

Incentivizing Safer Actions in Policy Optimization for Constrained Reinforcement Learning

Somnath Hazra¹, Pallab Dasgupta² and Soumyajit Dey¹

¹Indian Institute of Technology Kharagpur, India

²Synopsys, USA

somnathhazra@kgpian.iitkgp.ac.in, pallabd@synopsys.com, soumya@cse.iitkgp.ac.in

Abstract

Constrained Reinforcement Learning (RL) aims to maximize the return while adhering to predefined constraint limits, which represent domain-specific safety requirements. In continuous control settings, where learning agents govern system actions, balancing the trade-off between reward maximization and constraint satisfaction remains a significant challenge. Policy optimization methods often exhibit instability near constraint boundaries, resulting in suboptimal training performance. To address this issue, we introduce a novel approach that integrates an adaptive incentive mechanism in addition to the reward structure to stay within the constraint bound before approaching the constraint boundary. Building on this insight, we propose Incrementally Penalized Proximal Policy Optimization (IP3O), a practical algorithm that enforces a progressively increasing penalty to stabilize training dynamics. Through empirical evaluation on benchmark environments, we demonstrate the efficacy of IP3O compared to the performance of state-of-the-art Safe RL algorithms. Furthermore, we provide theoretical guarantees by deriving a bound on the worst-case error of the optimality achieved by our algorithm.

1 Introduction

Constraint satisfaction represents a critical challenge in Reinforcement Learning (RL), particularly due to its significant implications in real-world applications such as robotics [Levine *et al.*, 2016; Ono *et al.*, 2015], autonomous driving [Fisac *et al.*, 2018; Kiran *et al.*, 2021; Fernandez-Llorca and Gómez, 2023], healthcare [Yu *et al.*, 2021], finance [McNamarra, 2016], etc. These tasks involve learning a policy to address sequential decision-making problems to maximize return under predefined safety constraints and accounting for uncertainties in the environment. Such requirements are commonly formulated using the Constrained Markov Decision Process (CMDP) framework [Altman, 2021], which extends

the conventional Markov Decision Process (MDP) to accommodate safety-critical considerations. Consequently, the applicability of traditional RL methods becomes limited in these settings, necessitating the development of new approaches to address constraint satisfaction effectively.

Traditional RL algorithms focus solely on maximizing the return [Schulman *et al.*, 2017]. To satisfy the cost budget, the safe RL methods typically incorporate an additional term into the loss function. However, since the cost component is not inherently optimized, the imposed penalty significantly dictates the resulting policy’s behavior; it should be rewarding enough as far as the cost constraint limits are respected. For instance, Constrained Policy Optimization (CPO) [Achiam *et al.*, 2017] uses a second-order expansion term to approximate the cost but suffers from high computational overhead and suboptimal performance due to approximation errors. Primal-dual methods transform the constrained problem into an unconstrained dual via Lagrangian multipliers [Tessler *et al.*, 2018; Ding *et al.*, 2020; Yu *et al.*, 2019; Dai *et al.*, 2023], yet these approaches are prone to constraint violations and oscillatory behavior in practical scenarios [Stooke *et al.*, 2020]. Projection-based methods apply a secondary optimization to project unsafe actions into a feasible region [Zhang *et al.*, 2020; Yang *et al.*, 2020; Yang *et al.*, 2022]; however, these methods are computationally expensive and can underperform in certain environments [Dai *et al.*, 2023].

In contrast, penalty function approaches directly encode the constraints in the primal space itself. By introducing a penalty term into the loss function based on the constraint cost, these methods establish a penalty barrier to regulate policy behavior [Liu *et al.*, 2020; Zhang *et al.*, 2022; Zhang *et al.*, 2023]. The choice of barrier function is critical, as it determines the trade-off between performance and constraint satisfaction [Zhang *et al.*, 2023]. Despite their effectiveness, most existing approaches impose penalties only after a constraint violation occurs, offering no proactive incentive for the policy to remain within predefined safety limits. This strategy often compromises the stability of training and constraint satisfaction. To address this limitation, we propose a novel approach that transforms the cost function into a positive incentive within the safe region and gradually converts it into a penalty as the safety boundary is breached. The design ensures a more dynamic adherence to constraints, leading to

Code and Supplementary Material available here:
<https://github.com/somnathhazra/IP3O>.

more stable and efficient optimization.




Figure 1: Illustration of a risky trajectory and a conservation trajectory under unknown environmental dynamics.

To illustrate our hypothesis, consider a robot tasked with reaching the rock, i.e. the goal, located on the opposite side of a pond while avoiding falling into the water (Figure 1). The shortest path, marked in red, is the most efficient but also the riskiest due to environmental uncertainties. While on-policy algorithms in RL can mitigate such risks [Sutton and Barto, 2018], given the cost function, the reward function alone may fail to discourage unsafe actions effectively in a CMDP. Introducing an incentive derived from the cost function encourages safer behavior. However, excessive incentives for avoiding risk can result in overly conservative policies that fail to achieve the goal, as indicated by the blue path in the figure. To address this trade-off, we propose Incrementally Penalized Proximal Policy Optimization (IP3O), which incorporates an adaptive penalty mechanism to balance safety and performance under general model-free settings.

The primary contributions are summarized as follows:

- We design a novel penalty function to adaptively incentivize safe actions based on the cost function, ensuring that safety constraints are satisfied without compromising return optimization.
- We propose the Incrementally Penalized Proximal Policy Optimization (IP3O) algorithm, which integrates the penalty barrier function to achieve efficient learning of safe policies. Additionally, we provide theoretical guarantees by deriving a worst-case performance bound for the algorithm.
- We conduct extensive empirical evaluations of IP3O against multiple state-of-the-art approaches in safe RL across benchmark environments, demonstrating its superior performance in balancing safety and return.

These contributions collectively address key challenges in constrained RL by bridging the gap between safety and performance in real-world scenarios. We next present the relevant background in the Preliminaries section.

2 Preliminaries

RL problems are modeled using a Markov Decision Process (MDP), defined by the tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{R}, P, \rho, \gamma \rangle$, where \mathcal{S} denotes the state space, \mathcal{A} denotes the action space, $\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the reward function, $P : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ is transition probability function, $\rho : \mathcal{S} \rightarrow [0, 1]$ is the initial state distribution, and $\gamma \in (0, 1)$ is the discount factor

for calculating the return. In model-free settings, P is unknown. The objective is to find a policy $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$ that maximizes the expected discounted return: $\mathcal{J}_{\mathcal{R}}(\pi) = \mathbb{E}_{\tau \sim \pi}[\sum_{t=0}^{\infty} \gamma^t \mathcal{R}(s_t, a_t)]$, where τ is the trajectory. The utility of a state s is measured using the value function: $V_{\mathcal{R}}^{\pi}(s) = \mathbb{E}_{\tau \sim \pi}[\sum_{t=0}^{\infty} \gamma^t \mathcal{R}(s_t, a_t) | s_0 = s]$; and that after taking action a at state s is given by the state-action value function: $Q_{\mathcal{R}}^{\pi}(s, a) = \mathbb{E}_{\tau \sim \pi}[\sum_{t=0}^{\infty} \gamma^t \mathcal{R}(s_t, a_t) | s_0 = s, a_0 = a]$. The advantage of taking action a at s is given by the advantage function: $A_{\mathcal{R}}^{\pi}(s, a) = Q_{\mathcal{R}}^{\pi}(s, a) - V_{\mathcal{R}}^{\pi}(s)$.

The Constrained MDP (CMDP) introduces the cost function, $\mathcal{C} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, and is represented using the tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{C}, P, \rho, \gamma \rangle$. The allowable policy set $\Pi_{\mathcal{C}} \subset \Pi$ is restricted such that the expected cumulative discounted cost satisfies $\mathcal{J}_{\mathcal{C}}(\pi) \leq d$, where d is the constraint threshold. The value functions, $V_{\mathcal{C}}^{\pi}, Q_{\mathcal{C}}^{\pi}, A_{\mathcal{C}}^{\pi}$ are similarly defined as in MDPs using \mathcal{C} . Thus, constrained RL aims to find a policy $\pi \in \Pi_{\mathcal{C}}$, where:

$$\arg \max_{\pi} \mathcal{J}_{\mathcal{R}}(\pi) \quad \text{s.t.} \quad \mathcal{J}_{\mathcal{C}}(\pi) \leq d \quad (1)$$

A more general consideration is where \mathcal{C} represents a set of m constraints. In those cases, the objective is to find a policy where $\arg \max_{\pi} \mathcal{J}_{\mathcal{R}}(\pi) \quad \text{s.t.} \quad \mathcal{J}_{\mathcal{C}_i}(\pi) \leq d_i$. In the rest of the text we assume the multi-constraint formulation to describe our approach.

3 Related Works

Constrained RL in model-free settings remains a challenging problem due to the difficulty in converging to near-optimal solutions under high-dimensional approximations and environmental uncertainties. Recent research has introduced various approaches to address these challenges, focusing on effective constraint handling during policy optimization.

The constrained RL objective is often built as a dual problem using Lagrangian relaxation [Chow *et al.*, 2018], converting episodic constraints into an unconstrained dual objective. RCPO [Tessler *et al.*, 2018] penalizes policy updates based on constraint violations with a Lagrangian multiplier. However, dual optimization adds oscillatory behavior, mitigated in CPPOPID [Stooke *et al.*, 2020] with a PID controller, or using an Augmented Lagrangian in APPO [Dai *et al.*, 2023]. While sometimes effective, these methods are very sensitive to Lagrangian parameters, limiting their applicability.

On-policy RL algorithms, that update the existing policy within a trust region, have been extensively used in constrained RL to safely update the existing policy. Algorithms such as TRPO [Schulman, 2015] enforce trust-region updates through divergence constraints, while PPO [Schulman *et al.*, 2017] simplifies this using gradient clipping. Building on TRPO, CPO [Achiam *et al.*, 2017] integrates second-order approximation of constraints; but suffers from the computational overhead of Fisher Information Matrix inversion. Methods such as PCPO [Yang *et al.*, 2020] and CUP [Yang *et al.*, 2022] adopts two-phase optimization: policies are updated and then projected back into the feasible space. FO-COPS [Zhang *et al.*, 2020] introduces a non-parametric stage for feasible policy generation, followed by parametric projection. However, these projection-based approaches often fail to guarantee optimality, highlighting a critical limitation.

Penalty barrier methods impose a penalty based on the constraint violation, restricting optimization within predefined barriers. IPO [Liu *et al.*, 2020] uses a logarithmic penalty function; whereas P3O [Zhang *et al.*, 2022] employs a ReLU-based penalty. Yet, these methods activate penalties only after constraints are breached, leading to a non-smooth transition in learning dynamics. The reward function may not incentivize a safe action. As described in the next section, our work builds on this line of research by introducing a smooth transition from incentives to penalties, facilitating stable policy learning while ensuring efficient adherence to constraints.

4 Methodology

Policy optimization in reinforcement learning improves a policy iteratively by balancing exploration and exploitation of the current policy’s knowledge. The performance difference between two policies is given by the following equation [Kakade and Langford, 2002]:

$$\mathcal{J}_{\mathcal{R}}(\pi') - \mathcal{J}_{\mathcal{R}}(\pi) = \frac{1}{1 - \gamma} \mathbb{E}_{s \sim d^{\pi'}, a \sim \pi'} [A_{\mathcal{R}}^{\pi}(s, a)] \quad (2)$$

where $d^{\pi}(s) = (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t P(s_t = s | \pi)$ is the discounted state distribution under policy π . Extending this to cost functions \mathcal{C} , the constrained optimization objective (1) can be reformulated as follows:

$$\begin{aligned} \pi_{k+1} &= \arg \max_{\pi} \mathbb{E}_{s \sim d^{\pi}, a \sim \pi} [A_{\mathcal{R}}^{\pi}(s, a)] \\ \text{s.t. } \mathcal{J}_{\mathcal{C}_i}(\pi_k) + \frac{1}{1 - \gamma} \mathbb{E}_{s \sim d^{\pi}, a \sim \pi} [A_{\mathcal{C}_i}^{\pi}(s, a)] &\leq d_i \end{aligned} \quad (3)$$

where π_k is the current policy. Given the continuous state space, we use a parametric policy as $\pi(\theta_k) = \pi_k$. To approximate $A_{\mathcal{R}}^{\pi_k}$ and $A_{\mathcal{C}_i}^{\pi_k}$ using samples from the previous policy, π_k , the importance sampling ratio is used: $r(\theta) = \frac{\pi(\theta)}{\pi(\theta_k)}$. This modifies the optimization problem as follows.

$$\begin{aligned} \pi_{k+1} &= \arg \max_{\pi} \mathbb{E}_{s \sim d^{\pi_k}, a \sim \pi_k} [r(\theta) A_{\mathcal{R}}^{\pi_k}(s, a)] \\ \text{s.t. } \frac{1}{1 - \gamma} \mathbb{E}_{s \sim d^{\pi_k}, a \sim \pi_k} [r(\theta) A_{\mathcal{C}_i}^{\pi_k}(s, a)] + \mathcal{J}_{\mathcal{C}_i}(\pi_k) &\leq d_i \end{aligned} \quad (4)$$

To stabilize the policy update within a trust region, the ratio $r(\theta)$ is clipped within $(1 - \epsilon, 1 + \epsilon)$, where ϵ is the clipping threshold. In penalty function approaches, constraint violations are penalized through barrier functions added to the reward maximization objective. Existing works [Liu *et al.*, 2020; Zhang *et al.*, 2022; Gao *et al.*, 2024] introduce penalties only after constraints are violated, leading to abrupt learning transitions and suboptimal policies due to risk-averse updates.

Our approach addresses these limitations by introducing an *Incrementally Penalized Proximal Policy Optimization (IP3O)* algorithm. The key features of our approach are:

- Incentivizing safe behavior: When constraints are satisfied, an incentive is encoded in the penalty barrier function to guide the policy toward constraint satisfaction without abrupt penalties.
- Gradual transition to penalties: For unsafe trajectories, a smooth transition to penalization is incorporated, mitigating sudden changes in the learning dynamics.

This formulation ensures stable policy updates, enabling a smoother interplay between reward optimization and constraint satisfaction. In this section, we describe the IP3O algorithm in detail, including its implementation and theoretical guarantees.

4.1 Penalty Function Design

Penalty functions, such as Leaky ReLU, allow negative activation for values below zero, incentivizing constraint satisfaction. However, this has a major drawback; their inability to reduce the slope in the negative region leads to excessive incentives for remaining deeply within the constraint region, which may overshadow rewards and result in overly conservative policies that fail to reach the goal (illustrated by the blue path in Figure 1).

To address this issue, we propose using the Exponential Linear Unit (ELU) function [Clevert, 2015] as the penalty function. ELU meets two key requirements: • smoothly transitions from incentivizing constraint satisfaction to penalizing violations, ensuring stable policy updates. • mitigates the risk of overly conservative policies by nullifying excessive incentives within the constraint region. It is defined using the following function.

$$\text{ELU}(x, \alpha) = \begin{cases} x & \text{if } x \geq 0 \\ \alpha(\exp(x) - 1) & \text{otherwise} \end{cases} \quad (5)$$

Here, α determines the stagnation point ($-\alpha$); i.e., if $\alpha = 1$ the function gradually reduces to -1 for negative values of x and then stabilizes. So higher values promote safety by favoring constraint satisfaction. Despite its advantages, the ELU function introduces gradient discontinuity at $x = 0$ when $\alpha \neq 1$ [Barron, 2017]. To overcome this limitation, we employ the Continuously Differentiable ELU (CELU) function that proposes a modification for negative values of x as: $\text{CELU}(x, \alpha) = \alpha(\exp(x/\alpha) - 1)$, if $x < 0$. CELU maintains the properties of ELU while ensuring gradient continuity, making it more suitable for policy optimization. A comparison among the activation functions and their gradients are shown in Figure 2.




Figure 2: Comparison of activation functions and their gradients.

4.2 Practical Implementation

We integrate the proposed CELU-based penalty function into the cost critic, combining it with the reward critic to guide policy updates. These critics provide gradient signals for optimizing the policy network, as shown in Figure 3. The loss function for the reward is derived from Equation 4 as: $\mathcal{L}_{\mathcal{R}}(\pi_k) = \mathbb{E}_{s \sim d^{\pi_k}, a \sim \pi_k} [-r'(\theta) A_{\mathcal{R}}^{\pi_k}(s, a)]$,




Figure 3: The overall algorithm describing the loss function for training the policy network, with the sources of loss.

where $r'(\theta) = \min(r(\theta), \text{clip}(r(\theta), 1 - \epsilon, 1 + \epsilon))$. Similarly the loss due to the cost constraint is derived as: $\mathcal{L}_{C_i}(\pi_k) = \frac{1}{1-\gamma} \mathbb{E}_{s \sim d^{\pi_k}} [r''(\theta) A_{C_i}^{\pi_k}(s, a)] + (\mathcal{J}_{C_i}(\pi_k) - d_i)$, where $r''(\theta) = \max(r(\theta), \text{clip}(r(\theta), 1 - \epsilon, 1 + \epsilon))$. The combined loss function is stated as follows.

$$\mathcal{L}(\pi_k) = \mathcal{L}_{\mathcal{R}}(\pi_k) + \eta \sum_{i=1}^m \text{CELU}(\mathcal{L}_{C_i}(\pi_k)) \quad (6)$$

where $\eta > 0$ is the penalty factor. This formulation ensures a smooth transition from incentivizing constraint satisfaction to penalizing unsafe behaviors. Since the CELU function uses an exponential term, the gradient of the loss arising from CELU never reaches 0. To ensure no policy updates owing to residual gradients we can use a small positive number $h < \alpha$, such that $\max(\text{CELU}(\mathcal{L}_{C_i}(\pi_k)), -\alpha(1 - h))$ instead of directly using the CELU function in Equation 6. However in practical applications, we did not require this. To ensure equivalence between the solutions of the proposed loss function (6) and the constrained optimization objective (4), we establish the following theorem. This result guarantees that the proposed approach preserves the optimality of the original problem while promoting stable and robust policy updates. Hyper-parameter settings, including η and α are detailed in the Supplementary Material.

Theorem 1. *Given a sequence of policies $\{\pi_k\}$ obtained by minimizing $\mathcal{L}(\pi_k)$ and considering Slater’s condition for strong duality, let λ^* be the Lagrange multiplier for the optimal solution of (4). If η satisfies $\eta \geq \|\lambda^*\|_\infty$, the limit π^* of $\{\pi_k\}$ is also a solution to (4).*

Proof. Provided in the Supplementary Material. \square

The value evaluation approximation using the current policy π_k , instead of the currently evaluating policy, π , as done in the original problem; introduces a bias [Jiang and Li, 2016]. Moreover the approximation of the loss due to the cost function used may induce a penalty on the optimal reward resulting from the optimization of the original problem in Equation 4. Using the following theorem we establish the worst-case error bound for these approximations.

Theorem 2. *If the loss function $\mathcal{L}(\pi_k)$ is optimized instead of the original problem (4), the upper bound on the error is given by the following.*

$$\frac{\sqrt{2\delta}\gamma\epsilon_{\mathcal{R}}^{\pi}}{1-\gamma} + \eta \sum_{i=1}^m \left[\frac{\sqrt{2\delta}\gamma\epsilon_{C_i}^{\pi}}{1-\gamma} + |\alpha \log(h)| \right] \quad (7)$$

where $\epsilon_{\mathcal{R}}^{\pi} = \max_s |\mathbb{E}_{a \sim \pi} [A_{\mathcal{R}}^{\pi_k}(s, a)]|$, $\epsilon_{C_i}^{\pi} = \max_s |\mathbb{E}_{a \sim \pi} [A_{C_i}^{\pi_k}(s, a)]|$, and $\delta = \mathbb{E}_{s \sim d^{\pi_k}} [D_{KL}(\pi || \pi_k)[s]]$.

Proof. Provided in the Supplementary Material. \square

In Algorithm 1 below we outline the pseudocode for policy update using our penalty function.

Algorithm 1 Policy optimization using IP3O

Input: Initial policy π_0 , initial value function $V_{\mathcal{R}}^{\pi_0}$, initial cost value function/s $V_{C_i}^{\pi_0}$

- 1: **for** $k = 0, \dots, K - 1$ **do**
- 2: Sample training batch $\mathcal{D}_k = \{\tau_1, \dots, \tau_N\}$ consisting of N trajectories using π_k
- 3: Compute $V_{\mathcal{R}}^{\pi_k}, V_{C_i}^{\pi_k}$ for trajectories in \mathcal{D}_k
- 4: # Advantage calculation
- 5: Compute $A_{\mathcal{R}}^{\pi_k}(s, a) = Q_{\mathcal{R}}^{\pi_k}(s, a) - V_{\mathcal{R}}^{\pi_k}(s)$
- 6: Compute $A_{C_i}^{\pi_k}(s, a) = Q_{C_i}^{\pi_k}(s, a) - V_{C_i}^{\pi_k}(s)$
- 7: Update: $V_{\mathcal{R}}^{\pi_k}(s), V_{C_i}^{\pi_k}(s) \rightarrow V_{\mathcal{R}}^{\pi_{k+1}}(s), V_{C_i}^{\pi_{k+1}}(s)$
- 8: # Policy update
- 9: **for** $t = 0, \dots, T - 1$ **do**
- 10: Compute $\mathcal{L}_{\mathcal{R}}(\pi_k) = \mathbb{E}_{s \sim d^{\pi_k}} [-r'(\theta) A_{\mathcal{R}}^{\pi_k}(s, a)]$
- 11: Compute $\frac{\mathcal{L}_{C_i}(\pi_k)}{1-\gamma} \mathbb{E}_{s \sim d^{\pi_k}} [r''(\theta) A_{C_i}^{\pi_k}(s, a)] + (\mathcal{J}_{C_i}(\pi_k) - d_i) =$
- 12: Compute $\mathcal{L}(\pi_k)$ using Equation 6
- 13: $\pi = \pi_k + \omega \cdot \nabla \mathcal{L}(\pi_k)$
- 14: # Trust region criterion (Gradient clipping)
- 15: **if** $\mathbb{E}_{s \sim d^{\pi_k}} [D_{KL}(\pi || \pi_k)[s]] \notin [\delta^-, \delta^+]$ **then**
- 16: break
- 17: **end if**
- 18: **end for**
- 19: Update $\pi_k \rightarrow \pi_{k+1}$
- 20: **end for**

Return: Trained policy π_K

The output of the algorithm is the final policy. Here ω is the learning rate. For simplicity we have shown trust region updates using the KL divergence criterion, but in practice we use the PPO updates through gradient clipping [Schulman *et al.*, 2017]. In the next section we present the empirical details to demonstrate the efficacy of our approach on some benchmark environments in safe RL.

5 Experiments

In this section, we present the evaluation results of our proposed algorithm on benchmark safe RL environments and compare its performance with state-of-the-art approaches. For benchmarking, we consider the following methods.

- First order methods such as CUP [Yang *et al.*, 2022] and FOCOPS [Zhang *et al.*, 2020].
- Lagrangian methods such as CPPPID [Stooke *et al.*, 2020], and PPO [Schulman *et al.*, 2017] with a Lagrangian multiplier.
- Second order methods such as CPO [Achiam *et al.*, 2017], and PCPO [Yang *et al.*, 2020].

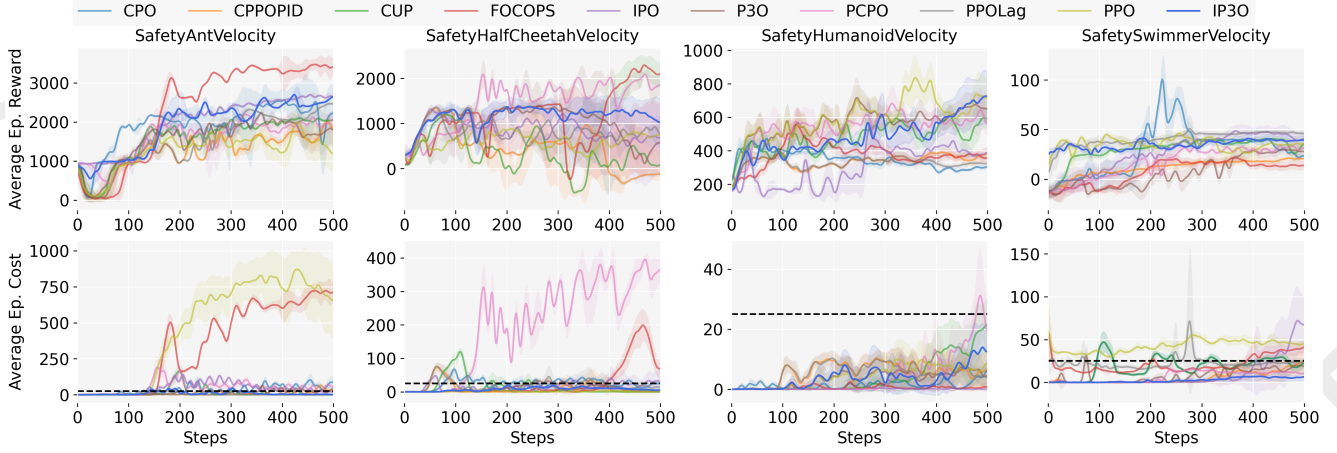


Figure 4: Comparison with the baselines using the MuJoCo safety velocity scenarios. Our method is marked as IP3O (blue). The dashed line indicates the cost constraint limit.

- Penalty function methods such as IPO [Liu *et al.*, 2020], and P3O [Zhang *et al.*, 2022].

All the above algorithms use an on-policy buffer for learning. Additionally, we include a comparison with the vanilla PPO algorithm [Schulman *et al.*, 2017] as a baseline. All baseline implementations are adapted from the OmniSafe repository [Ji *et al.*, 2024]. The evaluations were conducted across three widely-used environments: MuJoCo Safety Velocity [Ji *et al.*, 2023], Safety Gymnasium [Ray *et al.*, 2019], and Bullet Safety Gymnasium [Gronauer, 2022]. These environments provide diverse challenges that test the agent’s ability to maximize cumulative rewards while adhering to pre-defined safety constraints, as formulated in Equation 1. We also evaluate our approach on multi-agent environments, using the MetaDrive simulator [Li *et al.*, 2022], which are detailed later. Our approach allows tuning of the safety level via the hyper-parameter α , which is analyzed further in the Ablation Studies section.

5.1 MuJoCo Safety Velocity

The Safety Velocity environments simulate autonomous robots using MuJoCo physics simulator. The objective is to control the agent to run as fast as possible while ensuring compliance with the velocity constraints of the robot. We evaluate our approach on the Ant, Half-Cheetah, Humanoid, and Swimmer robots. Each environment has varying episode lengths. For these experiments, we set the constraint limit to 25 and use $\alpha = 0.5$. Detailed environment descriptions are provided in the Supplementary Material.

As shown in Figure 4, IP3O obtains better returns while maintaining strict compliance with the velocity constraints compared to state-of-the-art algorithms. Algorithms such as FOCOPS and PCPO shows better returns compared to IP3O for the Ant and HalfCheetah environments, but struggle to find a feasible policy. Other penalty based algorithms, such as IPO and P3O demonstrate sub-optimal returns owing to sharp gradient change near the constraint boundary. This demonstrates that our approach balances safety and performance effectively in dynamic control tasks.

5.2 Safety Gymnasium

Safety Gymnasium environments present the challenge of navigating the agent towards goal state using RL policy in continuous state-space environments. The agent needs to avoid unsafe states such as pre-defined hazardous regions, or avoiding interactions with certain unsafe moving objects; while navigating. Generally the hazards and goals are detected using lidar signals returned by the environments at each step, that are part of the observation space for the agents. The rewards can be maximized by going towards the goal states, while going away from them incurs negative rewards. We evaluate our method in the Goal and Button tasks, where each episode runs for 1,000 time-steps. For these experiments, the constraint limit is 25, and we set $\alpha = 0.1$.

Figure 5 shows that our algorithm consistently outperforms baseline methods in adhering to safety constraints within the policy space Π_C while achieving competitive or higher cumulative rewards. This highlights the robustness of our approach in navigation tasks with complex safety requirements. Algorithms such as PCPO and IPO obtains better returns but fail to find a safety compliant policy. Owing to lower value for α , the learnt policy frequently shows constraint violations, but obtains much better rewards, especially in the Goal tasks.

5.3 Bullet Safety Gymnasium

The Bullet Safety Gymnasium environments consist of situations similar to Safety Gymnasium, and feature tasks such as circular movement near boundaries or gathering objects, using robots like Ball and Car, while staying within constraint barriers. Observation spaces include sensor data for nearby obstacles and task-specific information, such as distances to goals. These environments test safety-critical behaviors in constrained settings. For our evaluations, we set $\alpha = 1.0$ and a constraint limit of 25 across all scenarios. Detailed descriptions of the tasks are available in the Supplementary Material.

From the results in Figure 6, it is evident that IP3O achieves the best compliance with safety constraints among other methods; except in the CarReach environment, the policy demonstrates negligible constraint violation. However,

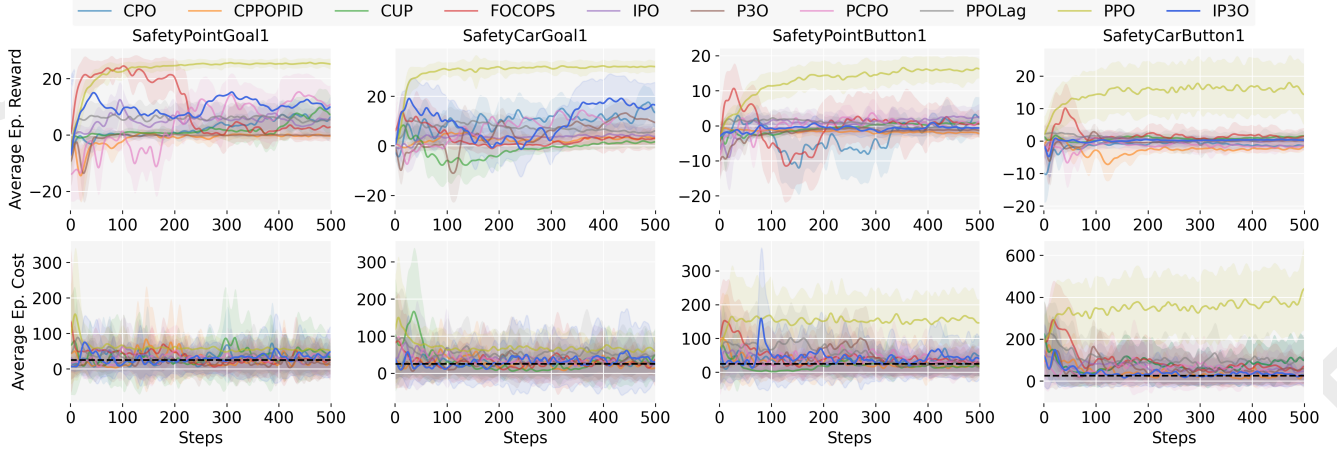


Figure 5: Comparison with the baselines using the Safety Gymnasium scenarios. Our method is marked as IP3O (blue). The dashed line indicates the cost constraint limit.

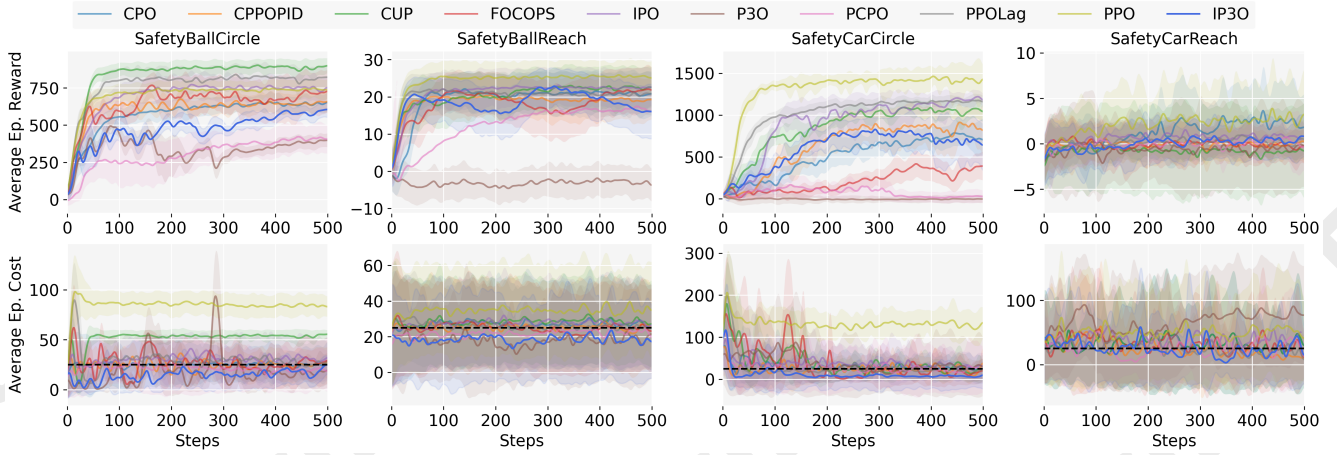


Figure 6: Comparison with the baselines using the Bullet Safety Gymnasium scenarios. Our method is marked as IP3O (blue). The dashed line indicates the cost constraint limit.

this slightly reduces reward maximization, demonstrating that our algorithm prioritizes constraint satisfaction when necessary. Most of the algorithms struggle to find a consistently safe policy owing to oscillations in the loss function near the constraint boundary.

The constraint violations for all the above scenarios are summarized in Figure 7. From the figure it can be seen that IP3O has the lowest cost violations overall, compared to the baselines for the MuJoCo velocity and Bullet Safety Gym environments; and marginally above the best performance in Safety Gym, since $\alpha = 0.1$.




Figure 7: Summary of constraint violations across all environments.




Figure 8: Evaluation results on the MetaDrive simulator

5.4 Multi-Agent Scenarios

The safety objective solved here is very relevant in the multi-agent automotive domain that consist of innate safety aspects. Our proposed approach is extendable to multi-agent scenarios. To demonstrate this, we evaluate it on cooperative driving tasks within the MetaDrive simulator [Li *et al.*, 2022], a lightweight and realistic platform designed for multi-agent decentralized reward settings. To train the policy using a cooperative feedback, the value error was estimated using the

summation of the decentralized rewards. MetaDrive provides challenging environments for controlling multiple vehicles in predefined driving scenarios, where agents must cooperate to maximize rewards by avoiding safety violations such as crashes or going off-road. Additional environment details are provided in the Supplementary Material.

Training spanned 2000 episodes, using an on-policy replay buffer. We benchmarked our approach against multi-agent versions of PPO, namely MAPPO [Yu *et al.*, 2022] and MAP3O [Zhang *et al.*, 2022]. Each trained policy was evaluated over 10 episodes per evaluation step. As shown in Figure 8, our algorithm, marked as MAIP3O (Multi-Agent IP3O), achieves comparable performance to baseline algorithms in terms of collective reward across all agents (vehicles). We focus on reward maximization as the primary objective, with safety integrated to facilitate reward optimization.

In the following sub-section we discuss the effect of some of our hyper-parameters on the outcome with respect to return and constraint satisfaction.

5.5 Ablation Studies

We conduct ablation studies to evaluate the effect of hyper-parameters on the performance of our approach, focusing on two key parameters: the α hyper-parameter, and the cost limit d . These experiments are carried out across different environments to provide a comprehensive understanding of how variations in these hyper-parameters impact performance.




Figure 9: Ablation experiments using α on the Half-Cheetah and the Humanoid environments.

The effect of α is studied in the Half-Cheetah and Humanoid tasks from the MuJoCo Safety Velocity benchmark. For these experiments, the cost limit d is fixed at 25. As shown in Figure 9, setting $\alpha = 0.5$ achieves optimal constraint satisfaction while maintaining a reasonable reward. Reducing α to 0.1, although resulting in higher rewards, leads to frequent constraint violations, as shown in Figure 9. On the other hand, increasing α beyond 0.5 does not significantly improve constraint satisfaction in environments where reward and constraints are intrinsically tied to velocity. In scenarios with tight safety constraints, prioritizing constraint satisfaction with higher α values is beneficial. However, in less

safety-critical tasks, relaxing the safety weight may allow for better reward maximization.

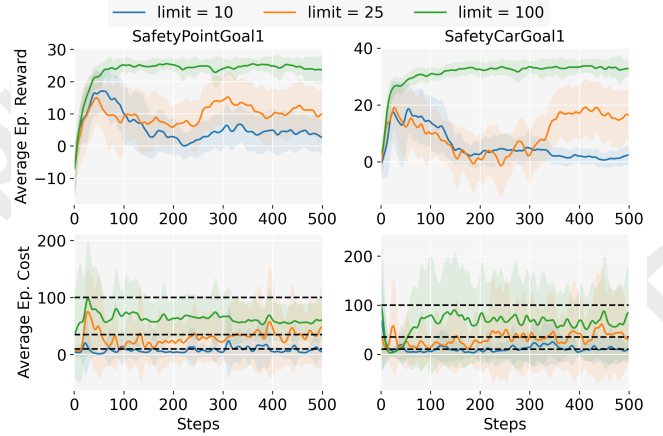


Figure 10: Ablation experiments using the cost limit d hyper-parameter on the PointGoal1 and the CarGoal1 environments.

In addition, we also analyze the robustness of IP3O across varying cost limit levels d in the PointGoal1 and CarGoal1 environments from the Safety Gymnasium benchmark, to better highlight the effect of cost limit on the policy learning. The results, presented in Figure 10, demonstrate that our approach effectively learns constraint satisfying policies, $\pi \in \Pi_C$, for a range of d values while maximizing rewards. This highlights the flexibility of our method in adapting to diverse safety requirements. Overall, these results underline the importance of careful tuning of hyper-parameters to achieve the desired balance between safety and reward.

6 Conclusion

In this paper, we introduced Incrementally Penalized Proximal Policy Optimization (IP3O), a novel policy optimization algorithm designed to solve safe RL tasks for combining constraint satisfaction with reward maximization. IP3O demonstrates adaptability to diverse safety requirements. Our experiments demonstrated the efficacy of IP3O across benchmark environments in safe RL, such as MuJoCo Safety Velocity, Safety Gymnasium, and Bullet Safety Gymnasium. The results highlight its superior ability to balance safety and reward compared to state-of-the-art approaches. Additionally, the algorithm’s scalability to multi-agent settings, as shown in MetaDrive scenarios, shows its potential for real-world applications like autonomous driving and robotics.

IP3O can be tailored to different safety-critical tasks as shown in the ablation studies that validate the impact of key hyper-parameters such as α and constraint limits. These findings demonstrate the algorithm’s practicality for safe RL tasks. In conclusion, IP3O provides a framework for advancing safe RL research through practical applications and theoretical guarantees. Future work could explore its scalability to larger multi-agent systems and further analyze its theoretical properties under varying safety thresholds.

Acknowledgements

The authors acknowledge partial support from AI4CPS IIT Kharagpur grant no TRP3RDTR001 for this work.

References

- [Achiam *et al.*, 2017] Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. Constrained policy optimization. In *International Conference on Machine Learning*, pages 22–31. PMLR, 2017.
- [Altman, 2021] Eitan Altman. *Constrained Markov decision processes*. Routledge, 2021.
- [Barron, 2017] Jonathan T Barron. Continuously differentiable exponential linear units. *arXiv preprint arXiv:1704.07483*, 2017.
- [Chow *et al.*, 2018] Yinlam Chow, Mohammad Ghavamzadeh, Lucas Janson, and Marco Pavone. Risk-constrained reinforcement learning with percentile risk criteria. *Journal of Machine Learning Research*, 18(167):1–51, 2018.
- [Clevert, 2015] Djork-Arné Clevert. Fast and accurate deep network learning by exponential linear units (elus). *arXiv preprint arXiv:1511.07289*, 2015.
- [Dai *et al.*, 2023] Juntao Dai, Jiaming Ji, Long Yang, Qian Zheng, and Gang Pan. Augmented proximal policy optimization for safe reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 7288–7295, 2023.
- [Ding *et al.*, 2020] Dongsheng Ding, Kaiqing Zhang, Tamer Basar, and Mihailo Jovanovic. Natural policy gradient primal-dual method for constrained markov decision processes. *Advances in Neural Information Processing Systems*, 33:8378–8390, 2020.
- [Fernandez-Llorca and Gómez, 2023] David Fernandez-Llorca and Emilia Gómez. Trustworthy artificial intelligence requirements in the autonomous driving domain. *Computer*, 56(2):29–39, 2023.
- [Fisac *et al.*, 2018] Jaime F Fisac, Anayo K Akametalu, Melanie N Zeilinger, Shahab Kaynama, Jeremy Gillula, and Claire J Tomlin. A general safety framework for learning-based control in uncertain robotic systems. *IEEE Transactions on Automatic Control*, 64(7):2737–2752, 2018.
- [Gao *et al.*, 2024] Shiqing Gao, Jiaxin Ding, Luoyi Fu, Xinning Wang, and Chenghu Zhou. Exterior penalty policy optimization with penalty metric network under constraints. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24*, 2024.
- [Gronauer, 2022] Sven Gronauer. Bullet-safety-gym: A framework for constrained reinforcement learning. 2022.
- [Ji *et al.*, 2023] Jiaming Ji, Borong Zhang, Jiayi Zhou, Xuehai Pan, Weidong Huang, Ruiyang Sun, Yiran Geng, Yifan Zhong, Josef Dai, and Yaodong Yang. Safety gymnasium: A unified safe reinforcement learning benchmark. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023.
- [Ji *et al.*, 2024] Jiaming Ji, Jiayi Zhou, Borong Zhang, Juntao Dai, Xuehai Pan, Ruiyang Sun, Weidong Huang, Yiran Geng, Mickel Liu, and Yaodong Yang. Omnisafe: An infrastructure for accelerating safe reinforcement learning research. *Journal of Machine Learning Research*, 25(285):1–6, 2024.
- [Jiang and Li, 2016] Nan Jiang and Lihong Li. Doubly robust off-policy value evaluation for reinforcement learning. In *International Conference on Machine Learning*, pages 652–661. PMLR, 2016.
- [Kakade and Langford, 2002] Sham Kakade and John Langford. Approximately optimal approximate reinforcement learning. In *Proceedings of the Nineteenth International Conference on Machine Learning*, pages 267–274, 2002.
- [Kiran *et al.*, 2021] B Ravi Kiran, Ibrahim Sobh, Victor Talpaert, Patrick Mannion, Ahmad A Al Sallab, Senthil Yogamani, and Patrick Pérez. Deep reinforcement learning for autonomous driving: A survey. *IEEE Transactions on Intelligent Transportation Systems*, 23(6):4909–4926, 2021.
- [Levine *et al.*, 2016] Sergey Levine, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. End-to-end training of deep visuomotor policies. *Journal of Machine Learning Research*, 17(39):1–40, 2016.
- [Li *et al.*, 2022] Quanyi Li, Zhenghao Peng, Lan Feng, Qihang Zhang, Zhenghai Xue, and Bolei Zhou. Metadrive: Composing diverse driving scenarios for generalizable reinforcement learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3):3461–3475, 2022.
- [Liu *et al.*, 2020] Yongshuai Liu, Jiaxin Ding, and Xin Liu. Ipo: Interior-point policy optimization under constraints. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 4940–4947, 2020.
- [McNamara, 2016] Steven McNamara. The law and ethics of high-frequency trading. *Minn. JL Sci. & Tech.*, 17:71, 2016.
- [Ono *et al.*, 2015] Masahiro Ono, Marco Pavone, Yoshiaki Kuwata, and J Balaram. Chance-constrained dynamic programming with application to risk-aware robotic space exploration. *Autonomous Robots*, 39:555–571, 2015.
- [Ray *et al.*, 2019] Alex Ray, Joshua Achiam, and Dario Amodei. Benchmarking safe exploration in deep reinforcement learning. *arXiv preprint arXiv:1910.01708*, 7(1):2, 2019.
- [Schulman *et al.*, 2017] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [Schulman, 2015] John Schulman. Trust region policy optimization. *arXiv preprint arXiv:1502.05477*, 2015.
- [Stooke *et al.*, 2020] Adam Stooke, Joshua Achiam, and Pieter Abbeel. Responsive safety in reinforcement learning by pid lagrangian methods. In *International Con-*

ference on Machine Learning, pages 9133–9143. PMLR, 2020.

[Sutton and Barto, 2018] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

[Tessler *et al.*, 2018] Chen Tessler, Daniel J Mankowitz, and Shie Mannor. Reward constrained policy optimization. *International Conference on Learning Representations*, 2018.

[Yang *et al.*, 2020] Tsung-Yen Yang, Justinian Rosca, Karthik Narasimhan, and Peter J Ramadge. Projection-based constrained policy optimization. *International Conference on Learning Representations*, 2020.

[Yang *et al.*, 2022] Long Yang, Jiaming Ji, Juntao Dai, Linrui Zhang, Binbin Zhou, Pengfei Li, Yaodong Yang, and Gang Pan. Constrained update projection approach to safe policy optimization. *Advances in Neural Information Processing Systems*, 35:9111–9124, 2022.

[Yu *et al.*, 2019] Ming Yu, Zhuoran Yang, Mladen Kolar, and Zhaoran Wang. Convergent policy optimization for safe reinforcement learning. *Advances in Neural Information Processing Systems*, 32, 2019.

[Yu *et al.*, 2021] Chao Yu, Jiming Liu, Shamim Nemati, and Guosheng Yin. Reinforcement learning in healthcare: A survey. *ACM Computing Surveys (CSUR)*, 55(1):1–36, 2021.

[Yu *et al.*, 2022] Chao Yu, Akash Velu, Eugene Vinitzky, Jiaxuan Gao, Yu Wang, Alexandre Bayen, and Yi Wu. The surprising effectiveness of ppo in cooperative multi-agent games. *Advances in neural information processing systems*, 35:24611–24624, 2022.

[Zhang *et al.*, 2020] Yiming Zhang, Quan Vuong, and Keith Ross. First order constrained optimization in policy space. *Advances in Neural Information Processing Systems*, 33:15338–15349, 2020.

[Zhang *et al.*, 2022] Linrui Zhang, Li Shen, Long Yang, Shixiang Chen, Bo Yuan, Xueqian Wang, and Dacheng Tao. Penalized proximal policy optimization for safe reinforcement learning. In *International Joint Conference on Artificial Intelligence*, 2022.

[Zhang *et al.*, 2023] Linrui Zhang, Qin Zhang, Li Shen, Bo Yuan, Xueqian Wang, and Dacheng Tao. Evaluating model-free reinforcement learning toward safety-critical tasks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 15313–15321, 2023.