

---

# e-COP : Episodic Constrained Optimization of Policies

---

**Akhil Agnihotri**  
University of Southern California  
agnihotri.akhil@gmail.com

**Rahul Jain**  
Google DeepMind and USC  
rahulajain@google.com

**Deepak Ramachandran**  
Google DeepMind  
ramachandrand@google.com

**Sahil Singla**  
Google DeepMind  
sasingla@google.com

## Abstract

In this paper, we present the e-COP algorithm, the first policy optimization algorithm for constrained Reinforcement Learning (RL) in episodic (finite horizon) settings. Such formulations are applicable when there are separate sets of optimization criteria and constraints on a system’s behavior. We approach this problem by first establishing a policy difference lemma for the episodic setting, which provides the theoretical foundation for the algorithm. Then, we propose to combine a set of established and novel solution ideas to yield the e-COP algorithm that is easy to implement and numerically stable, and provide a theoretical guarantee on optimality under certain scaling assumptions. Through extensive empirical analysis using benchmarks in the Safety Gym suite, we show that our algorithm has similar or better performance than SoTA (non-episodic) algorithms adapted for the episodic setting. The scalability of the algorithm opens the door to its application in safety-constrained Reinforcement Learning from Human Feedback for Large Language or Diffusion Models.

## 1 Introduction

RL problems may be formulated in order to satisfy multiple simultaneous objectives. These can include performance objectives that we want to maximize, and physical, operational or other objectives that we wish to constrain rather than maximize. For example, in robotics, we often want to optimize a task completion objective while obeying physical safety constraints. Similarly, in generative AI models, we want to optimize for human preferences while ensuring that the output generations remain safe (expressed perhaps as a threshold on an automatic safety score that penalizes violent or other undesirable content). Scalable policy optimization algorithms such as TRPO [31], PPO [33], etc have been central to the achievements of RL over the last decade [34, 38, 7]. In particular, these have found utility in generative models, e.g., in the training of Large Language Models (LLMs) to be aligned to human preferences through the RL with Human Feedback (RLHF) paradigm [2]. But these algorithms are designed primarily for the unconstrained infinite-horizon discounted setting: Their use for constrained problems via optimization of the Lagrangian often gives unsatisfactory constraint satisfaction results. This has prompted development of a number of constrained policy optimization algorithms for the infinite-horizon discounted setting [3, 39, 43, 41, 19, 6, 5], and for the average setting [4].

However, many RL problems are more accurately formulated as episodic, i.e., having a finite time horizon. For instance, in image diffusion models [10, 21], the denoising sequence is really a *finite* step trajectory optimization problem, better suited to be solved via RL algorithms for the episodic setting. When existing algorithms for infinite horizon discounted setting are used for such problem, they exhibit sub-optimal performance or fail to satisfy task-specific constraints by prioritizing short-term constraint satisfaction over episodic goals [11, 28, 17]. Furthermore, the episodic setting allows for objective functions to be time-dependent which the infinite-horizon formulations do not. Even when the objective functions are

time-invariant, there is a key difference: for non-episodic settings, a stationary policy that is optimal exists whereas for episodic settings, the optimal policy is always non-stationary and time-dependent. This necessitates development of policy optimization algorithms specifically for the episodic constrained setting. We note that such policy optimization algorithms do not exist even for the unconstrained episodic setting.

In this paper, we introduce e-COP, a policy optimization algorithm for episodic constrained RL problems. Inspired by PPO, it uses deep learning-based function approximation, a KL divergence-based proximal trust region and gradient clipping, along with several novel ideas specifically for the finite horizon episodic and constrained setting. The algorithm is based on a technical lemma (Lemma 3.1) on the difference between two policies, which leads to a new loss function for the episodic setting. We also introduce other ideas that obviate the need for matrix inversion thus improving scalability and numerical stability. The resulting algorithm improves performance over state-of-the-art baselines (after adapting them for the episodic setting). In sum, the e-COP algorithm has the following advantages: (i) *Solution equivalence*: We show that the solution set of our e-COP loss function is same as that of the original CMDP problem, leading to precise cost control during training and avoidance of approximation errors (see Theorem 3.3); (ii) *Stable convergence*: The e-COP algorithm converges tightly to safe optimal policies without the oscillatory behavior seen in other algorithms like PDO [13] and FOCOPS [43]; (iii) *Easy integration*: e-COP follows the skeleton structure of PPO-style algorithms using clipping instead of steady state distribution approximation, and hence can be easily integrated into existing RL pipelines; and (iv) *Empirical performance*: the e-COP algorithm demonstrates superior performance on several benchmark problems for constrained optimal control compared to several state-of-the-art baselines.

**Our Contributions and Novelty.** We introduce the first policy optimization algorithm for episodic RL problems, both with and without constraints (no constraints is a special case). While some of the policy optimization algorithms can be adapted for the constrained setting via a Lagrangian formulation, as we show, they don't work so well empirically in terms of constraint satisfaction and optimality. The algorithm is based on a policy difference (technical) lemma, which is novel. We have gotten rid of Hessian matrix inversion, a common feature of policy optimization algorithms (see, for example, PPO [33], CPO [3], P3O [41], etc.) and replaced it with a quadratic penalty term which also improves numerical stability near the edge of the constraint set boundary - a problem unique to constrained RL problems. We provide an instance-dependent hyperparameter tuning routine that generalizes to various testing scenarios. And finally, our extensive empirical results against an extensive suite of baseline algorithms (e.g., adapted PPO [33], FOCOPS [43], CPO [3], PCPO [39], and P3O [41]) show that e-COP performs the best or near-best on a range of Safety Gym [12] benchmark problems.

**Related Work.** A broad view of planning and model-free RL techniques for Constrained MDPs is provided in [27] and [18]. The development of SOTA policy optimization started with the TRPO algorithm [31], which used a trust region to update the current policy, and was further improved in PPO by use of proximal ideas [33]. This led to works like CPO [3], RCPO [36], and PCPO [39] for constrained RL problems in the infinite-horizon discounted setting. ACPO [4] extended CPO to the infinite-horizon average setting. These methods typically require inversion of a computationally-expensive Fischer information matrix at each update step, thus limiting scalability. Lagrangian-based algorithms [29, 30] showed that you could incorporate constraints but constrained satisfaction remained a concern. Algorithms like PDO [13] and RCPO [36] also use Lagrangian duality but to solve risk-constrained RL problems and suffer from computational overhead. Some other notable algorithms include IPO [26], P3O [41], APPO [15], etc. that use penalty terms, and hence do not suffer from computational overhead but have other drawbacks. For example, IPO assumes feasible intermediate iterations, which cannot be fulfilled in practice, P3O requires arbitrarily large penalty factors for feasibility which can lead to significant estimation errors. We note that all the above algorithms are for the infinite-horizon discounted (non-episodic) setting (except ACPO [4], which is for the average setting). We are not aware of any policy optimization algorithm for the episodic RL problem, with or without constraints.

## 2 Preliminaries

An episodic, or fixed horizon Markov decision process (MDP) is a tuple,  $m := (S, A, r, P, \mu, H)$ , where  $S$  is the set of states,  $A$  is the set of actions,  $r : S \times A \times S \rightarrow \mathbb{R}$  is the reward function,  $P : S \times A \times S \rightarrow [0, 1]$  is the transition probability function such that  $P(s'|s, a)$  is the probability of transitioning to state  $s'$  from state  $s$  by taking action  $a$ ,  $\mu : S \rightarrow [0, 1]$  is the initial state distribution, and  $H$  is the time horizon for each episode (characterized by a terminal state  $s_H$ ).

A policy  $\pi : S \rightarrow \Delta(A)$  is a mapping from states to probability distributions over the actions, with  $\pi(a|s)$  denoting the probability of selecting action  $a$  in state  $s$ , and  $\Delta(A)$  is the probability simplex over the action space  $A$ . However, due to the temporal nature of episodic RL, the optimal policies are generally not stationary, and we index the policy at time  $h$  by  $\pi_h$ , and denote  $\boldsymbol{\pi}_{1:H} = (\pi_h)_{h=1}^H$ . Then, the total undiscounted reward objective within an episode is defined as

$$J(\boldsymbol{\pi}_{1:H}) := \mathbb{E}_{\tau \sim \boldsymbol{\pi}_{1:H}} \left[ \sum_{h=1}^H r(s_h, a_h, s_{h+1}) \right]$$

where  $\tau$  refers to the sample trajectory  $(s_1, a_1, s_2, a_2, \dots, s_H)$  generated when following a policy sequence, i.e.,  $a_h \sim \pi_h(\cdot|s_h)$ ,  $s_{h+1} \sim P(\cdot|s_h, a_h)$ , and  $s_1 \sim \mu$ .

Let  $R_{h:H}(\tau)$  denote the total reward of a trajectory  $\tau$  starting from time  $h$  until episode terminal time  $H$  generated by following the policy sequence  $\boldsymbol{\pi}_{h:H}$ . We also define the state-value function of a state  $s$  at step  $h$  as  $V_h^\pi(s) := \mathbb{E}_{\tau \sim \pi} [R_{h:H}(\tau) | s_h = s]$  and the action-value function as  $Q_h^\pi(s, a) := \mathbb{E}_{\tau \sim \pi} [R_{h:H}(\tau) | s_h = s, a_h = a]$ . The advantage function is  $A_h^\pi(s, a) := Q_h^\pi(s, a) - V_h^\pi(s)$ . We also define  $\mathbb{P}_h^\pi(s | s_1) = \sum_{a \in A} \mathbb{P}_h^\pi(s, a | s_1)$ , where the term  $\mathbb{P}_h^\pi(s, a | s_1)$  is the probability of reaching  $(s, a)$  at time step  $h$  following  $\pi$  and starting from  $s_1$ .

**Constrained MDPs.** A constrained Markov decision process (CMDP) is an MDP augmented with constraints that restrict the set of allowable policies for that MDP. Specifically, we have  $m$  cost functions,  $C_1, \dots, C_m$  (with each function  $C_i : S \times A \times S \rightarrow \mathbb{R}$  mapping transition tuples to costs, similar to the reward function), and bounds  $d_1, \dots, d_m$ . And similar to the value function for the reward objective, we define the expected total cost objective for each cost function  $C_i$  (called cost value for the constraint) as

$$J_{C_i}(\boldsymbol{\pi}_{1:H}) := \mathbb{E}_{\tau \sim \boldsymbol{\pi}_{1:H}} \left[ \sum_{h=1}^H C_i(s_h, a_h, s_{h+1}) \right].$$

The goal then, in each episode, is to find a policy sequence  $\boldsymbol{\pi}_{1:H}^*$  such that

$$J(\boldsymbol{\pi}_{1:H}^*) := \max_{\boldsymbol{\pi}_{1:H} \in \Pi_C} J(\boldsymbol{\pi}_{1:H}), \text{ where } \Pi_C := \{\boldsymbol{\pi}_{1:H} \in \Pi : J_{C_i}(\boldsymbol{\pi}_{1:H}) \leq d_i, \forall i \in [1 : m]\} \quad (1)$$

is the set of feasible policies for a CMDP for some given class of policies  $\Pi$ . Lastly, analogous to  $V_h^\pi$ ,  $Q_h^\pi$ , and  $A_h^\pi$ , we can also define quantities for the cost functions  $C_i(\cdot)$  by replacing, and denote them by  $V_{C_i, h}^\pi$ ,  $Q_{C_i, h}^\pi$ , and  $A_{C_i, h}^\pi$ . Proofs of theorems and statements, if not given, are available in Appendix A.

**Notation.**  $[N]$  denotes  $\{1, \dots, N\}$  for some  $N \in \mathbb{N}$ .  $\pi_h$  refers to the policy at time step  $h$  within an episode. Denote  $\boldsymbol{\pi}_{s:t} := (\pi_s, \pi_{s+1}, \dots, \pi_t)$  for some  $s \leq t$  with  $s, t \in [H]$ . We shall only write  $\pi_{k, h}$  when it is necessary to differentiate policies from different episodes but at the same time  $h$ . It then naturally follows to define  $\boldsymbol{\pi}_{k, s:t}$  to be the sequence  $\boldsymbol{\pi}_{s:t}$  in episode  $k$ . We will denote  $\boldsymbol{\pi}_k \equiv \boldsymbol{\pi}_{k, 1:H}$ , and where not needed drop the index for the episode so that  $\boldsymbol{\pi} \equiv \boldsymbol{\pi}_k$ .

### 3 Episodic Constrained Optimization of Policies (e-COP)

In this section, we propose a constrained policy optimization algorithm for episodic MDPs. Policy optimization algorithms for MDPs have proven remarkably effective given their ability to computationally scale up to high dimensional continuous state and action spaces [31–33]. Such algorithms have also been proposed for infinite-horizon constrained MDPs with discounted criterion [1] as well as the average criterion [4] but not for the finite horizon (or as it is often called, the episodic) setting.

We note that finite horizon is not simply a special case of infinite-horizon discounted setting since the reward/cost functions in the former can be time-varying while the latter only allows for time-invariant objectives. Furthermore, even with time-invariant objectives, the optimal policy is time-dependent, while for the latter setting there an optimal policy that is stationary exists.

**A Policy Difference Lemma for Episodic MDPs.** Most policy optimization RL algorithms are based on a value or policy difference technical lemma [23]. Unfortunately, the policy difference lemmas that have been derived previously for the infinite-horizon discounted [31] and average case [4] are not applicable here and hence, we derive a new policy difference lemma for the episodic setting.

**Lemma 3.1.** For an episode of length  $H$  and two policies,  $\pi$  and  $\pi'$ , the difference in their performance assuming identical initial state distribution  $\mu$  (i.e.,  $s_1 \sim \mu$ ) is given by

$$J(\pi) - J(\pi') = \sum_{h=1}^H \mathbb{E}_{\substack{s_h, a_h \sim \mathbb{P}_h^\pi(\cdot, \cdot | s) \\ s_1 \sim \mu}} [A_h^{\pi'}(s_h, a_h)]. \quad (2)$$

The proof can be found in Appendix A.1. A key difference to note between the above and similar results for infinite-horizon settings [31, 4] is that considering stationary policies (and hence corresponding occupation measures) is not enough for the episodic setting since, in general, such a policy may be far from optimal. This explains why Lemma 3.1 looks so different (e.g., see (2) in [31], and Lemma 3.1 in [4]). Indeed, the lemma above indicates that policy updates do not have to recurse backwards from the terminal time as dynamic programming algorithms do for episodic settings, which is somewhat surprising.

**A Constrained Policy Optimization Algorithm for Episodic MDPs.** Iterative policy optimization algorithms achieve state of the art performance [33, 36, 39] on RL problems. Most such algorithms maximize the advantage function based on a suitable policy difference lemma, solving an unconstrained RL problem. Some additionally ensure satisfaction of infinite horizon expectation constraints [3, 4]. However, given that our policy lemma for the episodic setting (Lemma 3.1) is significantly different, we need to re-design the algorithm based on it. A first attempt is presented as Algorithm 1, where each iteration corresponds to an update with a full horizon  $H$  episode.

---

**Algorithm 1** Iterative Policy Optimization for Constrained Episodic (IPOCE) RL

---

- 1: **Input:** Initial policy  $\pi_0$ , number of episodes  $K$ , episode horizon  $H$ .
- 2: **for**  $k = 1, 2, \dots, K$  **do**
- 3:   Run  $\pi_{k-1}$  to collect trajectories  $\tau$ .
- 4:   Evaluate  $A_h^{\pi_{k-1}}$  and  $A_{C_i, h}^{\pi_{k-1}}$  for  $h \in [H]$  from  $\tau$ .
- 5:   **for**  $t = H, H-1, \dots, 1$  **do**

$$6: \quad \pi_{k,t}^* = \arg \min_{\pi_{k,t}} \sum_{h=t}^H \mathbb{E}_{\substack{s \sim \rho_{\pi_{k,h}} \\ a \sim \pi_{k,h}}} [-A_h^{\pi_{k-1}}(s, a)] \quad \text{s.t.} \quad J_{C_i}(\pi_{k-1}) + \sum_{h=t}^H \mathbb{E}_{\substack{s \sim \rho_{\pi_{k,h}} \\ a \sim \pi_{k,h}}} [A_{C_i, h}^{\pi_{k-1}}(s, a)] \leq d_i, \quad \forall i \quad (3)$$

- 7:   **end for**
  - 8:   Set  $\pi_k \leftarrow (\pi_{k,1}^*, \pi_{k,2}^*, \dots, \pi_{k,H}^*)$ .
  - 9: **end for**
- 

The iterative constrained policy optimization algorithm introduced above uses the current iterate of the policy  $\pi_k$  to collect a trajectory  $\tau$ , and use them to evaluate  $A_h^{\pi_{k-1}}$  and  $A_{C_i, h}^{\pi_{k-1}}$  for  $h \in [H]$ . At the end of the episode, we solve  $H$  optimization problems (one for each  $h \in [H]$ ) that result in a new sequence of policies  $\pi$ . As is natural in episodic problems, we do backward iteration in time, i.e., solve the problem in step (6) at  $h = H$ , and then go backwards towards  $h = 1$ .

Note that the expectation of advantage functions in equation (3) is with respect to the policy  $\pi$  (the optimization variable) and its corresponding *time-dependent* state occupation distribution  $\rho_{\pi_h}$ . In the infinite-horizon settings, the expectation is with respect to the steady state stationary distribution, but that does not exist in the episodic setting.

**Using current policy for action selection.** Algorithm 1 represents an exact principled solution to the constrained episodic MDP, but the intractable optimization performed in (3) makes it impractical (as in the case of infinite horizon policy optimization algorithms [31, 3, 39]). We proceed to introduce a sequence of ideas that make the algorithm practical (e.g., by avoiding computationally expensive Hessian matrix inversion for use with trust-region methods [33, 14, 41]). However, getting rid of trust regions leads to large updates on policies, but PPO [33] and P3O [41] successfully overcome this problem by clipping the advantage function and adding a ReLU-like penalty term to the minimization objective. Motivated by this, we rewrite the optimization problem in (3) as follows by parameterizing the policy  $\pi_{k,t}$  in episode  $k$  and time step  $t$  by  $\theta_{k,t}$ :

$$\pi_{k,t} = \arg \min_{\pi_{k,t}} \sum_{h=t}^H \mathbb{E}_{\substack{s \sim \rho_{\pi_{k,h}} \\ a \sim \pi_{k-1,h}}} [-\rho(\theta_h) A_h^{\pi_{k-1}}(s, a)] + \sum_i \lambda_{t,i} \max \left\{ 0, \sum_{h=t}^H \mathbb{E}_{\substack{s \sim \rho_{\pi_{k,h}} \\ a \sim \pi_{k-1,h}}} [\rho(\theta_h) A_{C_i, h}^{\pi_{k-1}}(s, a)] + J_{C_i}(\pi_{k-1}) - d_i \right\}, \quad (4)$$

where  $\rho(\theta_h) = \frac{\pi_{\theta_{k,h}}}{\pi_{\theta_{k-1,h}}}$  is the importance sampling ratio,  $\lambda_{t,i}$  is a penalty factor for constraint  $C_i$ , and  $\pi_{k,\theta_h} \equiv \pi_{k,h} \equiv \theta_{k,h}$ . Note that the ReLU-like penalty term above is different from the traditional first-order and second-order gradient approximations that are employed in trust-region methods [3, 39]. In essence, the penalty is applied when the agent breaches the associated constraint, while the objective remains consistent with standard policy optimization when all constraints are satisfied.

**Introducing quadratic damping penalty.** It has been noted in such iterative policy optimization algorithms that the behaviour of the learning agent when it nears the constraint threshold is quite volatile during training [3, 39, 42]. This is because the penalty term is active only when the constraints are violated which results in sharp behavior change for the agent. To alleviate this problem, we introduce an additional quadratic damping term to the objective above, which provides stable cost control to compliment the lagged Lagrangian multipliers. This has proved effective in physics-based control applications [16, 25, 20] resulting in improved convergence since the damping term provides stability, while keeping the solution set the same as for the original Problem (3) and the adapted Problem (4) (as we prove later).

For brevity, we denote the constraint term in Problem (4) as

$$\Psi_{C_i,t}(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k) := \sum_{h=t}^H \mathbb{E}_{\substack{s \sim \rho_{\pi_{k,h}} \\ a \sim \pi_{k-1,h}}} [\rho(\theta_h) A_{C_i,h}^{\boldsymbol{\pi}_{k-1}}(s, a)] + J_{C_i}(\boldsymbol{\pi}_{k-1}) - d_i.$$

Now introduce a slack variable  $x_{t,i} \geq 0$  for each constraint to convert the inequality constraint ( $\Psi_{C_i,t}(\cdot, \cdot) \leq 0$ ) to equality by letting

$$w_{t,i}(\boldsymbol{\pi}_k) := \Psi_{C_i,t}(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k) + x_{t,i} = 0.$$

With this notation, we restate Problem (4) as:

$$\pi_{k,t}^* = \min_{\boldsymbol{\pi}_{k,t}} \mathcal{L}_t(\boldsymbol{\pi}_k, \boldsymbol{\lambda}) := \sum_{h=t}^H \mathbb{E}_{\substack{s \sim \rho_{\pi_{k,h}} \\ a \sim \pi_{k-1,h}}} [-\rho(\theta_h) A_h^{\boldsymbol{\pi}_{k-1}}(s, a)] + \sum_i^m \lambda_{t,i} \max\{0, \Psi_{C_i,t}(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k)\}.$$

Now we introduce the quadratic damping term and the intermediate loss function then takes the form,

$$\mathcal{L}_t(\boldsymbol{\pi}_k, \boldsymbol{\lambda}, \boldsymbol{x}, \beta) := \sum_{h=t}^H \mathbb{E}_{\substack{s \sim \rho_{\pi_{k,h}} \\ a \sim \pi_{k-1,h}}} [-\rho(\theta_h) A_h^{\boldsymbol{\pi}_{k-1}}(s, a)] + \sum_i^m \lambda_{t,i} w_{t,i}(\boldsymbol{\pi}_k) + \frac{\beta}{2} \sum_i^m w_{t,i}^2(\boldsymbol{\pi}_k) \quad (5)$$

$$\text{Then, } (\pi_{k,t}^*, \boldsymbol{\lambda}_t^*, \boldsymbol{x}_t^*) = \max_{\boldsymbol{\lambda} \geq 0} \min_{\boldsymbol{\pi}_{k,t}, \boldsymbol{x}} \mathcal{L}_t(\boldsymbol{\pi}_k, \boldsymbol{\lambda}, \boldsymbol{x}, \beta),$$

where  $\beta$  is the damping factor,  $\boldsymbol{\lambda}_t = (\lambda_{t,i})_{i=1}^m$ , and  $\boldsymbol{x}_t = (x_{t,i})_{i=1}^m$ . We can then construct a primal-dual solution to the max-min optimization problem. The need for a slack variable  $\boldsymbol{x}$  can be obviated by setting the partial derivative of  $\mathcal{L}_t(\cdot)$  with respect to  $\boldsymbol{x}$  equal to 0. This leads to a ReLU-like solution:  $x_{t,i}^* = \max\left(0, -\Psi_{C_i,t}(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k) - \frac{\lambda_{t,i}}{\beta}\right)$ . The intermediate problem then takes the form as below.

**Proposition 3.2.** *The inner optimization problem in (5) with respect to  $\boldsymbol{x}$  is a convex quadratic program with non-negative constraints, which can be solved to yield the following intermediate problem:*

$$(\pi_{k,t}^*, \boldsymbol{\lambda}_t^*) = \max_{\boldsymbol{\lambda} \geq 0} \min_{\boldsymbol{\pi}_{k,t}} \mathcal{L}_t(\boldsymbol{\pi}_k, \boldsymbol{\lambda}, \beta), \quad \text{where} \\ \mathcal{L}_t(\boldsymbol{\pi}_k, \boldsymbol{\lambda}, \beta) = \sum_{h=t}^H \mathbb{E}_{\substack{s \sim \rho_{\pi_{k,h}} \\ a \sim \pi_{k-1,h}}} [-\rho(\theta_h) A_h^{\boldsymbol{\pi}_{k-1}}(s, a)] + \frac{\beta}{2} \sum_i^m \left( \max\left\{0, \Psi_{C_i,t}(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k) + \frac{\lambda_{t,i}}{\beta}\right\}^2 - \frac{\lambda_{t,i}^2}{\beta^2} \right). \quad (6)$$

The proof can be found in Appendix A.1. One can see that the cost penalty is active when  $\Psi_{C_i,t}(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k) \geq -\frac{\lambda_{t,i}}{\beta}$  rather than when  $\Psi_{C_i,t}(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k) \geq 0$ . This allows the agent to act in a constrained manner even before the constraint is violated. Further, as we show next, the introduction of the damping factor and the RELU-like penalty does not change the solution of the problem (under some suitable assumptions):

**Theorem 3.3.** *Let  $\pi^{(3)*}$  be a solution to Problem (3), and let  $(\pi^{(6)*}, \boldsymbol{\lambda}^{(6)*})$  be a solution to Problem (6). Then, for sufficiently large  $\beta > \bar{\beta}$  and  $\lambda_{t,i} > \bar{\lambda} \forall i$ ,  $\pi^{(3)*}$  is a solution to Problem (6), and  $\pi^{(6)*}$  is a solution to Problem (3).*

We refer the reader to Appendix A.1 for the proof. This theorem implies that we can search for the optimal feasible policies of the CMDP Problem (1) by iteratively solving Problem (6). Next, we make some further modifications to Problem (6) that give us our final tractable algorithm.

**Removing Lagrange multiplier dependency.** Problem (6) requires a primal-dual algorithm that will iteratively solve over the policies and the dual variable  $\lambda$ . But from the Lagrangian, we can actually take a derivative with respect to  $\lambda$ , and then solve for it, which yields the following update rule for it:

$$\lambda_{t,i}^{(k)} = \max \left( 0, \lambda_{t,i}^{(k-1)} + \beta \Psi_{C_i,t}(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_{k-1}) \right). \quad (7)$$

This update rule simplifies the optimization problem and updates the Lagrange multipliers in the  $k$ th episode based on the constraint violation in the  $(k-1)$ th episode.

**Clipping the advantage functions.** Solving the optimization problem presented in equation (6) is intractable since we do not know  $\rho_\pi$  beforehand. Hence, we replace  $\rho_\pi$  by the empirical distribution observed with the policy of the previous episode,  $\boldsymbol{\pi}_{k-1}$ , i.e.,  $\rho_{\pi_{k,h}} \approx \rho_{\pi_{k-1,h}} \forall h$ . Similar to [33] for PPO, we also use *clipped* surrogate objective functions for both the reward and cost advantage functions. Thus, the final problem combining equation (4) and equation (6) can be constructed as follows.

If we let

$$\begin{aligned} \mathcal{L}_t(\theta) &= \sum_{h=t}^H \mathbb{E}_{\substack{s \sim \rho_{\pi_{k-1,h}} \\ a \sim \pi_{k-1,h}}} \left[ -\min \left\{ \rho(\theta_h) A_h^{\pi_{k-1}}(s, a), \text{clip}(\rho(\theta_h), 1 - \epsilon, 1 + \epsilon) A_h^{\pi_{k-1}}(s, a) \right\} \right] \quad \text{and,} \\ \mathcal{L}_{C_i,t}(\theta) &= \sum_{h=t}^H \mathbb{E}_{\substack{s \sim \rho_{\pi_{k-1,h}} \\ a \sim \pi_{k-1,h}}} \left[ -\min \left\{ \rho(\theta_h) A_{C_i,h}^{\pi_{k-1}}(s, a), \text{clip}(\rho(\theta_h), 1 - \epsilon, 1 + \epsilon) A_{C_i,h}^{\pi_{k-1}}(s, a) \right\} \right] \end{aligned}$$

then, the final loss function  $\tilde{\mathcal{L}}_t(\cdot)$  of the final problem is:

$$\begin{aligned} \pi_{k,t}^* = \arg \min_{\pi_{k,t}} \tilde{\mathcal{L}}_t(\pi_\theta, \boldsymbol{\lambda}, \beta) &:= \arg \min_{\pi_{k,t}} \mathcal{L}_t(\theta) + \sum_i^m \lambda_{t,i} \max \left\{ 0, \mathcal{L}_{C_i,t}(\theta) + (J_{C_i}(\boldsymbol{\pi}_{k-1}) - d_i) \right\} \\ &+ \frac{\beta}{2} \sum_i^m \left( \max \left\{ 0, \mathcal{L}_{C_i,t}(\theta) + (J_{C_i}(\boldsymbol{\pi}_{k-1}) - d_i) + \frac{\lambda_{t,i}}{\beta} \right\}^2 - \frac{\lambda_{t,i}^2}{\beta^2} \right) \quad (8) \end{aligned}$$

Usually for experiments, Gaussian policies with means and variances predicted from neural networks are used [31, 3, 33, 39]. We employ the same approach and since we work in the finite horizon setting, the reward and constraint advantage functions can easily be calculated from any trajectory  $\tau \sim \boldsymbol{\pi}$ . The surrogate problem in equation (8) then includes the pessimistic bounds on Problem (6), which is unclipped.

**Adaptive parameter selection.** The value of  $\beta$  is required to be larger than the unknown  $\bar{\beta}$  according to Theorem 3.3, but we also know that too large a  $\beta$  decays the performance (as seen in harmonic oscillator kinetic energy formulations [25, 20, 40]). To manage this tradeoff, we provide an instance-dependent adaptive way to adjust the damping factor as a hyperparameter. In each episode  $k$ , we update the damping parameter whenever a secondary constraint cost value denoted by  $\mathcal{C}(\boldsymbol{\pi}_k)$  is larger than some threshold  $c_k$ . Using Proposition 3.2, we provide the following definitions.

$$\mathcal{C}(\boldsymbol{\pi}_k) := \sum_{t=1}^H \sum_i^m \max \left\{ J_{C_i}(\boldsymbol{\pi}_k) - d_i, -\frac{\lambda_{t,i}^{(k)}}{\beta} \right\} \quad \text{and} \quad c_k := \frac{\sqrt{m}}{\beta} \cdot \max_{t \in [H]} \|\boldsymbol{\lambda}_t^{(k)}\|_\infty$$

---

**Algorithm 2** Episodic Constrained Optimization of Policies (e-COP)

---

- 1: **Input:** Initial policy  $\theta_0 := \boldsymbol{\pi}_0 := \boldsymbol{\pi}_{\theta_0}$ , critic networks  $V^{\phi_0}$  and  $V_{C_i}^{\psi_0} \forall i$ , penalty factor  $\beta$ , number of episodes  $K$ , episode horizon  $H$ , learning rate  $\alpha$ , penalty update factor  $\rho$ .
  - 2: **for**  $k = 1, 2, \dots, K$  **do**
  - 3:   Collect a set of trajectories  $\mathcal{D}_{k-1}$  with policy  $\boldsymbol{\pi}_{k-1}$  and update the critic network.
  - 4:   Get updated  $\lambda^{(k)}$  using equation (7).
  - 5:   **for**  $t = H, H-1, \dots, 1$  **do**
  - 6:     Update the policy  $\theta_{k,t} \leftarrow \theta_{k,t+1} - \alpha \nabla_{\theta} \tilde{\mathcal{L}}_t(\theta, \lambda^{(k)}, \beta)$  using equation (8).
  - 7:   **end for**
  - 8:   **if**  $\mathcal{C}(\boldsymbol{\pi}_k) \geq c_k$  **then**
  - 9:      $\beta = \min(\beta_{\max}, \rho\beta)$
  - 10:   **end if**
  - 11: **end for**
-






Figure 1: The Humanoid, Circle, Reach, Grid, Bottleneck, and Navigation tasks. See Appendix A.2.1 for details.

Hence, we increase  $\beta$  by a constant factor  $\rho > 1$  after every episode if  $\mathcal{C}(\pi_k) \geq c_k$  until a stopping condition is fulfilled, typically a constant  $\beta_{\max}$ . This leads to constraint-satisfying iterations that are more stable, and we show that it enables a fixed  $\beta$  to generalize well across various tasks. The initial  $\beta$  can simply be selected by a quantified line-search to obtain a feasible  $\beta > \bar{\beta}$  [3, 4].

We note that the final loss function in equation (8) is differentiable almost everywhere, so we could easily solve it via any first-order optimizer [24]. The final practical algorithm, e-COP is given in Algorithm 2.

## 4 Experiments

We conducted extensive experimental evaluation on the relative empirical performance of the e-COP algorithm to arrive at the following conclusions: (i) The e-COP algorithm performs better or nearly as well as all baseline algorithms for infinite-horizon discounted safe RL tasks in maximizing episodic return while satisfying given constraints. (ii) It is more robust to stochastic and complex environments [30], even where previous methods struggle. (iii) It has stable behavior and more accurate cost control as compared to other baselines near the constraint threshold due to the damping term.

**Environments.** For a comprehensive empirical evaluation, we selected eight scenarios from well-known safe RL benchmark environments - Safe MuJoCo [43] and Safety Gym [30], as well as MuJoCo environments. These include: Humanoid, PointCircle, AntCircle, PointReach, AntReach, Grid, Bottleneck, and Navigation. See Figure 1 for an overview of the tasks and scenarios. Note that Navigation is a multi-constraint task and for the Reach environment, we set the reward as a function of the Euclidean distance between agent’s position and goal position. In addition, we make it impossible for the agent to reach the goal before the end of the episode. For more information see Appendix A.2.1.

**Baselines.** We compare our e-COP algorithm with the following baseline algorithms: CPO [3], PCPO [39], FOCOPS [43], PPO with Lagrangian relaxation [33, 35], and penalty-based P3O [41]. Since the above state-of-the-art baseline algorithms are already well understood to outperform other algorithms such as PDO [13], IPO [26], and CPPO-PID [35] in prior benchmarking studies, we do not compare against them. Moreover, since PPO does not originally incorporate constraints, for fair comparison, we introduce constraints using a Lagrangian relaxation (called PPO-L). In addition, for each algorithm, we report its performance with the discount factor that achieves the best performance. See Appendix A.3.1 for more details.

**Evaluation Details and Protocol.** For the Circle task, we use a point-mass with  $S \subseteq \mathbb{R}^9$ ,  $A \subseteq \mathbb{R}^2$  and for the Reach task, an ant robot with  $S \subseteq \mathbb{R}^{16}$ ,  $A \subseteq \mathbb{R}^8$ . The Grid task has  $S \subseteq \mathbb{R}^{56}$ ,  $A \subseteq \mathbb{R}^4$ . We use two hidden layer neural networks to represent Gaussian policies for the tasks. For Circle and Reach, size is (32,32) for both layers, and for Grid and Navigation the layer sizes are (16,16) and (25,25). We set the step size  $\delta$  to  $10^{-4}$ , and for each task, we conduct 5 independent runs of  $K = 500$  episodes each of horizon  $H = 200$ . Since there are two objectives (rewards in the objective and costs in the constraints), we show the plots which maximize the reward objective while satisfying the cost constraint.

### 4.1 Performance Analysis


Table 1 lists the numerical performance of all tested algorithms in seven single constraint scenarios, and one multiple constraint scenario. We find that overall, the e-COP algorithm in most cases outperforms (green) all other baseline algorithms in finding the optimal policy while satisfying the constraints, and in other cases comes a close second (light green).

From Figure 2, we can see how the e-COP algorithm is able to improve the reward objective over the baselines while having approximate constraint satisfaction. We also see that updates of e-COP are faster and smoother than other baselines due to the added damping penalty, which ensures smoother

Task	e-COP	FOCOPS [43]	PPO-L [30]	PCPO [39]	P3O [41]	CPO [3]	APPO [15]	IPO [26]	
Humanoid	R	1652.5 ± 13.4	<b>1734.1 ± 27.4</b>	1431.2 ± 25.2	1602.3 ± 10.1	1669.4 ± 13.7	1465.1 ± 55.3	1488.2 ± 29.3	1578.6 ± 25.2
	C (20.0)	17.3 ± 0.3	19.7 ± 0.6	18.8 ± 1.5	16.3 ± 1.4	20.1 ± 3.3	18.5 ± 2.9	20.0 ± 1.3	19.1 ± 2.5
PointCircle	R	<b>110.5 ± 9.3</b>	81.6 ± 8.4	57.2 ± 9.2	68.2 ± 9.1	89.1 ± 7.1	65.3 ± 5.3	91.2 ± 9.6	68.7 ± 15.2
	C (10.0)	9.8 ± 0.9	10.0 ± 0.4	9.8 ± 0.5	9.9 ± 0.4	9.9 ± 0.3	9.5 ± 0.9	10.2 ± 0.6	9.3 ± 0.5
AntCircle	R	<b>198.6 ± 7.4</b>	161.9 ± 22.2	134.4 ± 10.3	168.3 ± 13.3	182.6 ± 18.7	127.1 ± 12.1	155.5 ± 19.4	149.3 ± 33.6
	C (10.0)	9.8 ± 0.6	9.9 ± 0.5	9.6 ± 1.6	9.5 ± 0.6	9.8 ± 0.2	10.1 ± 0.7	10.0 ± 0.5	9.5 ± 1.0
PointReach	R	<b>81.5 ± 10.2</b>	65.1 ± 9.6	46.1 ± 14.8	73.2 ± 7.4	76.3 ± 6.4	89.2 ± 8.1	74.3 ± 6.7	49.1 ± 10.6
	C (25.0)	24.5 ± 6.1	24.8 ± 7.6	25.1 ± 6.1	24.9 ± 5.6	26.3 ± 6.9	33.3 ± 10.7	26.3 ± 8.1	24.7 ± 11.3
AntReach	R	70.8 ± 14.6	48.3 ± 5.6	54.2 ± 9.5	39.4 ± 5.3	<b>73.6 ± 5.1</b>	102.3 ± 7.1	61.5 ± 10.4	45.2 ± 13.3
	C (25.0)	24.2 ± 8.4	25.1 ± 11.9	21.9 ± 10.7	27.9 ± 12.2	24.8 ± 7.3	35.1 ± 10.9	24.5 ± 6.4	24.9 ± 9.2
Grid	R	258.1 ± 33.1	215.4 ± 45.6	<b>276.3 ± 57.9</b>	226.5 ± 29.2	201.5 ± 39.2	178.1 ± 23.8	184.4 ± 21.5	229.4 ± 32.8
	C (75.0)	71.3 ± 26.9	76.6 ± 29.8	71.8 ± 25.1	72.6 ± 16.5	79.3 ± 19.3	69.3 ± 19.8	79.5 ± 35.8	74.2 ± 24.6
Bottleneck	R	<b>345.1 ± 52.6</b>	251.3 ± 59.1	298.3 ± 71.2	264.2 ± 43.8	291.1 ± 26.7	388.1 ± 36.6	220.1 ± 30.1	279.3 ± 43.8
	C (50.0)	49.7 ± 15.1	46.6 ± 19.8	41.4 ± 17.6	49.8 ± 10.5	45.3 ± 8.2	54.3 ± 13.5	47.4 ± 12.3	48.2 ± 14.6
Navigation	R	<b>217.6 ± 11.5</b>		175.1 ± 3.7		153.5 ± 25.2		135.7 ± 19.2	164.1 ± 12.8
	C1 (10.0)	9.6 ± 1.5	n/a	9.9 ± 1.9	n/a	9.9 ± 1.7	n/a	9.9 ± 2.1	10.0 ± 0.5
	C2 (25.0)	23.7 ± 4.1		22.3 ± 2.1		24.5 ± 4.1		23.9 ± 3.8	24.6 ± 3.1

Table 1: Mean performance with normal 95% confidence interval over 5 independent runs on some tasks.

### Episodic Rewards:



### Constraint Costs:

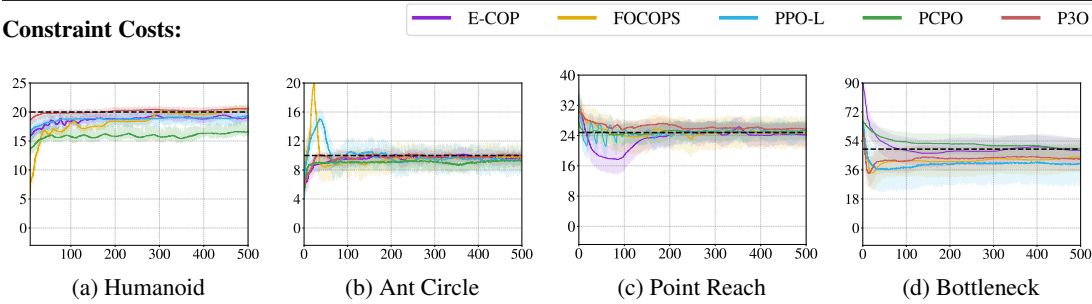


Figure 2: The cumulative episodic reward and constraint cost function values vs episode learning curves for some algorithm-task pairs. Solid lines in each figure are the means, while the shaded area represents 1 standard deviation, all over 5 runs. The dashed line in constraint plots is the constraint threshold.

convergence with only a few constraint-violating behaviors during training. In particular, e-COP is the *only* algorithm that best learns almost-constraint-satisfying maximum reward policies across *all* tasks: in simple Humanoid and Circle environments, e-COP is able to almost exactly track the cost constraint values to within the given threshold. However, for the high dimensional Grid environment we have more constraint violations due to complexity of the policy behavior, leading to higher variance in episodic rewards as compared to other environments. Regardless, overall in these environments, e-COP still outperforms *all* other baselines with the least episodic constraint violation. For the multiple constraint Navigation environment, see Figure 3.

## 4.2 Secondary Evaluation

In this section, we take a deeper dive into the empirical performance of e-COP. We discuss its dependence on various factors, and try to verify its merits.

**Generalizability.** From the discussion above, it’s clear that e-COP demonstrates accurate safety satisfaction across tasks of varying difficulty levels. From Figure 4, we further see that e-COP satisfies the






Figure 3: Navigation environment with multiple constraints: Episodic Rewards (left), Cost1 (center, for hazards) and Cost2 (right, for pillars) of e-COP. The dashed line in the cost plots is the cost threshold (10 for Cost1 and 25 for Cost2). C1/C2 constrained means only taking Cost1/Cost2 into the e-COP loss function and ignoring the other one.

constraints in all cases and precisely converges to the specified cost limit. Furthermore, the fluctuation observed in the baseline Lagrangian-based algorithms is shown not to be tied to a specific cost limit.

We also conducted a set of experiments wherein we study how e-COP effectively adapts to different cost thresholds. For this, we use the hyperparameters of a pre-trained e-COP agent, which is trained with a particular cost threshold in an environment, for learning on different cost thresholds within the same environment. Figure 6 in Appendix A.3.4 illustrates the training curves of these pre-trained agents, and we see that while e-COP can generalize well across different cost thresholds, other baseline algorithms may require further tuning to accommodate different constraint thresholds.




Figure 4: Cumulative episodic rewards and costs of baselines in two environments with two constraint thresholds.

Task		e-COP				P3O [41]
		$\beta = 5$ , fixed	$\beta = 5$ , adaptive	$\beta = 10$ , fixed	$\beta = 10$ , adaptive	
PointCircle	R	150.5 $\pm$ 11.1	<b>168.6 <math>\pm</math> 14.3</b>	145.2 $\pm$ 12.2	165.3 $\pm$ 11.4	162.4 $\pm$ 14.7
	C (20.0)	17.3 $\pm$ 1.3	<b>19.7 <math>\pm</math> 0.6</b>	18.8 $\pm$ 1.5	18.3 $\pm$ 1.4	19.1 $\pm$ 3.3
AntReach	R	48.2 $\pm$ 3.5	58.6 $\pm$ 5.1	53.2 $\pm$ 5.3	<b>65.2 <math>\pm</math> 8.1</b>	61.1 $\pm$ 5.6
	C (20.0)	19.8 $\pm$ 5.9	20.0 $\pm$ 4.4	20.6 $\pm$ 4.5	<b>19.2 <math>\pm</math> 6.2</b>	18.9 $\pm$ 7.3

Table 2: Performance of e-COP for different  $\beta$  settings on two tasks. Values are given with normal 95% confidence interval over 5 independent runs.

**Sensitivity.** The effectiveness and performance of e-COP would not be justified if it was not robust to the damping hyperparameter  $\beta$ , which varies across tasks depending on the values of  $\mathcal{C}(\cdot)$  and  $c_h$ . Since this damping penalty enables e-COP to have stable continuous cost control, we update it adaptively as described in Algorithm 2. As seen in Table 2, damping penalty indeed stabilizes the training process and helps in converging to an optimal safe policy.

## 5 Conclusion

In this paper, we have introduced an easy to implement, scalable policy optimization algorithm for episodic RL problems with constraints due to safety or other considerations. It is based on a policy difference lemma for the episodic setting, which surprisingly has quite a different form than the ones for infinite-horizon discounted or average settings. This provides the theoretical foundation for the algorithm, which is designed by incorporating several time-tested, practical as well as novel ideas. Policy optimization algorithms for Constrained MDPs tend to be numerical unstable and non-scalable due to the need for inverting the Fisher information matrix. We sidestep both of these issues by introducing a quadratic damping penalty term that works remarkably well. The algorithm development is well supported by theory, as well as with extensive empirical analysis on a range of Safety Gym and Safe MuJoCo benchmark environments against a suite of baseline algorithms adapted from their non-episodic roots.

## References

- [1] J. Achiam. UC Berkeley CS 285 (Fall 2017), Advanced Policy Gradients, 2017. URL: [http://rail.eecs.berkeley.edu/deeprlcourse-fa17/f17docs/lecture\\_13\\_advanced\\_pg.pdf](http://rail.eecs.berkeley.edu/deeprlcourse-fa17/f17docs/lecture_13_advanced_pg.pdf). Last visited on 2020/05/24.
- [2] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- [3] J. Achiam, D. Held, A. Tamar, and P. Abbeel. Constrained policy optimization. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 22–31. JMLR. org, 2017.
- [4] A. Agnihotri, R. Jain, and H. Luo. ACPO: A policy optimization algorithm for average MDPs with constraints. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 397–415. PMLR, 21–27 Jul 2024.
- [5] A. Agnihotri, R. Jain, D. Ramachandran, and Z. Wen. Online bandit learning with offline preference data. *arXiv preprint arXiv:2406.09574*, 2024.
- [6] A. Agnihotri, P. Saraf, and K. R. Bapnad. A convolutional neural network approach towards self-driving cars. In *2019 IEEE 16th India Council International Conference (INDICON)*, pages 1–4, 2019.
- [7] I. Akkaya, M. Andrychowicz, M. Chociej, M. Litwin, B. McGrew, A. Petron, A. Paino, M. Plappert, G. Powell, R. Ribas, et al. Solving rubik’s cube with a robot hand. *arXiv preprint arXiv:1910.07113*, 2019.
- [8] D. P. Bertsekas. *Constrained optimization and Lagrange multiplier methods*. Academic press, 2014.
- [9] E. G. Birgin and J. M. Martínez. *Practical augmented Lagrangian methods for constrained optimization*. SIAM, 2014.
- [10] K. Black, M. Janner, Y. Du, I. Kostrikov, and S. Levine. Training diffusion models with reinforcement learning. *arXiv preprint arXiv:2305.13301*, 2023.
- [11] H. Bojun. Steady state analysis of episodic reinforcement learning. *Advances in Neural Information Processing Systems*, 33:9335–9345, 2020.
- [12] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba. Openai gym, 2016.
- [13] Y. Chow, M. Ghavamzadeh, L. Janson, and M. Pavone. Risk-constrained reinforcement learning with percentile risk criteria. *The Journal of Machine Learning Research*, 18(1):6070–6120, 2017.
- [14] Y. Chow, O. Nachum, A. Faust, E. Duenez-Guzman, and M. Ghavamzadeh. Lyapunov-based safe policy optimization for continuous control. *arXiv preprint arXiv:1901.10031*, 2019.
- [15] J. Dai, J. Ji, L. Yang, Q. Zheng, and G. Pan. Augmented proximal policy optimization for safe reinforcement learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37:7288–7295, Jun. 2023.
- [16] A. P. Dowling and A. S. Morgans. Feedback control of combustion oscillations. *Annu. Rev. Fluid Mech.*, 37:151–182, 2005.
- [17] I. Greenberg and S. Mannor. Detecting rewards deterioration in episodic reinforcement learning. In *International Conference on Machine Learning*, pages 3842–3853. PMLR, 2021.
- [18] S. Gu, L. Yang, Y. Du, G. Chen, F. Walter, J. Wang, Y. Yang, and A. Knoll. A review of safe reinforcement learning: Methods, theory and applications. *arXiv preprint arXiv:2205.10330*, 2022.
- [19] H. Hu, Z. Liu, S. Chitlangia, A. Agnihotri, and D. Zhao. Investigating the impact of multi-lidar placement on object detection for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2550–2559, June 2022.
- [20] I. Imayoshi, A. Isomura, Y. Harima, K. Kawaguchi, H. Kori, H. Miyachi, T. Fujiwara, F. Ishidate, and R. Kageyama. Oscillatory control of factors determining multipotency and fate in mouse neural progenitors. *Science*, 342(6163):1203–1208, 2013.
- [21] M. Janner, Y. Du, J. B. Tenenbaum, and S. Levine. Planning with diffusion for flexible behavior synthesis. *arXiv preprint arXiv:2205.09991*, 2022.
- [22] J. Ji, J. Zhou, B. Zhang, J. Dai, X. Pan, R. Sun, W. Huang, Y. Geng, M. Liu, and Y. Yang. Omnisafe: An infrastructure for accelerating safe reinforcement learning research. *arXiv preprint arXiv:2305.09304*, 2023.

- [23] S. Kakade and J. Langford. Approximately optimal approximate reinforcement learning. In *International Conference on Machine Learning*, volume 2, pages 267–274, 2002.
- [24] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [25] W. Klimesch. Alpha-band oscillations, attention, and controlled access to stored information. *Trends in cognitive sciences*, 16(12):606–617, 2012.
- [26] Y. Liu, J. Ding, and X. Liu. Ipo: Interior-point policy optimization under constraints. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 4940–4947, 2020.
- [27] Y. Liu, A. Halev, and X. Liu. Policy learning with constraints in model-free reinforcement learning: A survey. In *The 30th international joint conference on artificial intelligence (ijcai)*, 2021.
- [28] G. Neu and C. Pike-Burke. A unifying view of optimism in episodic reinforcement learning. *Advances in Neural Information Processing Systems*, 33:1392–1403, 2020.
- [29] A. Ray, J. Achiam, and D. Amodei. Benchmarking Safe Exploration in Deep Reinforcement Learning. *arXiv preprint arXiv:1910.01708*, 2019.
- [30] A. Ray, J. Achiam, and D. Amodei. Benchmarking safe exploration in deep reinforcement learning. *arXiv preprint arXiv:1910.01708*, 7, 2019.
- [31] J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz. Trust region policy optimization. In *International Conference on Machine Learning*, pages 1889–1897, 2015.
- [32] J. Schulman, P. Moritz, S. Levine, M. Jordan, and P. Abbeel. High-dimensional continuous control using generalized advantage estimation. *International Conference on Learning Representations (ICLR)*, 2016.
- [33] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [34] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, et al. Mastering the game of go without human knowledge. *nature*, 550(7676):354–359, 2017.
- [35] A. Stooke, J. Achiam, and P. Abbeel. Responsive safety in reinforcement learning by pid lagrangian methods. In *International Conference on Machine Learning*, pages 9133–9143. PMLR, 2020.
- [36] C. Tessler, D. J. Mankowitz, and S. Mannor. Reward constrained policy optimization. *International Conference on Learning Representation (ICLR)*, 2019.
- [37] E. Vinitzky, A. Kreidieh, L. Le Flem, N. Kheterpal, K. Jang, F. Wu, R. Liaw, E. Liang, and A. M. Bayen. Benchmarks for reinforcement learning in mixed-autonomy traffic. In *Proceedings of Conference on Robot Learning*, pages 399–409, 2018.
- [38] O. Vinyals, I. Babuschkin, W. M. Czarnecki, M. Mathieu, A. Dudzik, J. Chung, D. H. Choi, R. Powell, T. Ewalds, P. Georgiev, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.
- [39] T.-Y. Yang, J. Rosca, K. Narasimhan, and P. J. Ramadge. Projection-based constrained policy optimization. In *International Conference on Learning Representation (ICLR)*, 2020.
- [40] J. Yuan and A. Lamperski. Online convex optimization for cumulative constraints. *Advances in Neural Information Processing Systems*, 31, 2018.
- [41] L. Zhang, L. Shen, L. Yang, S. Chen, B. Yuan, X. Wang, and D. Tao. Penalized proximal policy optimization for safe reinforcement learning. *arXiv preprint arXiv:2205.11814*, 2022.
- [42] S. Zhang, Y. Wan, R. S. Sutton, and S. Whiteson. Average-reward off-policy policy evaluation with function approximation. *arXiv preprint arXiv:2101.02808*, 2021.
- [43] Y. Zhang, Q. Vuong, and K. W. Ross. First order constrained optimization in policy space. *arXiv preprint arXiv:2002.06506*, 2020.

## A Appendix

### A.1 Proofs

**Lemma A.1.** For an episode of length  $H$  and two policies,  $\pi$  and  $\pi'$ , the difference in their performance assuming identical initial state distribution  $\mu$  (i.e.,  $s_1 \sim \mu$ ) is given by

$$J(\pi) - J(\pi') = \sum_{h=1}^H \mathbb{E}_{\substack{s_h, a_h \sim \mathbb{P}_h^\pi(\cdot, \cdot | s) \\ s_1 \sim \mu}} [A_h^{\pi'}(s_h, a_h)]. \quad (2)$$

*Proof.* Let us consider  $s_1 \sim \mu$  and categorize the value function difference between the two policies. Also define  $\mathbb{P}_h^\pi(s | s_1) = \sum_{a \in A} \mathbb{P}_h^\pi(s, a | s_1)$ , where the term  $\mathbb{P}_h^\pi(s, a | s_1)$  is the probability of reaching  $(s, a)$  at time step  $h$  following  $\pi$  and starting from  $s_1$ .

$$\begin{aligned} V_1^\pi(s_1) - V_1^{\pi'}(s_1) &= \mathbb{E}_{a_1, s_2} [r(s_1, a_1) + V_2^\pi(s_2) | s_1] + \mathbb{E}_{s_2} [V_2^{\pi'}(s_2) - V_2^\pi(s_2) | s_1] - V_1^{\pi'}(s_1) \\ &= \mathbb{E}_{s_2} [V_2^\pi(s_2) - V_2^{\pi'}(s_2) | s_1] + \mathbb{E}_{a_1, s_2} [r(s_1, a_1) + V_2^{\pi'}(s_2) - V_1^{\pi'}(s_1) | s_1] \\ &= \mathbb{E}_{s_2} [V_2^\pi(s_2) - V_2^{\pi'}(s_2) | s_1] + \mathbb{E}_{a_1} [Q_1^{\pi'}(s_1, a_1) - V_1^{\pi'}(s_1) | s_1] \\ &= \mathbb{E}_{s_2} [V_2^\pi(s_2) - V_2^{\pi'}(s_2) | s_1] + \mathbb{E}_{a_1} [A_1^{\pi'}(s_1, a_1) | s_1], \end{aligned}$$

where  $a_1 \sim \pi_1(\cdot | s_1)$ ,  $s_2 \sim P(\cdot | s_1, \pi_1(s_1))$  and  $s_1 \sim \mu$ , the initial state distribution.

Now recursively apply the same procedure to the term  $V_h^\pi(s_h) - V_h^{\pi'}(s_h) \quad \forall h \in \{2, \dots, H\}$  to obtain the following:

$$V_1^\pi(s) - V_1^{\pi'}(s) = \sum_{h=1}^H \mathbb{E}_{s_h, a_h \sim \mathbb{P}_h^\pi(\cdot, \cdot | s)} [A_h^{\pi'}(s_h, a_h) | s]$$

Now we know that  $J(\pi) = \mathbb{E}_{s \sim \mu} [V_1^\pi(s)]$ , this means that we combine this with the above to obtain the final result.  $\square$

**Proposition A.2.** The inner optimization problem in (5) with respect to  $\mathbf{x}$  is a convex quadratic program with non-negative constraints, which can be solved to yield the following intermediate problem:

$$\begin{aligned} (\pi_{k,t}^*, \lambda_t^*) &= \max_{\lambda \geq 0} \min_{\pi_{k,t}} \mathcal{L}_t(\pi_k, \lambda, \beta), \quad \text{where} \\ \mathcal{L}_t(\pi_k, \lambda, \beta) &= \sum_{h=t}^H \mathbb{E}_{\substack{s \sim \rho^{\pi_{k,h}} \\ a \sim \pi_{k-1,h}}} [-\rho(\theta_h) A_h^{\pi_{k-1}}(s, a)] + \frac{\beta}{2} \sum_i^m \left( \max \left\{ 0, \Psi_{C_i,t}(\pi_{k-1}, \pi_k) + \frac{\lambda_{t,i}}{\beta} \right\}^2 - \frac{\lambda_{t,i}^2}{\beta^2} \right). \quad (6) \end{aligned}$$

*Proof.* As in Equation (4), we have the following equivalent problem:

$$\pi_{k,t}^* = \arg \min_{\pi_{k,t}} \sum_{h=t}^H \mathbb{E}_{\substack{s \sim \rho^{\pi_{k,h}} \\ a \sim \pi_{k-1,h}}} [-\rho(\theta_h) A_h^{\pi_{k-1}}(s, a)] + \sum_i^m \lambda_{t,i} \max \left\{ 0, \sum_{h=t}^H \mathbb{E}_{\substack{s \sim \rho^{\pi_{k,h}} \\ a \sim \pi_{k-1,h}}} [\rho(\theta_h) A_{C_i,h}^{\pi_{k-1}}(s, a)] + (J_{C_i}(\pi_{k-1}) - d_i) \right\}.$$

Letting  $\Psi_{C_i,t}(\pi_{k-1}, \pi_k) := \sum_{h=t}^H \mathbb{E}_{\substack{s \sim \rho^{\pi_{k,h}} \\ a \sim \pi_{k-1,h}}} [\rho(\theta_h) A_{C_i,h}^{\pi_{k-1}}(s, a)] + (J_{C_i}(\pi_{k-1}) - d_i)$ , and introducing

slack variables  $x_{t,i} \geq 0$  and defining  $w_{t,i}(\pi_k) := \Psi_{C_i,t}(\pi_{k-1}, \pi_k) + x_{t,i} = 0$ , we get the quadratic damped problem same as Equation (5) below.

$$\begin{aligned} (\pi_{k,t}^*, \lambda_t^*, \mathbf{x}_t^*) &= \max_{\lambda \geq 0} \min_{\pi_{k,t}, \mathbf{x}} \mathcal{L}_t(\pi_k, \lambda, \mathbf{x}, \beta) \\ &= \max_{\lambda \geq 0} \min_{\pi_{k,t}, \mathbf{x}} \sum_{h=t}^H \mathbb{E}_{\substack{s \sim \rho^{\pi_{k,h}} \\ a \sim \pi_{k-1,h}}} [-\rho(\theta_h) A_h^{\pi_{k-1}}(s, a)] + \sum_i^m \lambda_{t,i} w_{t,i}(\pi_k) + \frac{\beta}{2} \sum_i^m w_{t,i}^2(\pi_k) \quad (9) \end{aligned}$$

Like the Lagrangian method, we can alternately update  $\pi$ ,  $\lambda$ , and  $\mathbf{x}$  to find the optimal triplet. Consider updating  $\pi$  and  $\mathbf{x}$  by minimizing  $\mathcal{L}_t(\pi, \lambda, \mathbf{x}, \beta)$  at any iteration:

$$(\pi_{k,t}^*, \mathbf{x}_t^*) = \arg \min_{\pi_{k,t}} \min_{x_i > 0} \sum_{h=t}^H \mathbb{E}_{s \sim \rho_{\pi_{k,h}}} [-\rho(\theta_h) A_h^{\pi_{k-1}}(s, a)] + \sum_{i=1}^m \lambda_{t,i} (\Psi_{C_i,t}(\pi_{k-1}, \pi_k) + x_{t,i}) + \frac{\beta}{2} \sum_{i=1}^m (\Psi_{C_i,t}(\pi_{k-1}, \pi_k) + x_{t,i})^2$$

The inner optimization problem with respect to  $\mathbf{x}$  is a convex quadratic programming problem with non-negative constraints,

$$\mathbf{x}_t^* = \arg \min_{x_i > 0} \sum_{i=1}^m \lambda_{t,i} (\Psi_{C_i,t}(\pi_{k-1}, \pi_k) + x_{t,i}) + \frac{\beta}{2} \sum_{i=1}^m (\Psi_{C_i,t}(\pi_{k-1}, \pi_k) + x_{t,i})^2$$

The optimal solution is  $x_{t,i}^* = \max \left\{ 0, -\frac{\lambda_{t,i}}{\beta} - \Psi_{C_i,t}(\pi_{k-1}, \pi_k) \right\}$ . Then,

$$\begin{aligned} w_{t,i}(\pi_k) &= \Psi_{C_i,t}(\pi_{k-1}, \pi_k) + x_{t,i}^* = \Psi_{C_i,t}(\pi_{k-1}, \pi_k) + \max \left\{ 0, -\frac{\lambda_{t,i}}{\beta} - \Psi_{C_i,t}(\pi_{k-1}, \pi_k) \right\} \\ &= \frac{\lambda_{t,i}}{\beta} + \Psi_{C_i,t}(\pi_{k-1}, \pi_k) + \max \left\{ 0, -\frac{\lambda_{t,i}}{\beta} - \Psi_{C_i,t}(\pi_{k-1}, \pi_k) \right\} - \frac{\lambda_{t,i}}{\beta} \\ &= \max \left\{ 0, \frac{\lambda_{t,i}}{\beta} + \Psi_{C_i,t}(\pi_{k-1}, \pi_k) \right\} - \frac{\lambda_{t,i}}{\beta} \end{aligned}$$

Substituting back into Equation (9), we get

$$\begin{aligned} \mathcal{L}_t(\pi_k, \boldsymbol{\lambda}, \mathbf{x}, \beta) &= \sum_{h=t}^H \mathbb{E}_{s \sim \rho_{\pi_{k,h}}} [-\rho(\theta_h) A_h^{\pi_{k-1}}(s, a)] + \sum_{i=1}^m \lambda_{t,i} w_{t,i}(\pi_k) + \frac{\beta}{2} \sum_{i=1}^m w_{t,i}^2(\pi_k) \\ &= \sum_{h=t}^H \mathbb{E}_{s \sim \rho_{\pi_{k,h}}} [-\rho(\theta_h) A_h^{\pi_{k-1}}(s, a)] + \sum_{i=1}^m \lambda_{t,i} \left( \max \left\{ 0, \frac{\lambda_{t,i}}{\beta} + \Psi_{C_i,t}(\pi_{k-1}, \pi_k) \right\} - \frac{\lambda_{t,i}}{\beta} \right) \\ &\quad + \frac{\beta}{2} \sum_{i=1}^m \left( \max \left\{ 0, \frac{\lambda_{t,i}}{\beta} + \Psi_{C_i,t}(\pi_{k-1}, \pi_k) \right\} - \frac{\lambda_{t,i}}{\beta} \right)^2 \\ &= \sum_{h=t}^H \mathbb{E}_{s \sim \rho_{\pi_{k,h}}} [-\rho(\theta_h) A_h^{\pi_{k-1}}(s, a)] + \frac{\beta}{2} \sum_{i=1}^m \left( \max \left\{ 0, \frac{\lambda_{t,i}}{\beta} + \Psi_{C_i,t}(\pi_{k-1}, \pi_k) \right\}^2 - \frac{\lambda_{t,i}^2}{\beta^2} \right) \end{aligned}$$

Hence, we finally get

$$\begin{aligned} (\pi_{k,t}^*, \boldsymbol{\lambda}_t^*) &= \max_{\boldsymbol{\lambda} \geq 0} \min_{\pi_{k,t}} \mathcal{L}_t(\pi_k, \boldsymbol{\lambda}, \beta) \\ &= \max_{\boldsymbol{\lambda} \geq 0} \min_{\pi_{k,t}} \sum_{h=t}^H \mathbb{E}_{s \sim \rho_{\pi_{k,h}}} [-\rho(\theta_h) A_h^{\pi_{k-1}}(s, a)] + \frac{\beta}{2} \sum_{i=1}^m \left( \max \left\{ 0, \Psi_{C_i,t}(\pi_{k-1}, \pi_k) + \frac{\lambda_{t,i}}{\beta} \right\}^2 - \frac{\lambda_{t,i}^2}{\beta^2} \right) \end{aligned}$$

□

**Lemma A.3.** Consider two problems, Problem (P) and Problem (Q) below. For sufficiently large  $\lambda_i > \bar{\lambda} \forall i$  and  $\beta > \bar{\beta}$  for some finite  $\bar{\lambda}$  and finite  $\bar{\beta}$ , the optimal solution set of Problem (Q) (equivalent version of Problem (6)) is identical to the optimal solution set of Problem (P).

Problem (P) :

$$\mathcal{L}_t^P(\pi_k, \boldsymbol{\lambda}, \mathbf{x}, \beta) := \sum_{h=t}^H \mathbb{E}_{s \sim \rho_{\pi_{k,h}}} [-\rho(\theta_h) A_h^{\pi_{k-1}}(s, a)] + \sum_{i=1}^m \lambda_{t,i} w_{t,i}(\pi_k) + \frac{\beta}{2} \sum_{i=1}^m w_{t,i}^2(\pi_k)$$

$$\text{Then, } (\pi_{k,t}^*, \boldsymbol{\lambda}_t^*, \mathbf{x}_t^*) = \max_{\boldsymbol{\lambda} \geq 0} \min_{\pi_{k,t}, \mathbf{x}} \mathcal{L}_t^P(\pi_k, \boldsymbol{\lambda}, \mathbf{x}, \beta) \quad (\text{P})$$

Problem (Q) :

$$\begin{aligned} \mathcal{L}_t^Q(\boldsymbol{\pi}_k, \boldsymbol{\lambda}, \mathbf{x}, \beta) &:= \sum_{h=t}^H \sum_{s \sim \rho^{\boldsymbol{\pi}_k, h}} \mathbb{E}_{a \sim \boldsymbol{\pi}_{k-1, h}} \left[ -\rho(\theta_h) A_h^{\boldsymbol{\pi}_{k-1}}(s, a) \right] + \sum_i^m \lambda_{t,i} \Psi_{C_i, t}^+(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k, \theta) \\ &\quad + \frac{\beta}{2} \sum_i^m \Psi_{C_i, t}^+(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k, \theta)^2 \end{aligned}$$

$$\text{Then, } (\boldsymbol{\pi}_{k,t}^*, \boldsymbol{\lambda}_t^*, \mathbf{x}_t^*) = \max_{\boldsymbol{\lambda} \geq 0} \min_{\boldsymbol{\pi}_k, \boldsymbol{\pi}_k, \boldsymbol{\theta}} \mathcal{L}_t^Q(\boldsymbol{\pi}_k, \boldsymbol{\lambda}, \mathbf{x}, \beta) \quad (\text{Q})$$

, where  $x^+ := \max(0, x)$ , and

$$\Psi_{C_i, t}(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k, \theta) := \sum_{h=t}^H \sum_{s \sim \rho^{\boldsymbol{\pi}_k, h}} \mathbb{E}_{a \sim \boldsymbol{\pi}_{k-1, h}} \left[ \rho(\theta_h) A_{C_i, h}^{\boldsymbol{\pi}_{k-1}}(s, a) \right] + (J_{C_i}(\boldsymbol{\pi}_{k-1}) - d_i).$$

*Proof.* This proof uses some ideas given in [41] for Part 1 below.

**Part 1** - Solution of Problem (P) is solution of Problem (Q).

Suppose  $\bar{\theta}_t$  is the optimum of the constrained Problem (P) augmented with the quadratic penalty. Let  $\bar{\boldsymbol{\lambda}}_t$  be the corresponding Lagrange multiplier vector for its dual problem, and  $\bar{\beta}$  be the additive quadratic penalty coefficient. Then for  $\lambda_{t,i} \geq \|\bar{\boldsymbol{\lambda}}\|_\infty \forall i$  and  $\beta \geq \|\bar{\beta}\|_\infty$ ,  $\bar{\theta}_t$  is also a minimizer of its ReLU-penalized optimization Problem (Q) as below. Let  $\Omega(\theta_t) := \sum_{h=t}^H \sum_{s \sim \rho^{\boldsymbol{\pi}_k, h}} \mathbb{E}_{a \sim \boldsymbol{\pi}_{k-1, h}} \left[ -\rho(\theta_h) A_h^{\boldsymbol{\pi}_{k-1}}(s, a) \right]$ . Then it follows

that:

$$\begin{aligned} \Omega(\theta_t) + \sum_i^m \lambda_{t,i} \Psi_{C_i, t}^+(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k, \theta) + \frac{\beta}{2} \sum_i^m \Psi_{C_i, t}^+(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k, \theta)^2 &\geq \Omega(\theta_t) + \sum_i^m \bar{\lambda}_i \Psi_{C_i, t}^+(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k, \theta) + \frac{\bar{\beta}}{2} \sum_i^m \Psi_{C_i, t}^+(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k, \theta)^2 \\ &\geq \Omega(\theta_t) + \sum_i^m \bar{\lambda}_i \Psi_{C_i, t}(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k, \theta) + \frac{\bar{\beta}}{2} \sum_i^m \Psi_{C_i, t}(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k, \theta)^2 \end{aligned}$$

By assumption,  $\bar{\theta}_t$  is a Karush-Kuhn-Tucker point in the constrained Problem (P), at which KKT conditions are satisfied with the Lagrange multiplier vector  $\bar{\boldsymbol{\lambda}}$  and  $\bar{\beta}$ . We then have:

$$\begin{aligned} \Omega(\theta_t) + \sum_i^m \bar{\lambda}_i \Psi_{C_i, t}(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k, \theta) + \frac{\bar{\beta}}{2} \sum_i^m \Psi_{C_i, t}(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k, \theta)^2 &\geq \Omega(\bar{\theta}_t) + \sum_i^m \bar{\lambda}_i \Psi_{C_i, t}(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k, \bar{\theta}) + \frac{\bar{\beta}}{2} \sum_i^m \Psi_{C_i, t}(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k, \bar{\theta})^2 \\ &= \Omega(\bar{\theta}_t) + \sum_i^m \bar{\lambda}_i \Psi_{C_i, t}^+(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k, \bar{\theta}) + \frac{\bar{\beta}}{2} \sum_i^m \Psi_{C_i, t}^+(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k, \bar{\theta})^2 \\ &= \Omega(\bar{\theta}_t) + \sum_i^m \lambda_{t,i} \Psi_{C_i, t}^+(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k, \bar{\theta}) + \frac{\beta}{2} \sum_i^m \Psi_{C_i, t}^+(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k, \bar{\theta})^2 \end{aligned}$$

, where the first line holds because  $\bar{\theta}_t$  minimizes the Lagrange function, and the second line is derived from the complementary slackness. Thus, we conclude that for the objective function of Problem (Q), call it  $\mathcal{L}^Q(\theta_t)$ , we have  $\mathcal{L}^Q(\theta_t) \geq \mathcal{L}^Q(\bar{\theta}_t)$  for all  $\theta_t \in \Theta$ , which means  $\bar{\theta}_t$  is a minimizer of the quadratic damped optimization Problem (Q).

**Part 2** - Solution of Problem (Q) is solution of Problem (P).

Let  $\tilde{\theta}_t$  be an optimal point of the quadratic damped Problem (Q), with  $\bar{\theta}_t$  and  $\bar{\boldsymbol{\lambda}}$  being the same as defined above. Then, if  $\tilde{\theta}_t$  is in the set of feasible solutions  $S_{\text{feasible}} = \{\theta \mid \Psi_{C_i, t}(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k, \theta) \leq 0 \forall i\}$ , we have:

$$\begin{aligned} \Omega(\tilde{\theta}_t) &= \Omega(\tilde{\theta}_t) + \sum_i^m \lambda_{t,i} \Psi_{C_i, t}^+(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k, \tilde{\theta}) + \frac{\beta}{2} \sum_i^m \Psi_{C_i, t}^+(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k, \tilde{\theta}) \\ &\leq \Omega(\theta_t) + \sum_i^m \lambda_{t,i} \Psi_{C_i, t}^+(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k, \theta) + \frac{\beta}{2} \sum_i^m \Psi_{C_i, t}^+(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k, \theta) \\ &= \Omega(\theta_t) \end{aligned}$$



The inequality above indicates  $\tilde{\theta}_t$  is also optimal in the constrained Problem (P). Now, if  $\tilde{\theta}$  is not feasible, we have:

$$\begin{aligned}
\Omega(\bar{\theta}_t) + \sum_i^m \lambda_{t,i} \Psi_{C_{i,t}}^+(\pi_{k-1}, \pi_k, \bar{\theta}) + \frac{\beta}{2} \sum_i^m \Psi_{C_{i,t}}^+(\pi_{k-1}, \pi_k, \bar{\theta})^2 &= \Omega(\bar{\theta}_t) + \sum_i^m \bar{\lambda}_i \Psi_{C_{i,t}}^+(\pi_{k-1}, \pi_k, \bar{\theta}) + \frac{\bar{\beta}}{2} \sum_i^m \Psi_{C_{i,t}}^+(\pi_{k-1}, \pi_k, \bar{\theta})^2 \\
&= \Omega(\bar{\theta}_t) + \sum_i^m \bar{\lambda}_i \Psi_{C_{i,t}}(\pi_{k-1}, \pi_k, \bar{\theta}) + \frac{\bar{\beta}}{2} \sum_i^m \Psi_{C_{i,t}}(\pi_{k-1}, \pi_k, \bar{\theta})^2 \\
&\leq \Omega(\tilde{\theta}_t) + \sum_i^m \bar{\lambda}_i \Psi_{C_{i,t}}(\pi_{k-1}, \pi_k, \tilde{\theta}) + \frac{\bar{\beta}}{2} \sum_i^m \Psi_{C_{i,t}}(\pi_{k-1}, \pi_k, \tilde{\theta})^2 \\
&\leq \Omega(\tilde{\theta}_t) + \sum_i^m \bar{\lambda}_i \Psi_{C_{i,t}}^+(\pi_{k-1}, \pi_k, \tilde{\theta}) + \frac{\bar{\beta}}{2} \sum_i^m \Psi_{C_{i,t}}^+(\pi_{k-1}, \pi_k, \tilde{\theta})^2 \\
&\leq \Omega(\tilde{\theta}_t) + \sum_i^m \lambda_{t,i} \Psi_{C_{i,t}}^+(\pi_{k-1}, \pi_k, \tilde{\theta}) + \frac{\beta}{2} \sum_i^m \Psi_{C_{i,t}}^+(\pi_{k-1}, \pi_k, \tilde{\theta})^2
\end{aligned}$$

, which is a contradiction to the assumption that  $\tilde{\theta}_t$  is a minimizer of the penalized optimization Problem (Q). Thus,  $\tilde{\theta}_t$  can only be the feasible optimal solution for Problem (P).  $\square$

**Lemma A.4.** Consider two problems, Problem (P') and Problem (R). For sufficiently large  $\beta > \bar{\beta}$  for some finite  $\bar{\beta}$ , the feasible optimal solution set of Problem (R) (equivalent version of Problem (3)) is identical to the solution set of Problem (P').

Problem (P') :

$$\begin{aligned}
\mathcal{L}_t^{P'}(\pi_k, \lambda, \mathbf{x}, \beta) &:= \sum_{h=t}^H \mathbb{E}_{s \sim \rho^{\pi_k, h}, a \sim \pi_{k-1, h}} [-\rho(\theta_h) A_h^{\pi_{k-1}}(s, a)] + \sum_i^m \lambda_{t,i} w_{t,i}(\pi_k) + \frac{\beta}{2} \sum_i^m w_{t,i}^2(\pi_k) \\
\text{Then, } (\pi_{k,t}^*, \lambda_t^*, \mathbf{x}_t^*) &= \max_{\lambda \geq 0} \min_{\pi_{k,t}, \mathbf{x}} \mathcal{L}_t^{P'}(\pi_k, \lambda, \mathbf{x}, \beta) \tag{P'}
\end{aligned}$$

Problem (R) :

$$\begin{aligned}
\mathcal{L}_t^R(\pi_k, \lambda, \mathbf{x}) &:= \sum_{h=t}^H \mathbb{E}_{s \sim \rho^{\pi_k, h}, a \sim \pi_{k-1, h}} [-\rho(\theta_h) A_h^{\pi_{k-1}}(s, a)] + \sum_i^m \lambda_{t,i} w_{t,i}(\pi_k) \\
\text{Then, } (\pi_{k,t}^*, \lambda_t^*, \mathbf{x}_t^*) &= \max_{\lambda \geq 0} \min_{\pi_{k,t}, \mathbf{x}} \mathcal{L}_t^R(\pi_k, \lambda, \mathbf{x}) \tag{R}
\end{aligned}$$

*Proof.* Recall that we are using parameterized policies, hence we overload notation as  $\theta \equiv \pi$  frequently. For brevity, denote  $\Omega_t(\pi) := \sum_{h=t}^H \mathbb{E}_{s \sim \rho^{\pi, h}, a \sim \pi_{k-1, h}} [\rho(\theta_h) A_h^{\pi_{k-1}}(s, a)]$ . We will also go back and forth between the equivalent problems of Problem (P') and Problem (6) of the main paper in Section 3.

**Part 1.** Solution of Problem (R) is solution of Problem (P').

Suppose that  $\pi^*$  is the optimal feasible policy for the primal Problem (R), which is a Lagrangian version of Problem (3). Consider the corresponding Lagrangian dual parameter  $\lambda^*$  of  $\pi^*$ , which satisfies the KKT condition,

$$\nabla_{\pi} \mathcal{L}_t^R(\pi_{k,t}^*, \lambda^*, \mathbf{x}^*) = -\nabla_{\pi} \Omega_t(\pi_{k,t}^*) + \sum_{i=1}^m \lambda_{t,i}^* \nabla_{\pi} w_{t,i}(\pi_{k,t}^*) = 0$$

and the second-order sufficient condition that for all non-zero vectors  $\mathbf{u}$  that satisfy  $\mathbf{u}^T \nabla_{\pi} w_{t,i}(\pi_{k,t}^*) = 0$ , we have

$$\mathbf{u}^T \nabla_{\pi}^2 \mathcal{L}_t^R(\pi_{k,t}^*, \lambda^*, \mathbf{x}^*) \mathbf{u} > 0 \tag{A}$$

Compare Equation (R) and Equation (P'), we have,

$$\begin{aligned}
\nabla_{\pi} \mathcal{L}_t^{P'}(\pi_{k,t}^*, \boldsymbol{\lambda}^*, \mathbf{x}^*, \beta) &= -\nabla_{\pi} \Omega_t(\pi_{k,t}^*) + \sum_{i=1}^m \lambda_{t,i}^* \nabla_{\pi} w_{t,i}(\pi_k^*) + \beta \sum_{i=1}^m w_{t,i}(\pi_k^*) \nabla_{\pi} w_{t,i}(\pi_k^*) \\
&= \nabla_{\pi} \mathcal{L}_t^R(\pi_{k,t}^*, \boldsymbol{\lambda}^*, \mathbf{x}^*) + \beta \sum_{i=1}^m w_{t,i}(\pi_k^*) \nabla_{\pi} w_{t,i}(\pi_k^*) \\
&= 0
\end{aligned}$$

, where we use  $w_{t,i}(\pi_k^*) := \Psi_{C_{i,t}}(\pi_{k-1}, \pi_k^*) + x_{t,i}^* = 0$  with the feasible policy  $\pi_k^*$ . Moreover,

$$\begin{aligned}
\nabla_{\pi}^2 \mathcal{L}_t^{P'}(\pi_{k,t}^*, \boldsymbol{\lambda}^*, \mathbf{x}^*, \beta) &= -\nabla_{\pi}^2 \Omega_t(\pi_{k,t}^*) \\
&\quad + \sum_{i=1}^m \lambda_{t,i}^* \nabla_{\pi}^2 w_{t,i}(\pi_k^*) + \beta \nabla_{\pi} w_t(\pi_k^*) \nabla_{\pi} w_t(\pi_k^*)^T \\
&= \nabla_{\pi}^2 \mathcal{L}_t^R(\pi_{k,t}^*, \boldsymbol{\lambda}^*, \mathbf{x}^*) + \beta \nabla_{\pi} w_t(\pi_k^*) \nabla_{\pi} w_t(\pi_k^*)^T.
\end{aligned}$$

To prove that  $(\pi_{k,t}^*, \boldsymbol{\lambda}^*)$  is a strict minimum solution to  $\mathcal{L}_t^{P'}(\pi_k, \boldsymbol{\lambda}, \mathbf{x}, \beta)$ , we only need to prove the following is true for sufficiently large  $\beta$ ,

$$\nabla_{\pi}^2 \mathcal{L}_t^{P'}(\pi_k^*, \boldsymbol{\lambda}^*, \mathbf{x}^*, \beta) > 0.$$

If the above is not true, then for any large  $\beta$ , there exists  $\mathbf{u}_t$  such that  $\|\mathbf{u}_t\| = 1$  and satisfies

$$\begin{aligned}
\mathbf{u}_t^T \nabla_{\pi}^2 \mathcal{L}_t^{P'}(\pi_k^*, \boldsymbol{\lambda}^*, \mathbf{x}^*, \beta) \mathbf{u}_t &= \mathbf{u}_t^T \nabla_{\pi}^2 \mathcal{L}_t^R(\pi_k^*, \boldsymbol{\lambda}^*, \mathbf{x}^*) \mathbf{u}_t + \beta \left\| \nabla_{\pi} w_t(\pi_k^*)^T \mathbf{u}_t \right\|^2 \leq 0 \\
\Rightarrow \left\| \nabla_{\pi} w_t(\pi_k^*)^T \mathbf{u}_t \right\|^2 &\leq -\frac{1}{\beta} \mathbf{u}_t^T \nabla_{\pi}^2 \mathcal{L}_t^R(\pi_k^*, \boldsymbol{\lambda}^*, \mathbf{x}^*) \mathbf{u}_t \rightarrow 0, \text{ as } \beta \rightarrow \infty.
\end{aligned}$$

Therefore,  $\{\mathbf{u}_h\}$  is a bounded sequence and there must be a limit point, denoted by  $\hat{\mathbf{u}}$ . Then

$$\begin{aligned}
\nabla_{\pi} w_t(\pi_k^*)^T \hat{\mathbf{u}} &= 0 \\
\hat{\mathbf{u}}^T \nabla_{\pi}^2 \mathcal{L}_t^R(\pi_k^*, \boldsymbol{\lambda}^*, \mathbf{x}^*) \hat{\mathbf{u}} &\leq 0.
\end{aligned}$$

The above contradicts Equation (A), so the conclusion. Hence,  $\pi_{k,t}^*$  is also the optimal feasible policy for the primal-dual Problem (P').

**Part 2.** Solution of Problem (P') is solution of Problem (R).

This part is straightforward since it is a standard result. Please see Chapter 2 and Chapter 9 of [9], and Chapter 2 and Chapter 4 of [8] for the proof. For completeness, we provide the result below.

Suppose  $\pi_{k,t}^*$  in the feasible optimal solution set of the primal-dual Problem (P'). Let  $\boldsymbol{\lambda}^*$  be the corresponding dual parameter of  $\pi_{k,t}^*$ . Consider Problem (6), which is an equivalent version of Problem (P'). For any feasible  $\pi_k$ , we have

$$\mathcal{L}_t(\pi_k^*, \boldsymbol{\lambda}^*, \beta) \leq \mathcal{L}_t(\pi_k, \boldsymbol{\lambda}^*, \beta).$$

Now we have two cases:

Case 1. When  $\frac{\lambda_{t,i}^*}{\beta} + \Psi_{C_{i,t}}(\pi_{k-1}, \pi_k^*) > 0$ , we have

$$\begin{aligned}
\mathcal{L}_t(\boldsymbol{\pi}_k, \boldsymbol{\lambda}^*, \beta) &= -\Omega_t(\pi_{k,t}) + \sum_{i=1}^m \lambda_{t,i}^* \Psi_{C_i,t}(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k) + \frac{\beta}{2} \sum_{i=1}^m \Psi_{C_i,t}^2(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k) \\
&= -\Omega_t(\pi_{k,t}) + \sum_{i=1}^m \beta \Psi_{C_i,t}(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k) \left( \frac{\lambda_{t,i}^*}{\beta} + \Psi_{C_i,t}(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k) \right) - \frac{\beta}{2} \sum_{i=1}^m \Psi_{C_i,t}^2(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k) \\
&\leq -\Omega_t(\pi_{k,t})
\end{aligned}$$

where the last step uses  $\beta > 0$ ,  $\Psi_{C_i,t}(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k) < 0$ , and  $\frac{\lambda_{t,i}^*}{\beta} + \Psi_{C_i,t}(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k) > 0$ .

Case 2. When  $\frac{\lambda_{t,i}^*}{\beta} + \Psi_{C_i,t}(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k^*) \leq 0$ , we have

$$\mathcal{L}_t(\boldsymbol{\pi}_k, \boldsymbol{\lambda}^*, \beta) = -\Omega_t(\pi_{k,t}) - \frac{1}{2\beta} \sum_{i=1}^m \lambda_{t,i}^{*2} \leq -\Omega_t(\pi_{k,t}).$$

Now, combining both cases above, we have  $\mathcal{L}_t(\boldsymbol{\pi}_k, \boldsymbol{\lambda}^*, \beta) \leq -\Omega_t(\pi_{k,t})$ . On the other hand,  $\mathcal{L}_t(\boldsymbol{\pi}_k^*, \boldsymbol{\lambda}^*, \beta) = \mathcal{L}_t(\boldsymbol{\pi}_k^*, \boldsymbol{\lambda}^*, \boldsymbol{x}^*, \beta) = -\Omega_t(\pi_{k,t}^*)$ . Thus, the combining all of the above we get

$$-\Omega_t(\pi_{k,t}^*) = \mathcal{L}_t(\boldsymbol{\pi}_k^*, \boldsymbol{\lambda}^*, \beta) \leq \mathcal{L}_t(\boldsymbol{\pi}_k, \boldsymbol{\lambda}^*, \beta) \leq -\Omega_t(\pi_{k,t}).$$

□

**Theorem 3.3.** Let  $\pi^{(3)*}$  be a solution to Problem (3), and let  $(\pi^{(6)*}, \boldsymbol{\lambda}^{(6)*})$  be a solution to Problem (6). Then, for sufficiently large  $\beta > \bar{\beta}$  and  $\lambda_{t,i} > \bar{\lambda} \forall i$ ,  $\pi^{(3)*}$  is a solution to Problem (6), and  $\pi^{(6)*}$  is a solution to Problem (3).

*Proof.* We prove this result as a two-step process.

First, we show that the solution sets of the below problems are identical. See Lemma A.3 for the proof.

Problem (P).

$$\begin{aligned}
\mathcal{L}_t^P(\boldsymbol{\pi}_k, \boldsymbol{\lambda}, \boldsymbol{x}, \beta) &:= \sum_{h=t}^H \mathbb{E}_{\substack{s \sim \rho_{\boldsymbol{\pi}_k, h} \\ a \sim \boldsymbol{\pi}_{k-1, h}}} [-\rho(\theta_h) A_h^{\boldsymbol{\pi}_{k-1}}(s, a)] + \sum_i^m \lambda_{t,i} w_{t,i}(\boldsymbol{\pi}_k) + \frac{\beta}{2} \sum_i^m w_{t,i}^2(\boldsymbol{\pi}_k) \\
\text{Then, } (\boldsymbol{\pi}_{k,t}^*, \boldsymbol{\lambda}_t^*, \boldsymbol{x}_t^*) &= \max_{\boldsymbol{\lambda} \geq 0} \min_{\boldsymbol{\pi}_{k,t}, \boldsymbol{x}} \mathcal{L}_t^P(\boldsymbol{\pi}_k, \boldsymbol{\lambda}, \boldsymbol{x}, \beta) \tag{P}
\end{aligned}$$

Problem (Q).

$$\begin{aligned}
\mathcal{L}_t^Q(\boldsymbol{\pi}_k, \boldsymbol{\lambda}, \boldsymbol{x}, \beta) &:= \sum_{h=t}^H \mathbb{E}_{\substack{s \sim \rho_{\boldsymbol{\pi}_k, h} \\ a \sim \boldsymbol{\pi}_{k-1, h}}} [-\rho(\theta_h) A_h^{\boldsymbol{\pi}_{k-1}}(s, a)] + \sum_i^m \lambda_{t,i} \Psi_{C_i,t}^+(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k, \theta) \\
&\quad + \frac{\beta}{2} \sum_i^m \Psi_{C_i,t}^+(\boldsymbol{\pi}_{k-1}, \boldsymbol{\pi}_k, \theta)^2 \\
\text{Then, } (\boldsymbol{\pi}_{k,t}^*, \boldsymbol{\lambda}_t^*, \boldsymbol{x}_t^*) &= \max_{\boldsymbol{\lambda} \geq 0} \min_{\boldsymbol{\pi}_{k,t}, \boldsymbol{x}} \mathcal{L}_t^Q(\boldsymbol{\pi}_k, \boldsymbol{\lambda}, \boldsymbol{x}, \beta) \tag{Q}
\end{aligned}$$

Second, we show that the solution sets of the below problems are identical. See Lemma A.4 for the proof.

Problem (P') :

$$\begin{aligned}
\mathcal{L}_t^{P'}(\boldsymbol{\pi}_k, \boldsymbol{\lambda}, \boldsymbol{x}, \beta) &:= \sum_{h=t}^H \mathbb{E}_{\substack{s \sim \rho_{\boldsymbol{\pi}_k, h} \\ a \sim \boldsymbol{\pi}_{k-1, h}}} [-\rho(\theta_h) A_h^{\boldsymbol{\pi}_{k-1}}(s, a)] + \sum_i^m \lambda_{t,i} w_{t,i}(\boldsymbol{\pi}_k) + \frac{\beta}{2} \sum_i^m w_{t,i}^2(\boldsymbol{\pi}_k) \\
\text{Then, } (\boldsymbol{\pi}_{k,t}^*, \boldsymbol{\lambda}_t^*, \boldsymbol{x}_t^*) &= \max_{\boldsymbol{\lambda} \geq 0} \min_{\boldsymbol{\pi}_{k,t}, \boldsymbol{x}} \mathcal{L}_t^{P'}(\boldsymbol{\pi}_k, \boldsymbol{\lambda}, \boldsymbol{x}, \beta) \tag{P'}
\end{aligned}$$




Figure 5: The Humanoid, Circle, Reach, Grid, Bottleneck, and Navigation tasks. (a) **Humanoid**: The agent is to run as fast as possible on a flat surface, while not exceeding a specified speed limit i.e. the cost constraint. (b) **Circle**: The agent is rewarded for moving in a specified circle but is penalized if the diameter of the circle is larger than some value [3]. (c) **Reach**: The agent is rewarded for reaching a goal while avoiding obstacles (cost constraints) that are placed to hinder the agent [30]. (d) **Grid**: The agent controls traffic lights in a 3x3 road network and is rewarded for high traffic throughput but is constrained to let lights be red for at most 5 consecutive seconds [37]. (e) **Bottleneck**: The agent controls vehicles (red) in a merging traffic situation and is rewarded for maximizing the number of vehicles that pass through but is constrained to ensure that white vehicles (not controlled by agent) have “low” speed for no more than 10 seconds [37]. (f) **Navigation**: The agent is rewarded for reaching the target area (green) but is constrained to avoid hazards (light purple) and impassible pillars (dark purple). The cost for hazards and pillars is different [30].

Problem (R) :

$$\mathcal{L}_t^R(\pi_k, \lambda, \mathbf{x}) := \sum_{h=t}^H \sum_{s \sim \rho_{\pi_k, h}} \mathbb{E}_{a \sim \pi_{k-1, h}} \left[ -\rho(\theta_h) A_h^{\pi_{k-1}}(s, a) \right] + \sum_i^m \lambda_{t, i} w_{t, i}(\pi_k)$$

Then,  $(\pi_{k, t}^*, \lambda_t^*, \mathbf{x}_t^*) = \max_{\lambda \geq 0} \min_{\pi_{k, t}, \mathbf{x}} \mathcal{L}_t^R(\pi_k, \lambda, \mathbf{x})$  (R)

Now, it follows from equivalency that the optimal solution of Problem (Q) and Problem (R), and hence Problem (6) and Problem (3), is the same. □

## A.2 Experiments Revisited

Below we detail the experimental attributes that we used in benchmarking. See Figure 5 for the environment details. All our experiments are run in the `omnisafe` module [22].

### A.2.1 Environment Details

Comprehensively, our experiments consist of eight tasks ranging from more superficial (Run and Circle tasks) to relatively more stochastic and sophisticated (Bottleneck and Grid tasks), each training different robots. They come from three well-known safe RL benchmark environments, Safe MuJoCo, Bullet-Safety-Gym, and Safety-Gym. For agents maneuvering on a two-dimensional plane, the cost is calculated as  $C(s, a) = \sqrt{v_x^2 + v_y^2}$ . For agents moving along a straight line, the cost is calculated as  $C(s, a) = |v_x|$ , where  $v_x$  and  $v_y$  are the velocities of the agent in the x and y directions.

**Circle** This environment is inspired by [3]. Reward is maximized by moving along a circle of radius  $d$ :

$$R = \frac{v^T[-y, x]}{1 + |\sqrt{x^2 + y^2} - d|},$$

but the safety region  $x_{\text{lim}}$  is smaller than the radius  $d$ :  $C = \mathbf{1}[x > x_{\text{lim}}]$ .

**Navigation** This environment is inspired by [30]. Reward is maximized by getting close to the destination  $R = \text{Dist}(\text{target}, s_{t-1}) - \text{Dist}(\text{target}, s_t)$ , but it yields a cost of +1 when the agent hits the hazard or the pillar. The two different types of cost functions are returned separately and have different thresholds. In our setting,  $d_1 = 25$  for the hazard constraint and  $d_2 = 20$  for the pillar constraint.

Since the main goal in MuJoCo is to train the robot to locomote on the plane, we call it the "Run" task in our article. Our chosen two robots are the relatively complex types in MuJoCo: Ant and Humanoid. **OpenAI Gym** is open source at <https://github.com/openai/gym>, and has a documentation

Hyperparameter	APPO	PDO	FOCOPS	CPPO-PID	IPO	P3O	CPO	TRPO-L	PCPO
Actor Net layers	(32, 32)	(32, 32)	(32, 32)	(32, 32)	(32, 32)	(32, 32)	(32, 32)	(32, 32)	(32, 32)
Critic Net layers	(32, 32)	(32, 32)	(32, 32)	(32, 32)	(32, 32)	(32, 32)	(32, 32)	(32, 32)	(32, 32)
Activation	tanh	tanh	tanh	tanh	tanh	tanh	tanh	tanh	tanh
Initial log std	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
Discount $\gamma$	0.99	0.95	0.995	0.995	0.99	0.99	0.99	0.99	0.99
Policy lr	$3e - 4$	$3e - 4$	$3e - 4$	$3e - 4$	$3e - 4$	$3e - 4$	$3e - 4$	$3e - 4$	$3e - 4$
Critic Net lr	$1e - 3$	$1e - 3$	$1e - 3$	$1e - 3$	$1e - 3$	$1e - 3$	$1e - 3$	$1e - 3$	$1e - 3$
No. of episodes	500	500	500	500	500	500	500	500	500
Steps per epochs	300	300	300	300	300	300	300	300	300
Target KL	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
KL early stop	True	True	True	True	True	True	False	False	False
Line Search Times	N/A	N/A	N/A	N/A	N/A	N/A	25	25	25
Line Search Decay	N/A	N/A	N/A	N/A	N/A	N/A	0.8	0.8	0.8
Proximal clip	0.2	0.2	0.2	0.2	0.2	0.2	N/A	N/A	N/A
Max horizon	200	200	200	200	200	200	200	200	200

Table 3: Hyperparameters used for each baseline.

at <https://www.gymnasium.ml/>. **Bullet Safety Gym**. The implementation of the Circle task comes from Bullet-safety-gym (Gronauer 2022), which Stooke, Achiam, and Abbeel (2020) first proposed. The reward is dense and increases by the agent’s velocity and the proximity to the boundary of the circle. Costs are received when the agent leaves the safety zone defined by the two yellow boundaries. The environment is open source at <https://github.com/SvenGronauer/Bullet-Safety-Gym>.

**Safety Gym**. The remaining two tasks, Goal and Button, are from Safety-Gym (Ray, Achiam, and Amodei 2019). Compared to Run and Circle tasks, they are more stochastic and sophisticated in that agents are challenged to maximize the return while satisfying the constraints.

The environment is open source at <https://github.com/openai/safety-gym>, and readers can see OpenAI’s blog at <https://openai.com/blog/safety-gym/> for more details.

### A.3 Agents

For single-constraint scenarios, Point agent is a 2D mass point ( $A \subseteq \mathbb{R}^2$ ) and Ant is an quadruped robot ( $A \subseteq \mathbb{R}^8$ ). For the multi-constraint scenario which is modified from OpenAI SafetyGym [30],  $S \subseteq \mathbb{R}^{28+16 \cdot m}$  where  $m$  is the number of pseudo-radar (one for each type of obstacles and we set two different types of obstacles in the Navigation task) and  $A \subseteq \mathbb{R}^2$  for a mass point or a wheeled car.

#### A.3.1 Experimental Details

To be fair in comparison, the proposed e-COP algorithm and FOCOPS [43] are implemented with same rules and tricks on the code-base of [30].

#### A.3.2 Hyperparameters

Table 3 shows the hyperparameters of baseline algorithms.

#### A.3.3 Runtime Environment

All experiments were implemented in Pytorch 1.7 .0 with CUDA 11.0 and conducted on an Ubuntu 20.04.2 LTS with 8 CPU cores (AMD Ryzen Threadripper PRO 3975WX 8-Coresz), 127G memory and 2 GPU cards (NVIDIA GeForce RTX 4060 Ti Cards).

#### A.3.4 Robustness to Cost Thresholds

We conducted a set of experiments wherein we study how e-COP effectively adapts to different cost thresholds. For this, we use a pre-trained e-COP agent, which is trained with a particular cost threshold in

an environment, and test its performance on different cost thresholds within the same environment. Figure 6 illustrates the training curves of these pre-trained agents, and we see that while e-COP can generalize well across different cost thresholds, other baseline algorithms may require further tuning to accommodate different constraint thresholds.




Figure 6: Cumulative episodic rewards and costs of baselines in two environments with two different constraint cost thresholds: 10 and 50 in Humanoid, and 3 and 15 in Point Circle. The hyperparameters are tuned at constraint limit of 20 in Humanoid and 10 in Point Circle.



## NeurIPS Paper Checklist

### 1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope?

Answer: [Yes]

Justification: We make claims on developing the theory for episodic policy optimization for CMDPs and on presenting experimental validation of our approach. Please see Sections 3 and 4.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

### 2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: While there are no algorithmic limitations, potential limitations can occur while applying our work to very high dimensional real-life scenarios, for instance, in large language models or the denoising process in image and video diffusion tasks. We have developed this work with the above two applications in mind, and plan to explore that direction of future research.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren’t acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

### 3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: All our lemmas and theorems have proofs, some which are available in the main text, and some in Appendix A.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

#### 4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We provide the evaluation protocol and baseline comparison metrics in Section 4. More information is also available in Appendix A.2.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

#### 5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We have provided the code.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

## 6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We detail our experimental procedure and details in Sections 4 and A.2.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

## 7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: All our results are reported with 1 standard deviation across multiple seeds. Please see Section 4.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).

- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

## 8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: In Section A.2 we detail the compute resources used.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

## 9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics [https://neurips.cc/public/EthicsGuidelines?](https://neurips.cc/public/EthicsGuidelines)

Answer: [Yes]

Justification: We have read the code of ethics and our work satisfies it in every respect.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

## 10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: In the Introduction Section 1 we detail the potential applications of our work.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point

out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.

- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

## 11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Our work needs no such safeguards for responsible release of assets.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

## 12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: We do not use existing assets. Proper references for existing results and algorithms are provided throughout the main text.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, [paperswithcode.com/datasets](https://paperswithcode.com/datasets) has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

## 13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: We do not release new assets. Our results, both theoretical, algorithmic, and empirical, are properly referenced and explained in the main text.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

#### 14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: We do not use crowdsourcing or research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

#### 15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: Our work does not involve crowdsourcing or research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.