

# Penalized Proximal Policy Optimization for Safe Reinforcement Learning

Linrui Zhang<sup>1\*</sup>, Li Shen<sup>2†</sup>, Long Yang<sup>3</sup>, Shixiang Chen<sup>2</sup>  
 Xueqian Wang<sup>1</sup>, Bo Yuan<sup>1</sup>, Dacheng Tao<sup>2</sup>

<sup>1</sup>Tsinghua University

<sup>2</sup>JD Explore Academy

<sup>3</sup>Peking University

zlr20@mails.tsinghua.edu.cn, {mathshenli,chenshixiang,dacheng.tao}@gmail.com,  
 yanglong001@pku.edu.cn, boyuan@ieee.org, wang.xq@sz.tsinghua.edu.cn

## Abstract

Safe reinforcement learning aims to learn the optimal policy while satisfying safety constraints, which is essential in real-world applications. However, current algorithms still struggle for efficient policy updates with hard constraint satisfaction. In this paper, we propose Penalized Proximal Policy Optimization (P3O), which solves the cumbersome constrained policy iteration via a single minimization of an equivalent unconstrained problem. Specifically, P3O utilizes a simple-yet-effective penalty function to eliminate cost constraints and removes the trust-region constraint by the clipped surrogate objective. We theoretically prove the exactness of the proposed method with a finite penalty factor and provide a worst-case analysis for approximate error when evaluated on sample trajectories. Moreover, we extend P3O to more challenging multi-constraint and multi-agent scenarios which are less studied in previous work. Extensive experiments show that P3O outperforms state-of-the-art algorithms with respect to both reward improvement and constraint satisfaction on a set of constrained locomotive tasks.

## 1 Introduction

Reinforcement Learning (RL) has achieved significant successes in playing video games [Vinyals *et al.*, 2019], robotic manipulation [Levine *et al.*, 2016], mastering Go [Silver *et al.*, 2017], etc. However, standard RL merely maximizes cumulative rewards, which may lead to undesirable behaviors in real-world applications especially for safety-critical tasks.

It is intractable to learn reasonable policies by directly penalizing unsafe interactions onto the reward function, since varied intensities of the punishment result in different Markov decision processes. There is still a lack of theories revealing the explicit relationship between policy improvement and safety satisfaction via reward shaping.

\*This work was done during Linrui Zhang’s internship at JD Explore Academy, Beijing, China.

†Corresponding Author.

Constrained Markov decision process (CMDP) [Altman, 1999] is a more practical and popular formulation that requires the agent to perform actions under given constraints. A common approach to solving such constrained sequential optimization is approximating the non-convex constraint function to a quadratic optimization problem via Taylor’s formulation [Achiam *et al.*, 2017; Yang *et al.*, 2020]. However, those algorithms still have the following drawbacks: (1) The convex approximation to a non-convex policy optimization results in non-negligible approximate errors and only learns near constraint-satisfying policies; (2) The closed-form solution to the primal problem involves the inversion of Hessian matrix, which is computationally expensive in large CMDPs with deep neural networks; (3) If there is more than one constraint, it is cumbersome to obtain the analytical solution without an extra inner-loop optimization, which limits their applications in multi-constraint and multi-agent scenarios; (4) If the primal problem is infeasible under certain initial policies, above algorithms require additional interactions for feasible recovery, which reduces their sample-efficiency.

To address above issues, we propose **Penalized Proximal Policy Optimization (P3O)** algorithm that contains two key techniques. Firstly, we employ exact penalty functions to derive an equivalent unconstrained optimization problem that is naturally compatible with multiple constraints and arbitrary initial policies. Secondly, we extend the clipped surrogate objective in Proximal Policy Optimization (PPO) [Schulman *et al.*, 2017] to CMDPs, which eliminates the trust-region constraint on both reward term and cost term. As a consequence, our method removes the need for quadratic approximation and Hessian matrix inversion. By minimizing the unconstrained objective instead, the solution to the primal constrained problem can be obtained.

Conclusively, the proposed P3O algorithm has three main strengths: (1) *accuracy*: P3O algorithm is accurate and sample-efficient with first-order optimization over a theoretically equivalent unconstrained objective instead of solving an approximate quadratic optimization; (2) *feasibility*: P3O algorithm could admit arbitrary initialization with a consistent objective function and doesn’t need additional recovery methods for infeasible policies; (3) *scalability*: P3O is naturally scalable to multiple constraints or agents without increasing complexity due to the exact penalty reformulation.

In the end, we summarize our contributions as four-fold:

- We propose the Penalized Proximal Policy Optimization (P3O) algorithm for safe RL, which employs first-order optimization over an equivalent unconstrained objective.
- We theoretically prove the exactness of the penalized method when the penalty factor is sufficiently large (doesn't have to go towards positive infinity) and propose an adaptive penalty factor tuning algorithm.
- We extend clipped surrogate objectives and advantage normalization tricks to CMDPs, which are easy to implement and enable a fixed penalty factor for general good results across different tasks.
- We conduct extensive experiments to show that P3O outperforms several state-of-the-art algorithms with respect to both reward improvement and constraint satisfaction and demonstrate its efficacy in more difficult multi-constraint and multi-agent scenarios which are less studied in previous safe RL algorithms.

## 2 Related Work

Safety in reinforcement learning is a challenging topic formally raised by Garcia and Fernández [2015]. Readers can refer to the survey [Liu *et al.*, 2021] for recent advances in safe RL. In this section, we only summarize the most related studies to our algorithm.

**Primal-Dual solution.** Considering the success of Lagrangian relaxation in solving constrained optimization problems, Primal-Dual Optimization (PDO) [Chow *et al.*, 2017] and its variants [Tessler *et al.*, 2018; Stooke *et al.*, 2020] leverage the Lagrangian Duality in constrained reinforcement learning. Although plenty of work has provided rigorous analysis on duality gap [Paternain *et al.*, 2019], non-asymptotic convergence rate [Ding *et al.*, 2020] and regret bound [Ding *et al.*, 2021], Primal-dual methods are still hard to be implemented and applied in practical use due to their sensitivity to the initialization as well as the learning rate of Lagrangian multipliers.

**Primal solution.** Constrained Policy Optimization (CPO) [Achiam *et al.*, 2017] directly searches the feasible policy in the trust region and guarantees a monotonic performance improvement while ensuring constraint satisfaction by solving a quadratic optimization problem with appropriate approximations. Projection-based CPO [Yang *et al.*, 2020] updates the policy in two stages by firstly performing regular TRPO [Schulman *et al.*, 2015] and secondly projecting the policy back into the constraint set. Those methods based on local policy search mainly suffer from approximate errors, the inversion of high-dimensional Hessian matrices, and the poor scalability to multiple constraints. To address drawbacks of the quadratic approximation, FOCOPS [Zhang *et al.*, 2020] solves the optimum of constrained policy optimization within the non-parametric space and then derives the first-order gradients of the  $\ell_2$  loss function within the parameter space. Nevertheless, FOCOPS has more auxiliary variables to learn than our first-order optimization objective, and the analytical solutions of FOCOPS are not straightforward with the increasing number of constraints.

The most similar work to our proposed algorithm is Interior-point Policy Optimization (IPO) [Liu *et al.*, 2020] which uses log-barrier functions as penalty terms to restrict policies into the feasible set. However, the interior-point method requires a feasible policy upon initialization which is not necessarily fulfilled and needs a further recovery. Moreover, the log-barrier function possibly leads to numerical issues when the penalty term goes towards infinity, or the solution is not exactly the same as the primal problem. By contrast, we employ the exact penalty function to derive an equivalent unconstrained objective and restrict policy updates in the trust region by clipping the important sampling ratio on both reward and cost terms for the approximate error reduction, which is different from the prior work.

## 3 Preliminaries

A Markov decision process (MDP) [Sutton and Barto, 2018] is a tuple  $(S, A, R, P, \mu)$ , where  $S$  is the state space,  $A$  is the action space,  $R : S \times A \times S \mapsto \mathbb{R}$  is the reward function,  $P : S \times A \times S \mapsto [0, 1]$  is the transition probability function to describe the dynamics of the environment, and  $\mu : S \mapsto [0, 1]$  is the initial state distribution. A stationary policy  $\pi : S \mapsto \mathcal{P}(A)$  maps the given states to probability distributions over action space and  $\Pi$  denotes the set of all stationary policies  $\pi$ . The optimal policy  $\pi^*$  maximizes the expected discounted return  $J_R(\pi) = \mathbb{E}_{\tau \sim \pi} [\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t, s_{t+1})]$  where  $\tau = \{(s_t, a_t)\}_{t \geq 0}$  is a sample trajectory and  $\tau \sim \pi$  indicates the distribution over trajectories depending on  $s_0 \sim \mu, a_t \sim \pi(\cdot | s_t), s_{t+1} \sim P(\cdot | s_t, a_t)$ . The value function is defined as  $V_R^\pi(s) = \mathbb{E}_{\tau \sim \pi} [\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t, s_{t+1}) | s_0 = s]$ , the action-value function is defined as  $Q_R^\pi(s, a) = \mathbb{E}_{\tau \sim \pi} [\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t, s_{t+1}) | s_0 = s, a_0 = a]$  and the advantage function is defined as  $A_R^\pi(s, a) = Q_R^\pi(s, a) - V_R^\pi(s)$ .

A constrained Markov decision process (CMDP) [Altman, 1999] extends MDP with a constraint set  $C$  consisting of cost functions  $C_i : S \times A \times S \mapsto \mathbb{R}, i = 1, 2, \dots, m$ . We define the expected discounted cost-return  $J_{C_i}(\pi) = \mathbb{E}_{\tau \sim \pi} [\sum_{t=0}^{\infty} \gamma^t C_i(s_t, a_t, s_{t+1})]$  and  $V_{C_i}^\pi, Q_{C_i}^\pi, A_{C_i}^\pi$  similarly in MDPs. The set of feasible stationary policies for a CMDP is denoted as  $\Pi_C = \{\pi \in \Pi | J_{C_i}(\pi) \leq d_i, \forall i\}$ , where  $d_i$  is the upper limit of the corresponding safety or cost constraint. The goal of safe RL is to find an optimal policy over the hard safe constraint, i.e.,  $\pi^* = \arg \max_{\pi \in \Pi_C} J_R(\pi)$ .

## 4 Methodology

Iterative search for optimal policy is commonly used in literature [Schulman *et al.*, 2015; Achiam *et al.*, 2017; Zhang *et al.*, 2020]. We formulate the unbiased constrained policy optimization problem over  $s \sim d^\pi$  and  $a \sim \pi$  as:

$$\begin{aligned} \pi_{k+1} &= \arg \max_{\pi} \mathbb{E}_{\substack{s \sim d^\pi \\ a \sim \pi}} [A_R^{\pi_k}(s, a)] \\ \text{s.t. } J_{C_i}(\pi_k) &+ \frac{1}{1 - \gamma} \mathbb{E}_{\substack{s \sim d^\pi \\ a \sim \pi}} [A_{C_i}^{\pi_k}(s, a)] \leq d_i, \forall i. \end{aligned} \quad (1)$$

where  $d^\pi(s) = (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t P(s_t = s | \pi)$  denote the discounted future state distribution.

**Proposition 1.** *The new policy  $\pi_{k+1}$  obtained from the current policy  $\pi_k$  via problem (1) yields a monotonic return improvement and hard constraint satisfaction.*

*Proof.* See the supplemental material.  $\square$

In this paper, we consider the parametric policy  $\pi(\theta)$ , i.e., using the neural network. Let  $r(\theta) = \frac{\pi(\theta)}{\pi(\theta_k)}$  be the importance sampling ratio, then we rewrite the optimization problem (1) as follows:

$$\begin{aligned} \theta_{k+1} &= \arg \min_{\theta} \mathbb{E}_{\substack{s \sim d^\pi \\ a \sim \pi_k}} [-r(\theta) A_R^{\pi_k}(s, a)] \\ \text{s.t. } &\mathbb{E}_{\substack{s \sim d^\pi \\ a \sim \pi_k}} [r(\theta) A_{C_i}^{\pi_k}(s, a)] + (1 - \gamma)(J_{C_i}(\pi_k) - d_i) \leq 0, \forall i. \end{aligned} \quad (2)$$

Different from previous methods [Achiam *et al.*, 2017; Yang *et al.*, 2020] that approximate non-convex problem (2) to convex optimization via Taylor’s formulation on the trust region, we penalize constraints with ReLU operators to the objective function which yields an unconstrained problem:

$$\begin{aligned} \theta_{k+1} &= \arg \min_{\theta} \mathbb{E}_{\substack{s \sim d^\pi \\ a \sim \pi_k}} [-r(\theta) A_R^{\pi_k}(s, a)] + \kappa \sum_i \max\{0, \\ &\mathbb{E}_{\substack{s \sim d^\pi \\ a \sim \pi_k}} [r(\theta) A_{C_i}^{\pi_k}(s, a)] + (1 - \gamma)(J_{C_i}(\pi_k) - d_i)\}. \end{aligned} \quad (3)$$

Intuitively, the penalized term takes effect when the agent violates the corresponding constraint; otherwise, the objective equals the standard policy optimization when all constraints are fulfilled. Below, we theoretically analyze the equivalence between problem (2) and problem (3), i.e., the ReLU operator constructs an exact penalty function with a finite penalty factor  $\kappa$ . To our best knowledge, it is the first work to solve the constrained policy optimization in the perspective of the exact penalty method.

**Theorem 1.** *Suppose  $\bar{\lambda}$  is the corresponding Lagrange multiplier vector for the optimum of problem (2). Let the penalty factor  $\kappa$  be a sufficiently large constant ( $\kappa \geq \|\bar{\lambda}\|_\infty$ ), problem (2) and problem (3) share the same optimal solution set.*

*Proof.* See the supplemental material.  $\square$

Notably, the finiteness of penalty factor  $\kappa$  in Theorem 1 is critical for policy updates because popular methods like square loss function and log-barrier function require the penalty term to go towards positive infinity for an infeasible solution. Otherwise, the optimal solution is not exactly the same as problem (2) [Liu *et al.*, 2020]. However, this would bring about numerical issues in practice and result in unbounded approximate error which will be discussed later.

By now, we construct an equivalent unconstrained objective with ReLU operators and a finite penalty factor  $\kappa$ . It is still intractable to solve problem (3) directly over the unknown future state distribution  $s \sim d^\pi$  except for cumbersome off-policy evaluation [Jiang and Li, 2016] from current policy  $\pi_k$ . We replace  $s \sim d^\pi$  with  $s \sim d^{\pi_k}$  and minimize

---

### Algorithm 1 P3O: Penalized Proximal Policy Optimization

---

**Input:** initial policy  $\pi_{\theta_0}$ , value function  $V_R^{\phi_0}$  and each cost-value function  $V_{C_i}^{\psi_i^0}, \forall i$ .

- 1: **for**  $k$  in  $0, 1, 2, \dots, K$  **do**
- 2:   Collect set of trajectories  $\mathcal{D}_k$  with policy  $\pi_{\theta_k}$ .
- 3:   Compute  $\hat{R}_t = \sum_{k=0}^{T-t} \gamma^k R_{t+k}$  and  $\hat{A}_R^{\pi_{\theta_k}}(s_t, a_t)$ .
- 4:   Compute  $\hat{C}_t^i = \sum_{k=0}^{T-t} \gamma^k C_{t+k}^i$  and  $\hat{A}_{C_i}^{\pi_{\theta_k}}(s_t, a_t), \forall i$ .
- 5:   Update  $\pi_{\theta_{k+1}}$  using Algorithm 2.
- 6:    $\phi_{k+1} = \arg \min_{\phi} \frac{1}{|\mathcal{D}_k|T} \sum (V_R^{\phi_k}(s_t) - \hat{R}_t)^2$ .
- 7:    $\phi'_{k+1} = \arg \min_{\phi'} \frac{1}{|\mathcal{D}_k|T} \sum (V_{C_i}^{\psi_i^k}(s_t) - \hat{C}_t^i)^2, \forall i$ .
- 8: **end for**

**Output:** Optimal policy  $\pi_{\theta_K}$ .

---



---

### Algorithm 2 Exact Penalized Policy Search Algorithm

---

**Input:** current policy  $\pi_{\theta_k}$ .

- 1: **for**  $n$  in  $0, 1, 2, \dots, N$  **do**
- 2:   Compute  $\mathcal{L}_R^{\text{CLIP}}(\theta)$  in Eq. (6).
- 3:   Compute  $\mathcal{L}_{C_i}^{\text{CLIP}}(\theta), \forall i = 1, 2, \dots, m$  in Eq. (7).
- 4:    $\theta \leftarrow \theta - \eta \nabla \mathcal{L}^{\text{P3O}}(\theta)$  in Eq. (5).
- 5:    $\kappa \leftarrow \min\{\rho\kappa, \kappa_{\max}\}$  ( $\rho > 1$ ).
- 6:   **if**  $\mathbb{E}_{s \sim d^\pi} [\text{D}_{KL}(\pi_\theta || \pi_{\theta_k})[s]] \notin [\delta_k^-, \delta_k^+]$  **then**
- 7:     Break.
- 8:   **end if**
- 9: **end for**
- 10:  $\theta_{k+1} \leftarrow \theta$ .

**Output:** next policy  $\pi_{\theta_{k+1}}$ .

---

the unconstrained objective over collected trajectories:

$$\begin{aligned} \theta_{k+1} &= \arg \min_{\theta} \mathbb{E}_{\substack{s \sim d^{\pi_k} \\ a \sim \pi_k}} [-r(\theta) A_R^{\pi_k}(s, a)] + \kappa \sum_i \max\{0, \\ &\mathbb{E}_{\substack{s \sim d^{\pi_k} \\ a \sim \pi_k}} [r(\theta) A_{C_i}^{\pi_k}(s, a)] + (1 - \gamma)(J_{C_i}(\pi_k) - d_i)\}. \end{aligned} \quad (4)$$

This results in biases w.r.t problem (3), and we further provide the analysis of approximate error for that replacement.

**Theorem 2.** *For any policy  $\pi$  and the given  $\pi_k$ , define*

$$\delta = \mathbb{E}_{s \sim d^{\pi_k}} [\text{D}_{KL}(\pi || \pi_k)[s]].$$

*The worst-case approximate error is  $\mathcal{O}(\kappa m \sqrt{\delta})$  if we replace problem (3) with problem (4).*

*Proof.* See the supplemental material.  $\square$

Theorem 2 indicates the necessity of a finite penalty factor  $\kappa$ , or the approximate error would not be controlled. On the other hand, we have to restrict policy updates within the neighborhood of last  $\pi_k$  (i.e.,  $\delta \rightarrow 0$ ) and improve the policy iteratively. Here, we incorporate the trust-region constraint via clipped surrogate objective [Schulman *et al.*, 2017] in the approximate exactly penalized problem (4) to guarantee a proximal policy iteration. Then, the final optimization objective, abbreviated as P3O (penalized PPO), is derived as:

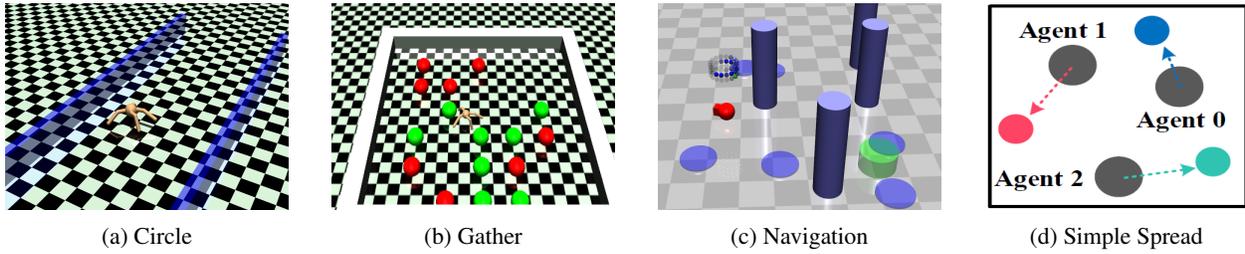


Figure 1: Experimental benchmarks. (a) Circle: The agent is rewarded for moving in a specified wide circle, but is constrained to stay within a safe region smaller than the radius of the circle. (b) Gather: The agent is rewarded for gathering green apples, but is constrained to avoid red bombs. (c) Navigation: The agent is rewarded for reaching the target area (green) but is constrained to avoid virtual hazards (light purple) and impassible pillars (dark purple). Note that the cost for hazards and pillars are calculated separately and hold different upper limits. (d) Simple Spread: Agents are rewarded for reaching corresponding destinations, but are constrained to the mutual collision. The observation of each agent is not shared in the CMDP execution.

$$\mathcal{L}^{\text{P3O}}(\theta) = \mathcal{L}_R^{\text{CLIP}}(\theta) + \kappa \cdot \sum_i \max\{0, \mathcal{L}_{C_i}^{\text{CLIP}}(\theta)\}, \quad (5)$$

where

$$\mathcal{L}_R^{\text{CLIP}}(\theta) = \mathbb{E}_{\substack{s \sim d^{\pi_k} \\ a \sim \pi_k}} \left[ -\min \left\{ r(\theta) A_R^{\pi_k}(s, a), \right. \right. \\ \left. \left. \text{clip}(r(\theta), 1 - \epsilon, 1 + \epsilon) A_R^{\pi_k}(s, a) \right\} \right]; \quad (6)$$

$$\mathcal{L}_{C_i}^{\text{CLIP}}(\theta) = \mathbb{E}_{\substack{s \sim d^{\pi_k} \\ a \sim \pi_k}} \left[ \max \left\{ r(\theta) A_{C_i}^{\pi_k}(s, a), \text{clip}(r(\theta), \right. \right. \\ \left. \left. 1 - \epsilon, 1 + \epsilon) A_{C_i}^{\pi_k}(s, a) \right\} + (1 - \gamma)(J_{C_i}(\pi_k) - d_i) \right]. \quad (7)$$

Based on the clipped surrogate objective technique, we solve the difficult constrained problem (2) by alternately estimating and minimizing the unconstrained loss function (5).

The remaining problem is how to obtain the finite penalty factor  $\kappa$  due to the variety of the magnitude of  $\mathcal{L}_R^{\text{CLIP}}(\theta)$  and each  $\mathcal{L}_{C_i}^{\text{CLIP}}(\theta)$  that depends on different tasks and up-to-date policies. The value of  $\kappa$  is required to be larger than the unknown greatest Lagrange multiplier according to Theorem 1, but Theorem 2 implies too large  $\kappa$  decays the performance. Thus, there is a trade-off among the hyper-parameter tuning. As shown in Algorithm 2, we increase  $\kappa$  at every time step, and the early stopping condition is fulfilled when the distance between solutions of two adjacent steps is small enough or the current policy is out of the trust region. In practice, we utilize the normalization trick that maps the advantage estimation to an approximate standard normal distribution regardless of the tasks themselves. We find this technique enables a fixed  $\kappa$  for general good results across different tasks. Experimental results show that both of above algorithms work effectively and the learning processes are stable in a wide range of  $\kappa$ .

Note that the first term in the P3O objective (5) is the same as standard PPO and the rest terms are the clipped constraint objectives that prevent the new policy from jumping out of the trust region during the penalty reduction. The loss function in P3O is differentiable almost everywhere and can be solved easily via the simple first-order optimizer [Kingma and Ba, 2014]. Conclusively, P3O inherits the benefits of PPO in CMDPs such as better accuracy and sample efficiency

than the approximating quadratic optimization [Achiam *et al.*, 2017; Yang *et al.*, 2020]. Additionally, the P3O algorithm is naturally scalable to multi-constraint scenarios by adding penalized terms onto the objective (5). The pseudo-code of the P3O algorithm for solving general multi-constraint CMDPs is shown in Algorithm 1. In the end, we extend the P3O algorithm to the decentralized multi-agent setting in which the CMDP is partially observable and the observation can not be shared between each agent. Due to the space limitation, detailed algorithms are provided in the supplementary material.

## 5 Experiments

In this section, we empirically demonstrate the efficacy of the P3O algorithm from the following four perspectives:

- P3O outperforms the state-of-the-art algorithms, e.g. CPO [Achiam *et al.*, 2017], PPO-Lagrangian [Ray *et al.*, 2019] and FOCOPS [Zhang *et al.*, 2020] in safe RL benchmarks with a single constraint.
- P3O is robust to more stochastic and complex environments [Ray *et al.*, 2019] where previous methods fail to simultaneously improve the return effectively and satisfy the constraint strictly.
- P3O introduces only one additional hyper-parameter (i.e., the penalty factor  $\kappa$ ), which is easy to tune and insensitive to different settings and constraint thresholds.
- P3O can be extended to multi-constraint and multi-agent scenarios that are barely tested in the previous work.

We design and conduct experiments in 2 single-constraint (*Circle* and *Gather*), 1 multi-constraint (*Navigation*) and 1 multi-agent (*Simple Spread*) safe RL environments respectively, as illustrated in Figure 1.

To be fair in comparison, the proposed P3O algorithm and FOCOPS [Zhang *et al.*, 2020] are implemented with same rules and tricks on the code-base of Ray *et al.* [2019]<sup>1</sup> for benchmarking safe RL algorithms. Also, we take standard PPO as the reference which ignores any constraint and serves as an upper bound baseline on the reward performance.

More information about experiment environments and detailed parameters are provided in the supplementary material.

<sup>1</sup><https://github.com/openai/safety-starter-agents>

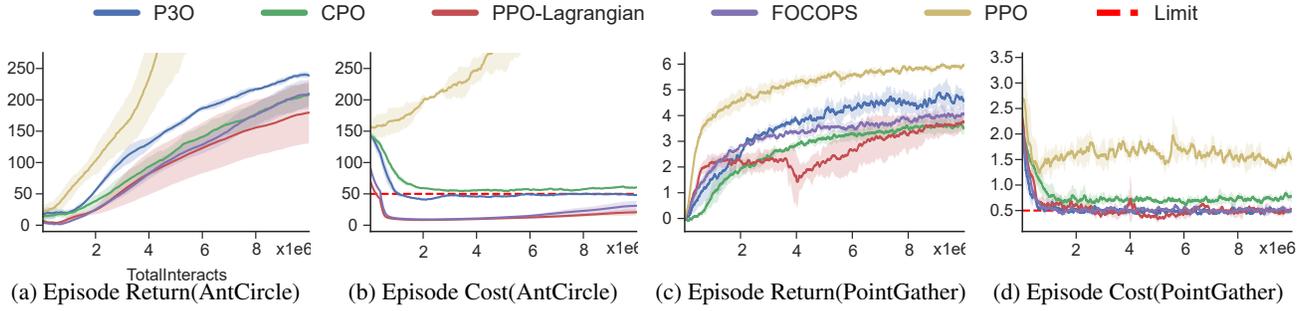


Figure 2: Average episode return and cost in the single-constraint scenario. The x-axis is the number of interactions with the emulator. The y-axis is the average reward/cost-return. The solid line is the mean and the shaded area is the standard deviation. Each tested algorithm runs over five different seeds. The dashed line in the cost plot is the constraint threshold which is 50 for Circle and 0.5 for Gather.

Task		P3O (Ours)	CPO [Achiam <i>et al.</i> , 2017]	PPO-L [Ray <i>et al.</i> , 2019]	FOCOPS [Zhang <i>et al.</i> , 2020]
AntCircle	Reward	253.92 $\pm$ 11.67	214.45 $\pm$ 26.32	189.82 $\pm$ 52.12	232.08 $\pm$ 19.87
	Cost( $\leq 50$ )	46.18 $\pm$ 7.49	57.66 $\pm$ 9.42	27.71 $\pm$ 7.33	34.50 $\pm$ 11.21
PointGather	Reward	4.57 $\pm$ 0.54	3.11 $\pm$ 0.29	3.49 $\pm$ 0.33	3.14 $\pm$ 0.79
	Cost( $\leq 0.5$ )	0.51 $\pm$ 0.10	0.73 $\pm$ 0.17	0.48 $\pm$ 0.09	0.51 $\pm$ 0.14

Table 1: Mean performance with normal 95% confidence interval on single-constraint tasks.

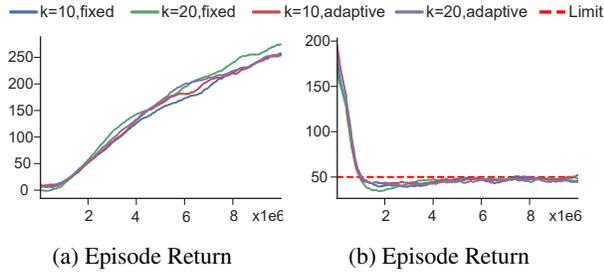


Figure 3: Performance of P3O for different  $\kappa$  settings on AntCircle.

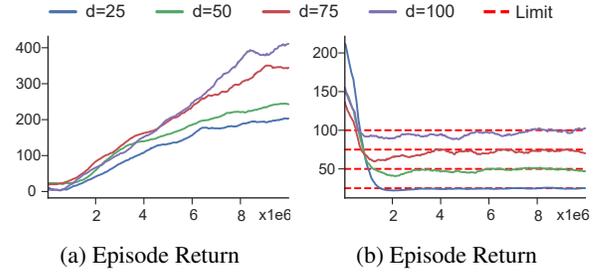


Figure 4: Performance of P3O for different cost limit  $d$  on AntCircle.

**Single-Constraint Scenario.** The numerical results of different algorithms are listed in Table 1. The proposed P3O algorithm has +15%, +24%, +9% higher reward improvement over CPO, PPO-Lagrangian, FOCOPS on the AntCircle task, and +34%, +23%, +31% on the PointGather task respectively. Meanwhile, P3O converges to the upper limit of the safety constraint more tightly which enlarges the parametric search space, as well as satisfies the hard constraint.

As illustrated in Figure 2, P3O is the best for improving the policy while satisfying the given constraint. Conversely, CPO has more constraint violations and even fails to satisfy the constraint after convergence possibly due to Taylor’s approximation. This is especially prominent in complex environments like SafetyGym [Ray *et al.*, 2019] shown in Figure 7. The non-negligible approximate error prevents the CPO agent from fully ensuring constraints on virtually over these environments whereas it enjoys a meaningless high return. As for PPO-Lagrangian, it is too conservative for the trade-off between reward improvement and constraint satisfaction. Besides, it fluctuates mostly on the learning curves which means it is less steady on the learning process. FOCOPS is better-performed than the two baselines above. However, the hyper-

parameter  $\lambda$  takes a great effect on the final performance. At last, both PPO-Lagrangian and FOCOPS are more or less sensitive to the initial value as well as the learning rate of the Lagrange multiplier. Improper hyper-parameters would cause catastrophic performance on reward and constraint value.

**Sensitivity Analysis.** The penalty factor  $\kappa$  plays a critical role in P3O, which is supposed to vary across different tasks depending on the specific value of  $A_R$  and each  $A_{c_i}$ . As shown in Algorithm 2, we can increase  $\kappa$  at every time steps. Otherwise we can map the advantage estimation to an approximate standard normal distribution by applying advantage normalizing tricks and use a fixed  $\kappa$  for general good results. The experimental results are stable in a wide range of  $\kappa$ , as illustrated in Figure 3.

We further verify that P3O works effectively for different threshold levels by changing the cost limit  $d$ . Figure 4 shows P3O can effectively learn constraint-satisfying policies for all the cases. Notably, P3O precisely converges to the cost upper limit. By contrast, the baseline algorithms are too conservative for policy updates and have more overshoot in the process of cost violation reduction.

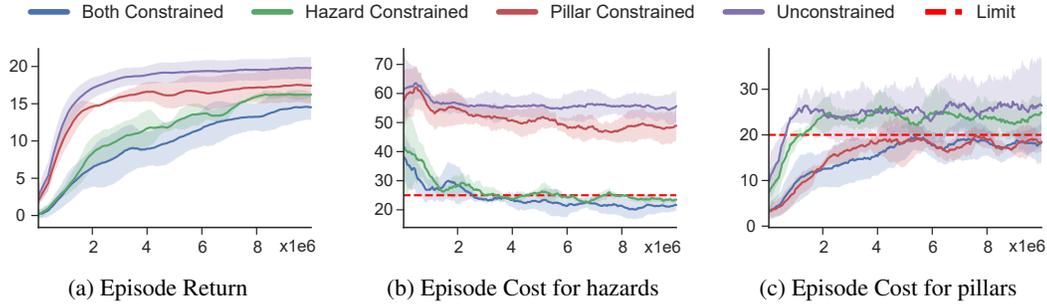


Figure 5: Average episode return(left), cost1(center, for hazards) and cost2(right, for pillars) in the multiple-constraint scenario. The dashed line in the cost plot is the constraint threshold which is 25 for cost1 and 20 for cost2. Hazard/Pillar constrained means only taking cost1/cost2 into P3O loss function whereas ignoring the other one.

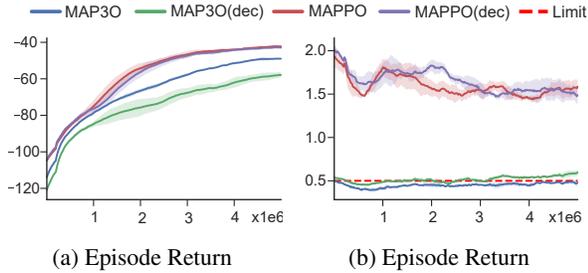


Figure 6: Average episode return and cost in the multi-agent scenario. The algorithm is decentralized with notation (dec).

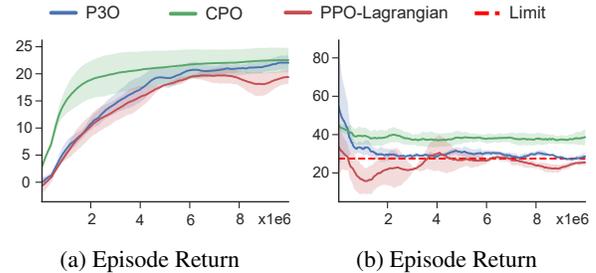


Figure 7: Average performance for P3O, CPO and PPO-Lagrangian on Safetygym PointGoal.

**Multi-Constraint Scenario.** It is convenient to extend P3O to multi-constraint scenarios by simply adding exact penalty functions with the minor modification on the loss function.

Learning curves of average episode return and two independent costs for the navigation task are shown in Figure 5. The unconstrained algorithm is the original PPO. The hazard (pillar) constrained algorithm only takes the cost of hazard (pillar) into P3O loss function whereas ignoring the other one. It can be seen that those single-constrained algorithms maintain the specified constraint under its upper limit but fail to control the other one. When both constraints are incorporated into our algorithm, each one is well satisfied at the convergence. The episode return is lower than unconstrained and single-constrained settings, which is reasonable because the safety constraints are more strict.

It is remarkable that at the beginning of training, the agent violates the hazard constraint but barely hits pillars because there are more small hazards around the starting point and the pillars are nearly unreachable by random moving due to the initialization mechanism. However, with the increasing velocity, the agent is less likely to satisfy the pillar constraint. Satisfaction on the two types of constraints shows that P3O admits both feasible and infeasible initial policies and finds a constraint-satisfying solution in the end.

**Multi-Agent Scenario.** We also study Multi-Agent Penalized Policy Optimization (MAP3O) in the collaborative task and take unconstrained Multi-Agent PPO (MAPPO) [Chao *et al.*, 2021] as the reference. To tackle the issue of partial observability, recurrent networks are used in both algorithms.

The given task is a decentralized partially observable Markov decision process (DEC-POMDP) in which each agent has an exclusive observation and a shared global reward/cost. We replace the MAPPO loss function with Eq. (5) and conduct a fully decentralized training and execution. The experimental results show that MAP3O can improve the total reward while ensuring the mutual collision constraint. Inspired by Chao *et al.*[2021], we enhance MAP3O by learning a centralized critic for both global reward and cost that takes in the concatenation of each agent’s local observation. The centralized version of MAP3O has better reward performance and satisfies the constraint more strictly. The learning curves of average episode return and cost are shown in Figure 6.

## 6 Conclusion

In this paper, we propose Penalized Proximal Policy Optimization (P3O), an effective and easy-to-implement algorithm for safe reinforcement learning. P3O has an unconstrained optimization objective due to the exact penalty reformulation and guarantees better performance on reward and constraint value empirically. Meanwhile, it is naturally compatible with multi-constraint and multi-agent scenarios. P3O adopts first-order minimization and avoids the inversion of high-dimensional Hessian Matrix, which is especially beneficial to large CMDPs with deep neural networks. Thus we consider the future work as end-to-end visual tasks which are less studied in safe reinforcement learning.

## Acknowledgements

This work is supported by Science and Technology Innovation 2030 –“Brain Science and Brain-like Research” Major Project (No. 2021ZD0201402 and No. 2021ZD0201405).

## References

- [Achiam *et al.*, 2017] Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. Constrained policy optimization. In *International Conference on Machine Learning*, pages 22–31. PMLR, 2017.
- [Altman, 1999] Eitan Altman. *Constrained Markov decision processes*, volume 7. CRC Press, 1999.
- [Chao *et al.*, 2021] YU Chao, A VELU, E VINITSKY, et al. The surprising effectiveness of ppo in cooperative, multi-agent games. *arXiv preprint arXiv:2103.01955*, 2021.
- [Chow *et al.*, 2017] 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.
- [Ding *et al.*, 2020] Dongsheng Ding, Kaiqing Zhang, Tamer Basar, and Mihailo R Jovanovic. Natural policy gradient primal-dual method for constrained markov decision processes. In *NeurIPS*, 2020.
- [Ding *et al.*, 2021] 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.
- [Garcia and Fernández, 2015] Javier Garcia and Fernando Fernández. A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research*, 16(1):1437–1480, 2015.
- [Jiang and Li, 2016] Nan Jiang and Lihong Li. Doubly robust off-policy value evaluation for reinforcement learning. In *International Conference on Machine Learning*, pages 652–661. PMLR, 2016.
- [Kingma and Ba, 2014] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [Levine *et al.*, 2016] 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.
- [Liu *et al.*, 2020] Yongshuai Liu, Jiaxin Ding, and Xin Liu. Ipo: Interior-point policy optimization under constraints. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 4940–4947, 2020.
- [Liu *et al.*, 2021] Yongshuai Liu, Avishai Halev, and Xin Liu. Policy learning with constraints in model-free reinforcement learning: A survey. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*, 2021.
- [Paternain *et al.*, 2019] Santiago Paternain, Luiz FO Chamon, Miguel Calvo-Fullana, and Alejandro Ribeiro. Constrained reinforcement learning has zero duality gap. *arXiv preprint arXiv:1910.13393*, 2019.
- [Ray *et al.*, 2019] Alex Ray, Joshua Achiam, and Dario Amodei. Benchmarking safe exploration in deep reinforcement learning. *arXiv preprint arXiv:1910.01708*, 7, 2019.
- [Schulman *et al.*, 2015] 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.
- [Schulman *et al.*, 2017] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [Silver *et al.*, 2017] 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.
- [Stooke *et al.*, 2020] 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.
- [Sutton and Barto, 2018] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- [Tessler *et al.*, 2018] Chen Tessler, Daniel J Mankowitz, and Shie Mannor. Reward constrained policy optimization. *arXiv preprint arXiv:1805.11074*, 2018.
- [Vinyals *et al.*, 2019] Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.
- [Yang *et al.*, 2020] Tsung-Yen Yang, Justinian Rosca, Karthik Narasimhan, and Peter J Ramadge. Projection-based constrained policy optimization. *arXiv preprint arXiv:2010.03152*, 2020.
- [Zhang *et al.*, 2020] Yiming Zhang, Quan Vuong, and Keith W Ross. First order constrained optimization in policy space. *arXiv preprint arXiv:2002.06506*, 2020.