L}(\theta_k, \lambda_k), \\ \lambda_{k+1} = [\lambda_k + \zeta g(\theta_{k+1})]_+, \end{cases} \quad (6)$$

$$\lambda_{k+1} = [\lambda_k + \zeta g(\theta_{k+1})]_+, \quad (7)$$

where $g(\cdot)$ is the constraint function, i.e.,

$$g(\theta_{k+1}) \doteq J_C(\pi_{\theta_{k+1}}) - \mathbf{d}, \quad (8)$$

and η_k, ζ denote the primal and dual LR, respectively. Note the dependence of the primal LR η_k on iteration k in (6); this is because in our adaptive primal-dual algorithm, described next, η_k depends on the LM λ_k .

We have reached the phase where we articulate the convergence of (7). This is formally established in the following theorem.

THEOREM 1. Consider the dual function $d(\cdot)$ defined in (5), the constraint function $g(\cdot)$ and cost limit \mathbf{d} in (8). Let $\lambda^* \in \arg \max_{\lambda \in \mathbb{R}_+^m} d(\lambda)$ and define $D^* = d(\lambda^*)$. Let θ_k and λ_k be the sequences generated by (6) and (7). Denote by ϵ_k the primal error of updating the Lagrangian given λ_k , i.e., $\epsilon_k = \mathcal{L}(\theta_{k+1}, \lambda_k) - \min_\theta \mathcal{L}(\theta, \lambda_k)$. Define $\lambda_{best} = \arg \max_{\lambda \in \{\lambda_k\}_{k=0}^K} d(\lambda)$, it holds that

$$0 \leq D^* - d(\lambda_{best}) \leq \frac{\|\lambda_0 - \lambda^*\|^2}{2\zeta K} + \frac{\zeta(B + (1-\gamma)\|\mathbf{d}\|)^2}{2(1-\gamma)^2} + \frac{1}{K} \sum_{k=0}^{K-1} \epsilon_k. \quad (9)$$

Proof. The proof follows the standard stochastic gradient descent analysis. See the supplementary material of [32] for details. \square

Note that the first two terms on the right-hand side of (9) depend exclusively on the problem and on the dual LR. The last term, however, depends on the primal error ϵ_k . Our goal is to minimize this error through an adaptive selection of the primal LR η_k . We resort to two different analyses to bound these errors (see Lemma 1). From these, we derive two adaptive LR. Before stating Lemma 1, we require the following assumptions.

ASSUMPTION 2. Assume the gradients of the objective and cost functions in (3) are Lipschitz continuous with constants $L_R, L_{C_1}, \dots, L_{C_m}$. Let $L_C = [L_{C_1}, \dots, L_{C_m}]^T$. This implies that the gradient of Lagrangian (4) is Lipschitz on θ

$$\|\nabla_\theta \mathcal{L}(\theta_1, \lambda) - \nabla_\theta \mathcal{L}(\theta_2, \lambda)\| \leq L(\lambda) \|\theta_1 - \theta_2\|. \quad (10)$$

where $L(\lambda) = L_R + \lambda^T L_C$.

ASSUMPTION 3. Assume the Lagrangian function $\mathcal{L}(\theta, \lambda)$ is μ -strongly convex on θ . This is, $\mathcal{L}(\theta_1, \lambda_k) \geq \mathcal{L}(\theta_2, \lambda_k) + \nabla_\theta \mathcal{L}(\theta_2, \lambda_k)^T (\theta_1 - \theta_2) + \mu/2 \|\theta_1 - \theta_2\|^2$.

LEMMA 1. Assume the Lagrangian function $\mathcal{L}(\theta_k, \lambda_k)$ is L' -Lipschitz continuous. Define $\delta_k = \|\theta_k - \theta_{k+1}^*\|^2$. Then, the following bounds hold for ϵ_k given the assumptions that the term inside the square root is non-negative.

$$\epsilon_k \leq L' \sqrt{\frac{2}{\mu} \left(L(\lambda) \delta_k + \left(L(\lambda) (\eta_k^1)^2 - \eta_k^1 \right) \|\nabla_{\theta} \mathcal{L}(\theta_k, \lambda_k)\|^2 \right)}, \quad (11)$$

$$\epsilon_k \leq L' \sqrt{\left(1 + (\eta_k^2)^2 L(\lambda)^2 - \eta_k^2 \mu \right) \delta_k}. \quad (12)$$

Proof. We start the proof by employing the L' -Lipschitz continuity of the Lagrangian to bound ϵ_k

$$\epsilon_k = \mathcal{L}(\theta_{k+1}, \lambda_k) - \mathcal{L}(\theta_{k+1}^*, \lambda_k) \leq L' \|\theta_{k+1} - \theta_{k+1}^*\|. \quad (13)$$

Hence, to bound ϵ_k , one can instead investigate $\|\theta_{k+1} - \theta_{k+1}^*\|$.

We first focus on (11). It follows from the strong convexity of the Lagrangian with respect to θ that

$$\begin{aligned} \mathcal{L}(\theta_{k+1}, \lambda_k) &\geq \mathcal{L}(\theta_{k+1}^*, \lambda_k) + \nabla_{\theta} \mathcal{L}(\theta_{k+1}^*, \lambda_k)^T (\theta_{k+1} - \theta_{k+1}^*) \\ &\quad + \frac{\mu}{2} \|\theta_{k+1} - \theta_{k+1}^*\|^2. \end{aligned} \quad (14)$$

Since $\nabla_{\theta} \mathcal{L}(\theta_{k+1}^*, \lambda_k) = 0$ the previous inequality reduces to

$$\|\theta_{k+1} - \theta_{k+1}^*\|^2 \leq \frac{2}{\mu} \left(\mathcal{L}(\theta_{k+1}, \lambda_k) - \mathcal{L}(\theta_{k+1}^*, \lambda_k) \right). \quad (15)$$

We apply the Taylor expansion on both terms on the right-hand side of the previous expression around θ_k

$$\mathcal{L}(\theta_{k+1}, \lambda_k) = \mathcal{L}(\theta_k, \lambda_k) + \nabla_{\theta} \mathcal{L}(\theta_c, \lambda_k)^T (\theta_{k+1} - \theta_k), \quad (16)$$

where $\theta_c = \xi \theta_k + (1 - \xi) \theta_{k+1}$, $\xi \in [0, 1]$, and

$$\mathcal{L}(\theta_{k+1}^*, \lambda_k) = \mathcal{L}(\theta_k, \lambda_k) + \nabla_{\theta} \mathcal{L}(\theta'_c, \lambda_k)^T (\theta_{k+1}^* - \theta_k), \quad (17)$$

where $\theta'_c = \xi' \theta_k + (1 - \xi') \theta_{k+1}^*$, $\xi' \in [0, 1]$. Substituting (16) and (17) into (15) reduces to

$$\begin{aligned} \|\theta_{k+1} - \theta_{k+1}^*\|^2 &\leq \frac{2}{\mu} \left(\nabla_{\theta} \mathcal{L}(\theta'_c, \lambda_k)^T (\theta_k - \theta_{k+1}^*) \right. \\ &\quad \left. + \nabla_{\theta} \mathcal{L}(\theta_c, \lambda_k)^T (\theta_{k+1} - \theta_k) \right). \end{aligned} \quad (18)$$

By virtue of $\nabla_{\theta} \mathcal{L}(\theta_{k+1}^*, \lambda_k) = 0$ (18) reduces to

$$\begin{aligned} \|\theta_{k+1} - \theta_{k+1}^*\|^2 &\leq \frac{2}{\mu} \left(\left(\nabla_{\theta} \mathcal{L}(\theta'_c, \lambda_k) - \nabla_{\theta} \mathcal{L}(\theta_{k+1}^*, \lambda_k) \right)^T \right. \\ &\quad \left. (\theta_k - \theta_{k+1}^*) + \nabla_{\theta} \mathcal{L}(\theta_c, \lambda_k)^T (\theta_{k+1} - \theta_k) \right). \end{aligned} \quad (19)$$

where the last inequality follows from the Lipschitz continuity of the gradient of Lagrangian (see Assumption 2).

By adding and subtracting $\nabla_{\theta} \mathcal{L}(\theta_k, \lambda_k)^T (\theta_{k+1} - \theta_k)$ we obtain

$$\begin{aligned} \|\theta_{k+1} - \theta_{k+1}^*\|^2 &\leq \frac{2}{\mu} \left(L(\lambda) \|\theta_k - \theta_{k+1}^*\|^2 + \right. \\ &\quad \left(\nabla_{\theta} \mathcal{L}(\theta_c, \lambda_k) - \nabla_{\theta} \mathcal{L}(\theta_k, \lambda_k) \right)^T (\theta_{k+1} - \theta_k) \\ &\quad \left. + \nabla_{\theta} \mathcal{L}(\theta_k, \lambda_k)^T (\theta_{k+1} - \theta_k) \right). \end{aligned} \quad (20)$$

Employing Lipschitz continuity of the gradient of Lagrangian yields

$$\begin{aligned} \|\theta_{k+1} - \theta_{k+1}^*\|^2 &\leq \frac{2}{\mu} \left(L(\lambda) \|\theta_k - \theta_{k+1}^*\|^2 + L(\lambda) \|\theta_{k+1} - \theta_k\|^2 \right. \\ &\quad \left. + \nabla_{\theta} \mathcal{L}(\theta_k, \lambda_k)^T (\theta_{k+1} - \theta_k) \right). \end{aligned} \quad (21)$$

By definition of δ_k the previous inequality can be rewritten as

$$\begin{aligned} \|\theta_{k+1} - \theta_{k+1}^*\|^2 &\leq \frac{2}{\mu} \left(L(\lambda) \delta_k + L(\lambda) (\eta_k^1)^2 \|\nabla_{\theta} \mathcal{L}(\theta_k, \lambda_k)\|^2 \right. \\ &\quad \left. - \eta_k^1 \|\nabla_{\theta} \mathcal{L}(\theta_k, \lambda_k)\|^2 \right). \end{aligned} \quad (22)$$

Taking the square root on the previous inequality and combining with (13) completes the proof of (11).

We then turn our attention to proving (12). Note that

$$\begin{aligned} \|\theta_{k+1} - \theta_{k+1}^*\|^2 &= \|\theta_k - \eta_k^2 \nabla_{\theta} \mathcal{L}(\theta_k, \lambda_k) - \theta_{k+1}^*\|^2 \\ &= \|\theta_k - \theta_{k+1}^*\|^2 + (\eta_k^2)^2 \|\nabla_{\theta} \mathcal{L}(\theta_k, \lambda_k)\|^2 \\ &\quad - 2\eta_k^2 \nabla_{\theta} \mathcal{L}(\theta_k, \lambda_k)^T (\theta_k - \theta_{k+1}^*). \end{aligned} \quad (23)$$

By $\nabla_{\theta} \mathcal{L}(\theta_{k+1}^*, \lambda_k) = 0$ the previous equation is equivalent to

$$\begin{aligned} \|\theta_{k+1} - \theta_{k+1}^*\|^2 &= \|\theta_k - \theta_{k+1}^*\|^2 \\ &\quad + (\eta_k^2)^2 \left\| \nabla_{\theta} \mathcal{L}(\theta_k, \lambda_k) - \nabla_{\theta} \mathcal{L}(\theta_{k+1}^*, \lambda_k) \right\|^2 \\ &\quad - 2\eta_k^2 \nabla_{\theta} \mathcal{L}(\theta_k, \lambda_k)^T (\theta_k - \theta_{k+1}^*). \end{aligned} \quad (24)$$

Assumption 2 indicates that the gradient of Lagrangian function is $L(\lambda)$ -Lipschitz continuous, and we thus obtain

$$\begin{aligned} \|\theta_{k+1} - \theta_{k+1}^*\|^2 &\leq \|\theta_k - \theta_{k+1}^*\|^2 \\ &\quad + (\eta_k^2)^2 L(\lambda)^2 \|\theta_k - \theta_{k+1}^*\|^2 \\ &\quad - 2\eta_k^2 \nabla_{\theta} \mathcal{L}(\theta_k, \lambda_k)^T (\theta_k - \theta_{k+1}^*). \end{aligned} \quad (25)$$

In addition, the strong convexity of the Lagrangian with respect to θ shows that

$$\begin{aligned} \nabla_{\theta} \mathcal{L}(\theta_k, \lambda_k)^T (\theta_{k+1}^* - \theta_k) &\leq \mathcal{L}(\theta_{k+1}^*, \lambda_k) - \mathcal{L}(\theta_k, \lambda_k) \\ &\quad - \frac{\mu}{2} \|\theta_{k+1}^* - \theta_k\|^2 \\ &\leq -\frac{\mu}{2} \|\theta_{k+1}^* - \theta_k\|^2, \end{aligned} \quad (26)$$

where the last inequality follows from the fact that θ_{k+1}^* is the minimizer of the Lagrangian. Subsequently, combining (25) and (26) yields

$$\begin{aligned} \|\theta_{k+1} - \theta_{k+1}^*\|^2 &\leq \|\theta_k - \theta_{k+1}^*\|^2 - \eta_k^2 \mu \|\theta_k - \theta_{k+1}^*\|^2 \\ &\quad + (\eta_k^2)^2 L(\lambda)^2 \|\theta_k - \theta_{k+1}^*\|^2 \\ &= \left(1 + (\eta_k^2)^2 L(\lambda)^2 - \eta_k^2 \mu \right) \|\theta_k - \theta_{k+1}^*\|^2. \end{aligned} \quad (27)$$

Taking the square root and combining with (13) completes the proof of the result. \square

By virtue of Assumptions 2, 3 and Lemma 1, we are able to derive two optimal LR's that minimize the two bounds in (11) and (12). We formalize this claim in the following theorem.

THEOREM 2. *Let Assumption 2 and Assumption 3 hold. Then, the learning rates $\eta_k^1 = \frac{1}{2L(\lambda)}$ and $\eta_k^2 = \frac{\mu}{2L(\lambda)^2}$ minimize the bounds in (11) and (12), respectively. Denote by ϵ_k^1 and ϵ_k^2 the bounds on the primal error using η_k^1 and η_k^2 . Consider the L' -Lipschitz continuity of the Lagrangian $\mathcal{L}(\theta_k, \lambda_k)$ and δ_k in Lemma 1. Then, it holds that*

$$\epsilon_k^1 \leq L' \sqrt{\frac{2\delta_k}{\mu} \left(L(\lambda) - \frac{\mu^2}{16L(\lambda)} \right)}, \quad (28)$$

$$\epsilon_k^2 \leq L' \sqrt{\delta_k \left(1 - \frac{\mu^2}{4L(\lambda)^2} \right)}. \quad (29)$$

Proof. We start by establishing two LR η_k^1 and η_k^2 . As observed in (11), the right-hand side is convex with respect to η_k^1 since $L', L(\lambda), \delta_k$ are independent of η_k^1 . To minimize it one can compute the derivative and make it zero. Then the minimizer of the right-hand side of (11) is given by

$$\eta_k^1 = \frac{1}{2L(\lambda)}. \quad (30)$$

Likewise, the right-hand side of (12) is convex with respect to η_k^2 , and the optimal LR (minimizer) takes the form of

$$\eta_k^2 = \frac{\mu}{2L(\lambda)^2}. \quad (31)$$

Having established η_k^1 and η_k^2 , we are now in the stage of proving ϵ_k^1 and ϵ_k^2 . For ϵ_k^1 , substituting η_k^1 into (11) yields the tightest upper bound, i.e.,

$$\epsilon_k^1 \leq L' \sqrt{\frac{2}{\mu} \left(L(\lambda)\delta_k - \frac{\|\nabla_{\theta} \mathcal{L}(\theta_k, \lambda_k)\|^2}{4L(\lambda)} \right)}. \quad (32)$$

Strong convexity indicates that

$$-\nabla_{\theta} \mathcal{L}(\theta_k, \lambda_k)^T (\theta_{k+1}^* - \theta_k) \geq \frac{\mu}{2} \|\theta_{k+1}^* - \theta_k\|^2. \quad (33)$$

Squaring both sides of the previous inequality, using the Cauchy-Schwartz inequality and the definition of δ_k it follows that

$$\|\nabla_{\theta} \mathcal{L}(\theta_k, \lambda_k)\|^2 \geq \frac{\mu^2}{4} \|\theta_{k+1}^* - \theta_k\|^2 = \frac{\mu^2}{4} \delta_k. \quad (34)$$

Combining the previous inequality with (32) yields

$$\epsilon_k^1 \leq L' \sqrt{\frac{2\delta_k}{\mu} \left(L(\lambda) - \frac{\mu^2}{16L(\lambda)} \right)}. \quad (35)$$

Our focus now shifts towards the derivation of ϵ_k^2 . Substituting η_k^2 into (12) yields the tightest upper bound on ϵ_k^2

$$\epsilon_k^2 \leq L' \sqrt{\delta_k \left(1 - \frac{\mu^2}{4L(\lambda)^2} \right)}. \quad (36)$$

These complete the proof of Theorem 2. \square

Notice that both LR proposed by Theorem 2 demonstrate an inverse relationship with respect to the LM, which is also empirically validated by our preliminary observations in Figure 1. Moreover, η_k^1 has an inverse-linear dependence on LM while η_k^2 has an inverse-quadratic dependence on LM. We therefore term them *InvLin* and *InvQua*, respectively. It is important to point out that we assume that the Lagrangian is strongly convex. This is generally not the case for RL problems. However, one can assume local convexity

Algorithm 1 Adaptive Primal-Dual (APD)

- 1: **Input:** $\theta_0, \lambda_0, \zeta, L_R, L_C, \mu$ (optional)
 - 2: **for** $k = 0, 1, \dots$, **do**
 - 3: Choose primal LR from (30) or (31)
 - 4: Update primal variable (policy parameter) as in (6)
 - 5: Update dual variable (LM) as in (7)
 - 6: **end for**
-

Algorithm 2 Practical Adaptive Primal-Dual (PAPD)

- 1: **Input:** $\theta_0, H_1(H'_1), H_2(H'_2), K_P, K_I, K_D$
 - 2: **for** $k = 0, 1, \dots$, **do**
 - 3: Choose primal LR from (39)
 - 4: Update primal variable (policy parameter) as in (6)
 - 5: Update dual variable (LM) as in (40)
 - 6: **end for**
-

around a local minimum. We chose this stronger assumption for simplicity in the exposition.

Having established *InvLin* and *InvQua* as well as corresponding bounds ϵ_k^1 and ϵ_k^2 , we are able to propose an adaptive primal-dual (APD) algorithm, which is summarized under Algorithm 1. Notice that Theorem 1 indicating the (approximate) convergence of the LM also holds for the APD algorithm. Furthermore, with ϵ_k^1 and ϵ_k^2 , we can further establish the guarantee of proximity to the primal optimum $J_R(\pi_{\theta^*})$. The following theorem addresses this aspect explicitly.

THEOREM 3. *Consider θ^* and $J_R(\pi_{\theta^*})$ in (3). Let the hypotheses of Theorem 1 hold, and consider the sequence of the LM generated by Algorithm 1. Then, it holds that*

$$\liminf_{K \rightarrow \infty} \frac{1}{K} \sum_{k=0}^{K-1} J_R(\pi_{\theta_{k+1}}) \geq J_R(\pi_{\theta^*}) - \frac{1}{K} \left(\sum_{k=0}^{K-1} \epsilon_k \right) - \frac{\zeta (B + (1 - \gamma) \|\mathbf{d}\|)^2}{2(1 - \gamma)^2}. \quad (37)$$

Proof. See the supplementary material of [32]. \square

Notice that ϵ_k in (37) can be bounded by (28) or (29), depending on *InvLin* or *InvQua* is selected. Theorem 3 demonstrates that the limit inferior of the average of the sequence derived by Algorithm 1 approximates well the value of $J_R(\pi_{\theta^*})$. In principle, the limit superior has the potential to be significantly larger than $J_R(\pi_{\theta^*})$, which would lead to constraint violations. In the following result, we prove that this is not the case, i.e., the sequence generated by Algorithm 1 is feasible on average.

THEOREM 4. *Let hypotheses of Theorem 1 hold. It holds that*

$$\limsup_{K \rightarrow \infty} \frac{1}{K} \sum_{k=0}^{K-1} J_{C_i}(\pi_{\theta_k}) \leq d_i, \quad i = 1, 2, \dots, m. \quad (38)$$

Proof. See the supplementary material of [32]. \square

Despite the theoretical guarantees on APD (Algorithm 1) regarding the convergence, optimality, and feasibility, estimating the Lipschitz constant $L(\lambda)$ and the strongly convex constant μ in the LR (30) and (31) is, in general, computationally expensive. We

thus consider a practical version of adaptive LRs. Under the single constraint, the adaptive LRs in (30) and (31) can be written as

$$\eta_k^1 = H_1/(\lambda_k + H_2), \quad \eta_k^2 = H_1'/(\lambda_k + H_2')^2, \quad (39)$$

where H_1, H_2, H_1', H_2' are hyper-parameters.

Note that the update on the policy parameter is updated by running a step of any RL algorithm with the adaptive LRs in (39). In particular, we consider state-of-the-art algorithms such as PPOL and DDPGL. However, the dual variables frequently exhibit tendencies to overshoot and oscillate, thereby hindering performance. To address these challenges, we adopt the Proportional Integral Derivative (PID) Lagrangian strategies as introduced by [34]. Drawing inspiration from the feedback control, the LM is updated as

$$\lambda_k = \left(K_P(J_C^k - d) + K_I I^k + K_D(J_C^k - J_C^{k-1}) \right)_+, \quad (40)$$

where I^k is recursively defined as $I^k = (I^{k-1} + J_C^k - d)_+$ and subscript $+$ indicates projection onto the non-negative orthant. The three hyper-parameters $K_P, K_I,$ and K_D represent the proportional, integral, and derivative gains. In particular, selecting $K_P = K_D = 0$ the update reduces to gradient descent on the dual domain.

With the practical adaptive LRs (*InvLin* and *InvQua*) defined as in (39), as well as the dual update rule (40) using PID-Lagrangian, the practical version of APD (PAPD) algorithm, which we implement and evaluate in the subsequent section, is described in Algorithm 2.

4 EXPERIMENTS

The experimental details are deferred to the supplementary material of [32].

4.1 Environment

To validate our findings, we consider the Bullet-Safety-Gym environments [30] and the Fast Safe Reinforcement Learning (FSRL) framework [35] in this work. The Bullet-Safety-Gym is a platform designed to train and evaluate safety features in constrained RL scenarios. Meanwhile, the FSRL library offers structured modules for implementing SRL algorithms including PPOL and DDPGL.

Table 1: Running time (seconds) using PPOL for all experiments. Each case contains five independent runs.

Environment	<i>BallRun</i>	<i>CarRun</i>	<i>BallCircle</i>	<i>CarCircle</i>
<i>InvLin</i>	128.1 ± 5.2	159.0 ± 5.6	223.8 ± 5.6	861.0 ± 28.2
<i>InvQua</i>	128.1 ± 2.7	161.9 ± 14.5	228.2 ± 11.1	993.2 ± 33.8
LR=0.0001	127.1 ± 1.5	162.0 ± 13.8	234.4 ± 2.6	984.5 ± 56.1
LR=0.00025	125.8 ± 4.3	164.0 ± 18.4	230.3 ± 9.1	961.9 ± 61.1
LR=0.0005	124.4 ± 3.9	170.9 ± 17.7	222.8 ± 13.6	908.8 ± 35.1
LR=0.001	120.5 ± 6.3	163.6 ± 22.7	214.7 ± 14.5	868.9 ± 27.9

4.2 Results

We compare Algorithm 2 with the constant LR primal-dual algorithm in four environments of Bullet-Safety-Gym: *SafetyBallRun-v0*, *SafetyCarRun-v0*, *SafetyBallCircle-v0*, *SafetyCarCircle-v0* (described in the supplementary material of [32]). For a fair comparison, we

maintain uniformity in all parameters and hyper-parameters (except for the LR) across each case (see the supplementary material of [32] for details of the hyper-parameters). While Algorithm 2 relies on five hyper-parameters (other than θ_0), its performance is fairly robust to the choice of these parameter values. We employ the default values of K_P, K_I, K_D in FSRL across all experiments. In addition, we numerically substantiate the robustness of the algorithm performance against variations in $H_1(H_1'), H_2(H_2')$ values in Section 4.3.

Figure 2 depicts the training curves of return, cost, and LR using PPOL over five random seeds. The solid line illustrates the mean and the shaded area depicts the minimum and maximum values across seeds. More specifically, the smallest constant LR (LR = 0.0001) has the worst performance in all four environments. Despite the stable training process, this LR makes significant sacrifices in both aspects of convergence rate and optimal value (return). To some extent, the above issue can be alleviated by using a larger LR. The purple curves in Figure 2 show a better performance achieved by LR = 0.00025. Nevertheless, the performance is not maintained as one keeps increasing LR. Indeed, the experiments in *SafetyBallCircle-v0* present that the training processes become more unstable (larger variance) and/or converge to a worse solution when LR = 0.0005 is selected. Ultimately, if LR is continuously raised until it reaches 0.001, all experiments exhibit either a significant fluctuation in return and cost, or in some cases, an even worse average performance compared to using a LR of 0.0005.

On the contrary, PAPD with *InvLin* and *InvQua* outperform all constant-LR cases in *SafetyBallCircle-v0* and *SafetyCarCircle-v0*, and achieve comparable performance in terms of return and cost with the best constant-LR trials in *SafetyBallRun-v0* and *SafetyCarRun-v0*. This stems from the common criterion in the optimization literature, where the solutions with higher returns but violating constraints (infeasible) are regarded as having “inferior performance”. Moreover, we present the running time employing PPOL for all experiments, as detailed in Table 1, where *BallRun*, *CarRun*, *BallCircle*, *CarCircle* serve as succinct abbreviations for the corresponding environments, namely *SafetyBallRun-v0*, *SafetyCarRun-v0*, *SafetyBallCircle-v0*, *SafetyCarCircle-v0*. Table 1 unveils the absence of substantial difference in running time between our PAPD algorithm and the constant-LR baselines. In addition, it is noteworthy that we use the same hyper-parameters $H_1 = 0.001, H_2 = 3$ in *InvLin* and same $H_1' = 0.015, H_2' = 6$ in *InvQua* across all experiments.

We also apply DDPGL to update policy parameters in the PAPD algorithm. Similar to what we observe in Figure 2, *InvLin* and *InvQua* employing DDPGL surpasses or matches the best performance of all constant-LR cases. Given the limited space, the results of DDPGL experiments are meticulously detailed and analyzed in the supplementary material of [32]. Likewise, we maintain uniformity in hyper-parameter values throughout all conducted experiments. This consistency, despite the varied experimental scenarios, speaks to the robustness of our PAPD approach. Further evidence supporting this robustness statement will be discussed in the next subsection.

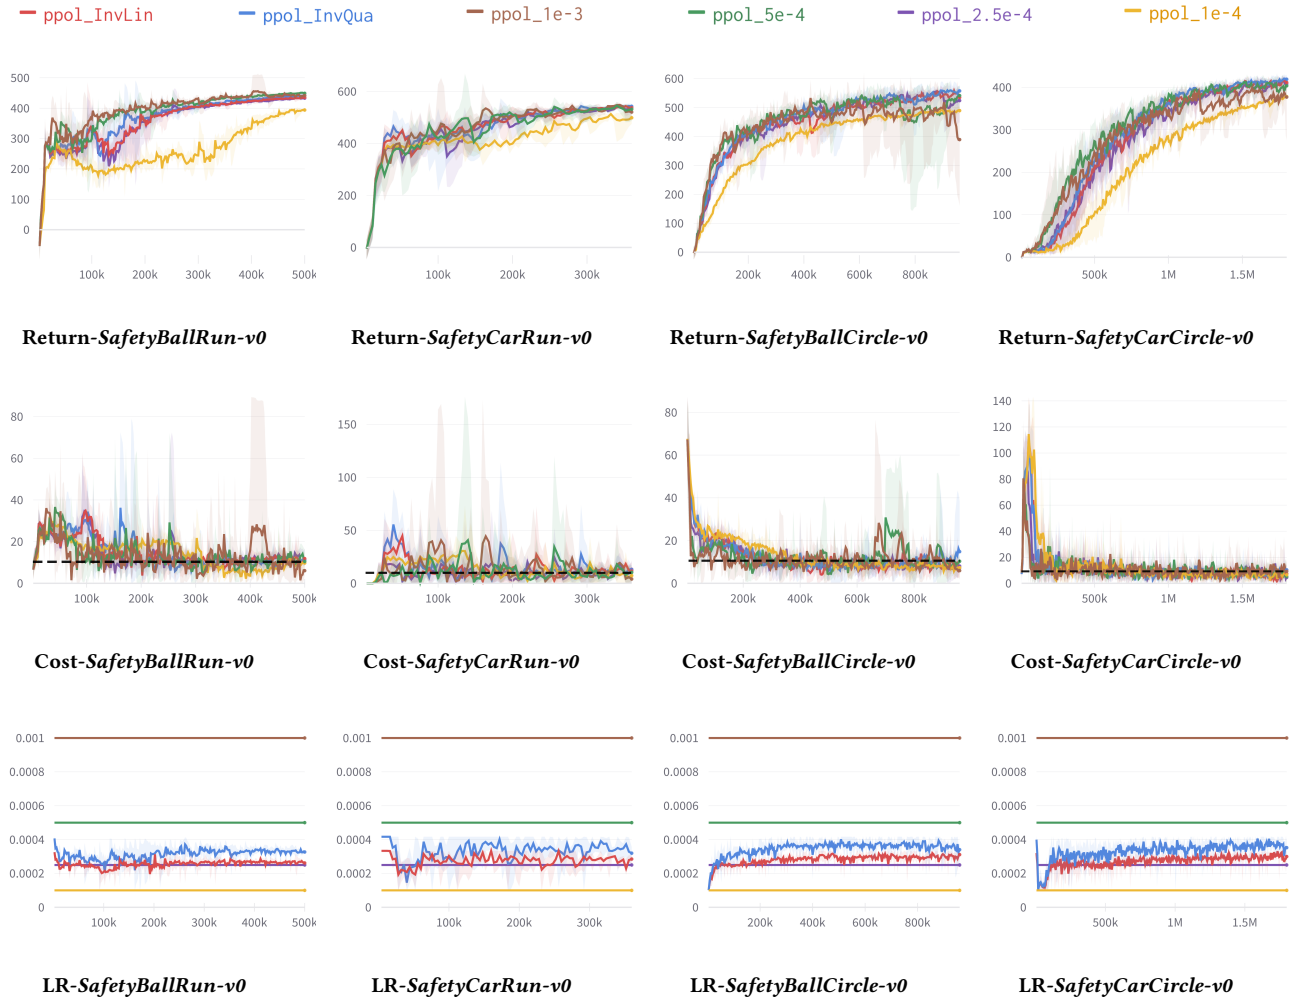


Figure 2: Learning curves for PPOL over four environments with five independent runs. In all figures, the horizontal axis is the number of time step. The solid line illustrates the mean and the shaded area depicts the maximum and the minimum. In all experiments, $H_1 = 0.001, H_2 = 3$ for *InvLin*, $H'_1 = 0.015, H'_2 = 6$ for *InvQua*, and cost limit $d = 10$ (black dashed line).

4.3 Robustness Verification

Having observed the sensitivity of experimental results with respect to the constant-LR, one might naturally find themselves intrigued by the sensitivity exhibited by $H_1(H'_1), H_2(H'_2)$. As shown in (39), $H_1(H'_1)$ plays a similar role to constant-LR in a natural way. Therefore, a compelling point of interest would involve their comparison with the constant-LR. On the other hand, $H_2(H'_2)$ is added to the LM λ_k . Indeed, for large values of λ_k , $H_2(H'_2)$ becomes more negligible, whereas for small values of λ_k , the LR will focus on $H_2(H'_2)$ itself. The LM describes the level of complexity involved in solving the problem, which indicates that $H_2(H'_2)$ generally depends on the problems/tasks. In this subsection, we substantiate the robustness of $H_1(H'_1)$ and $H_2(H'_2)$, respectively. The results are summarized in Tables 2, 3 and 4.

Table 2 summarizes the experimental results of PPOL algorithm in *SafetyCarRun-v0*, where each case contains five independent runs.

More concretely, the parameters in the first block are the same as in Figure 2, where constant LR (0.00025) has the most comparable performance with *InvLin* ($H_1 = 0.001$) and *InvQua* ($H'_1 = 0.015$), and we thus select it for robustness verification. In the next three blocks, the constant LR, H_1 , and H'_1 are reduced by the same proportion, and we also increase the three parameters by the same scale in the last four blocks. These 8 blocks contain a wide range of constant LR, H_1 , H'_1 , and thus enable us to fairly compare their sensitivity.

The goal of problem (3) is to maximize the expected return while satisfying the constraint, i.e., $\text{cost} < 10$ in our experiments. With this in mind, Table 2 shows that *InvLin* achieves the best performance in the first, second, and fifth blocks, while *InvQua* stands out as the epitome of exceptional performance in the rest of blocks. The fifth block is the specialty, where all three cases are infeasible. In this case, *InvLin* is selected to be the best due to the smallest constraint violations and almost the largest return. To summarize,

Table 2: Robustness verification for H_1/H'_1 using *SafetyCarRun-v0* and PPOL. Each case contains five independent runs. In all experiments, $H_2 = 3/H'_2 = 6$ are fixed as in Figure 2. The first block showcases three baselines: constant LR = 0.00025 (Baseline1), *InvLin* $H_1 = 0.001$ (Baseline2), *InvQua* $H'_1 = 0.015$ (Baseline3), as selected in Figure 2. In the rest of blocks, constant LR, H_1, H'_1 are decreased/increased by the same proportion.

Parameter	Return	Cost	Parameter	Return	Cost
Baseline1	538.10 ± 8.71	12.48 ± 10.96	Baseline1 × 0.2	421.26 ± 40.82	9.64 ± 4.90
Baseline2	537.59 ± 11.40	9.98 ± 5.98	Baseline2 × 0.2	502.48 ± 27.43	8.80 ± 5.15
Baseline3	545.22 ± 4.93	13.40 ± 4.78	Baseline3 × 0.2	513.93 ± 7.35	11.50 ± 4.55
Baseline1 × 0.32	492.65 ± 15.78	7.66 ± 5.34	Baseline1 × 0.4	510.41 ± 13.62	9.22 ± 6.14
Baseline2 × 0.32	511.41 ± 22.70	11.46 ± 7.22	Baseline2 × 0.4	526.35 ± 4.30	18.42 ± 8.38
Baseline3 × 0.32	509.91 ± 29.01	8.94 ± 5.74	Baseline3 × 0.4	519.52 ± 11.61	8.04 ± 5.66
Baseline1 × 1.6	545.26 ± 4.50	15.22 ± 7.48	Baseline1 × 2.4	539.62 ± 19.01	10.02 ± 10.87
Baseline2 × 1.6	548.04 ± 8.59	14.32 ± 7.21	Baseline2 × 2.4	528.23 ± 28.56	4.74 ± 5.68
Baseline3 × 1.6	548.45 ± 7.52	15.96 ± 11.93	Baseline3 × 2.4	542.24 ± 8.92	8.28 ± 8.27
Baseline1 × 3.2	536.08 ± 12.57	6.42 ± 8.25	Baseline1 × 4.0	529.97 ± 15.72	1.26 ± 1.08
Baseline2 × 3.2	544.97 ± 8.28	10.70 ± 5.69	Baseline2 × 4.0	533.70 ± 19.98	9.56 ± 10.51
Baseline3 × 3.2	538.08 ± 11.93	5.66 ± 6.96	Baseline3 × 4.0	534.07 ± 20.06	7.70 ± 11.52

Table 3: Summary of Table 2

Paramter	Best Performance	Overall Return	Overall Cost
Constant LR	536.08/6.42	514.17 ± 41.45	8.99 ± 4.15
<i>InvLin</i> - H_1	537.59/9.98	529.10 ± 15.74	10.99 ± 4.03
<i>InvQua</i> - H'_1	542.24/8.28	531.43 ± 14.92	9.94 ± 3.41

Table 4: Robustness verification for H_2/H'_2 using *SafetyCarRun-v0* and PPOL. Each case contains five independent runs. In all experiments, $H_1 = 0.001/H'_1 = 0.015$ are fixed as in Figure 2. The two baselines, denoted as BS1 and BS2, correspond to *InvLin* $H_2 = 3$ and *InvQua* $H'_2 = 6$ as selected in Figure 2. Starting from the first block, H_2/H'_2 are increased by 10%.

H_2/H'_2	Return	Cost	H_2/H'_2	Return	Cost
BS1 × 0.8	543.05 ± 10.21	12.40 ± 8.66	BS1 × 0.9	531.73 ± 12.61	5.80 ± 6.15
BS2 × 0.8	545.09 ± 3.82	6.56 ± 5.05	BS2 × 0.9	543.58 ± 6.84	8.04 ± 5.26
BS1 × 1.1	538.27 ± 5.70	8.86 ± 6.90	BS1 × 1.2	531.27 ± 8.97	7.90 ± 6.14
BS2 × 1.1	547.12 ± 2.44	13.28 ± 2.49	BS2 × 1.2	537.37 ± 7.68	9.50 ± 4.40

our proposed *InvLin* and *InvQua* outperform the constant LR in a wide range of values.

We further analyze Table 2 from the perspective of all blocks. Indeed, Table 3 extracts the best performance of three parameters across all blocks. All of them attain feasible solutions, with *InvQua* emerging as the optimal choice by achieving return = 542.24, cost = 8.28. On the other hand, we collect the mean values of the return and the cost in each block of Table 2, and compute their mean and standard deviation across all eight blocks. These are shown in the last two columns: “Overall Return” and “Overall Cost”. Note that the best “Overall Return-Cost” pair goes to *InvQua* again. In addition, *InvLin* and *InvQua* yield smaller standard deviation than the constant LR case in terms of both return and cost.

Furthermore, Table 4 tests the robustness of $H_2(H'_2)$. We also use the same $H_1(H'_1), H_2(H'_2)$ as in Figure 2 as our baselines. Then, we increase $H_2(H'_2)$ by 10% starting from the first block in Table 4, while keeping $H_1(H'_1)$ constant. Note that we maintain the satisfactory performance in three of the blocks and encounter small loss in the worst case while changing $H_2(H'_2)$. Moreover, *InvQua* shows larger sensitivity to H'_2 , which could be naturally explained by its quadratic format.

In summary, we are able to employ the same $H_1(H'_1), H_2(H'_2)$ in four different Bullet-Safety-Gym environments for both PPOL and DDPGL algorithms. In addition, Tables 2 and 3 indicate that *InvLin* and *InvQua* outperform the constant LR baselines in a wide range of $H_1(H'_1)$ that are selected, and Table 4 validates the robustness of $H_2(H'_2)$ as well. Therefore, these numerical results substantiate the robustness of our proposed PAPD algorithm.

5 CONCLUDING REMARKS

In this work, we propose the Adaptive Primal-Dual (APD) algorithm and its practical version (PAPD) that leverage adaptive learning rates for safe reinforcement learning (SRL). Theoretically, we provide the analyses of the APD algorithm in terms of the convergence, optimality and feasibility. We also numerically evaluate the PAPD algorithm across four well-known SRL environments in the Bullet-Safety-Gym. Our experiments show that the PAPD algorithm outperforms the primal-dual algorithm with constant learning rates. In addition, we validate the robustness of our proposed algorithm by testing across various environments, different RL methods, and a wide range of hyper-parameters associated with the two adaptive learning rates.

ACKNOWLEDGMENTS

This work is supported by the Rensselaer-IBM AI Research Collaboration. The authors thank Zuxin Liu for his thoughtful suggestions and valuable insights.

REFERENCES

- [1] John Schulman, Barret Zoph, Christina Kim, Jacob Hilton, Jacob Menick, Jiayi Weng, Juan Felipe Ceron Uribe, Liam Fedus, Luke Metz, Michael Pokornyy, et al. Chatgpt: Optimizing language models for dialogue. *OpenAI blog*, 2022.
- [2] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- [3] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.
- [4] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pages 1928–1937. PMLR, 2016.
- [5] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *nature*, 550(7676):354–359, 2017.
- [6] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharmashankar Kumar, Thore Graepel, et al. A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science*, 362(6419):1140–1144, 2018.
- [7] Hai Nguyen and Hung La. Review of deep reinforcement learning for robot manipulation. In *2019 Third IEEE International Conference on Robotic Computing (IRC)*, pages 590–595. IEEE, 2019.
- [8] Sergey Levine, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. End-to-end training of deep visuomotor policies. *The Journal of Machine Learning Research*, 17(1):1334–1373, 2016.
- [9] Javier Garcia and Fernando Fernández. A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research*, 16(1):1437–1480, 2015.
- [10] Shangding Gu, Long Yang, Yali Du, Guang Chen, Florian Walter, Jun Wang, Yaodong Yang, and Alois Knoll. A review of safe reinforcement learning: Methods, theory and applications. *arXiv preprint arXiv:2205.10330*, 2022.
- [11] Weiqin Chen, Dharmashankar Subramanian, and Santiago Paternain. Policy gradients for probabilistic constrained reinforcement learning. In *2023 57th Annual Conference on Information Sciences and Systems (CISS)*, pages 1–6. IEEE, 2023.
- [12] Gregory Kahn, Adam Villafior, Bosen Ding, Pieter Abbeel, and Sergey Levine. Self-supervised deep reinforcement learning with generalized computation graphs for robot navigation. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5129–5136. IEEE, 2018.
- [13] Kai Zhu and Tao Zhang. Deep reinforcement learning based mobile robot navigation: A review. *Tsinghua Science and Technology*, 26(5):674–691, 2021.
- [14] Peter Geibel. Reinforcement learning for mdps with constraints. In *European Conference on Machine Learning*, pages 646–653. Springer, 2006.
- [15] Yoshinobu Kadota, Masami Kurano, and Masami Yasuda. Discounted markov decision processes with utility constraints. *Computers & Mathematics with Applications*, 51(2):279–284, 2006.
- [16] Yinlam Chow, Mohammad Ghavamzadeh, Lucas Janson, and Marco Pavone. Risk-constrained reinforcement learning with percentile risk criteria. *The Journal of Machine Learning Research*, 18(1):6070–6120, 2017.
- [17] Eitan Altman. *Constrained Markov decision processes*, volume 7. CRC press, 1999.
- [18] Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. Constrained policy optimization. In *International conference on machine learning*, pages 22–31. PMLR, 2017.
- [19] Tsung-Yen Yang, Justinian Rosca, Karthik Narasimhan, and Peter J Ramadge. Projection-based constrained policy optimization. *arXiv preprint arXiv:2010.03152*, 2020.
- [20] Yiming Zhang, Quan Vuong, and Keith Ross. First order constrained optimization in policy space. *Advances in Neural Information Processing Systems*, 33:15338–15349, 2020.
- [21] Chen Tessler, Daniel J Mankowitz, and Shie Mannor. Reward constrained policy optimization. *arXiv preprint arXiv:1805.11074*, 2018.
- [22] Santiago Paternain, Luiz Chamon, Miguel Calvo-Fullana, and Alejandro Ribeiro. Constrained reinforcement learning has zero duality gap. *Advances in Neural Information Processing Systems*, 32, 2019.
- [23] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International conference on machine learning*, pages 1889–1897. PMLR, 2015.
- [24] Linrui Zhang, Li Shen, Long Yang, Shixiang Chen, Bo Yuan, Xueqian Wang, and Dacheng Tao. Penalized proximal policy optimization for safe reinforcement learning. *arXiv preprint arXiv:2205.11814*, 2022.
- [25] Shalabh Bhatnagar and K Lakshmanan. An online actor-critic algorithm with function approximation for constrained markov decision processes. *Journal of Optimization Theory and Applications*, 153:688–708, 2012.
- [26] Vivek S Borkar. An actor-critic algorithm for constrained markov decision processes. *Systems & control letters*, 54(3):207–213, 2005.
- [27] Dongsheng Ding, Kaiqing Zhang, Tamer Basar, and Mihailo Jovanovic. Natural policy gradient primal-dual method for constrained markov decision processes. *Advances in Neural Information Processing Systems*, 33:8378–8390, 2020.
- [28] Dongsheng Ding, Xiaohan Wei, Zhuoran Yang, Zhaoran Wang, and Mihailo Jovanovic. Provably efficient safe exploration via primal-dual policy optimization. In *International Conference on Artificial Intelligence and Statistics*, pages 3304–3312. PMLR, 2021.
- [29] Weiqin Chen, Dharmashankar Subramanian, and Santiago Paternain. Probabilistic constraint for safety-critical reinforcement learning. *arXiv preprint arXiv:2306.17279*, 2023.
- [30] Sven Gronauer. Bullet-safety-gym: A framework for constrained reinforcement learning. 2022.
- [31] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- [32] Weiqin Chen, James Onyejizu, Long Vu, Lan Hoang, Dharmashankar Subramanian, Koushik Kar, Sandipan Mishra, and Santiago Paternain. Adaptive primal-dual method for safe reinforcement learning. *arXiv preprint arXiv:2402.00355*, 2024.
- [33] Yair Censor. Pareto optimality in multiobjective problems. *Applied Mathematics and Optimization*, 4(1):41–59, 1977.
- [34] Adam Stooke, Joshua Achiam, and Pieter Abbeel. Responsive safety in reinforcement learning by pid lagrangian methods. In *International Conference on Machine Learning*, pages 9133–9143. PMLR, 2020.
- [35] Zuxin Liu, Zijian Guo, Haohong Lin, Yihang Yao, Jiacheng Zhu, Zhepeng Cen, Hanjiang Hu, Wenhao Yu, Tingnan Zhang, Jie Tan, et al. Datasets and benchmarks for offline safe reinforcement learning. *arXiv preprint arXiv:2306.09303*, 2023.