
Central Path Proximal Policy Optimization

Nikola Milosevic

Max Planck Institute for Human Cognitive and Brain Sciences, Leipzig
Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI), Dresden/Leipzig
nmilosevic@cbs.mpg.de

Johannes Müller

Institut für Mathematik, Technische Universität Berlin, 10623 Berlin, Germany

Nico Scherf

Max Planck Institute for Human Cognitive and Brain Sciences, Leipzig
Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI), Dresden/Leipzig

Abstract

In constrained Markov decision processes, enforcing constraints during training is often thought of as decreasing the final return. Recently, it was shown that constraints can be incorporated directly into the policy geometry, yielding an optimization trajectory close to the central path of a barrier method, which does not compromise final return. Building on this idea, we introduce Central Path Proximal Policy Optimization (C3PO), a simple modification of the PPO loss that produces policy iterates, that stay close to the central path of the constrained optimization problem. Compared to existing on-policy methods, C3PO delivers improved performance with tighter constraint enforcement, suggesting that central path-guided updates offer a promising direction for constrained policy optimization.

1 Introduction

Reinforcement learning (RL) has demonstrated impressive capabilities across a wide range of domains, yet real-world applications increasingly demand more than just reward maximization. In many real-world high-stakes environments agents must also avoid violating domain-specific safety or resource constraints. This motivates the study of constrained Markov decision processes (CMDPs), an extension of the standard RL framework that imposes expected cost constraints alongside the goal of reward maximization [Altman, 1999]. By treating feasibility and reward objectives separately, CMDPs provide a principled framework for specifying agent behavior in complex environments.

CMDPs are especially relevant in deep reinforcement learning settings, where the design of reward functions that lead to safe behavior is difficult. Prior work has emphasized the importance of explicit constraint modeling in reinforcement learning for safe exploration [Ray et al., 2019] and complex task specification [Roy et al., 2022], where constraints can be more natural and easier to design, including finetuning LLMs for harmlessness [Dai et al., 2024].

Despite their relatively low sample efficiency, model-free on-policy algorithms continue to play a foundational role in constrained RL. They offer conceptual clarity, support rigorous theoretical analysis, and provide strong baselines for studying the balance between performance and constraint satisfaction. As the field moves toward more scalable and sample-efficient approaches, insights developed in the on-policy setting remain central to both algorithm design and our broader understanding of safe learning, such as the policy improvement guarantees and constraint violation bounds introduced by Achiam et al. [2017].

In this context, there is a growing need for simple, scalable, and effective algorithms for solving CMDPs, ideally with properties similar to widely used algorithms such as proximal policy optimization (PPO; Schulman et al. [2017b]). PPO’s robustness, ease of implementation, and scalability have made it the method of choice in many deep RL and RLHF pipelines [Ouyang et al., 2022]. We aim to extend these strengths to the constrained setting by developing an algorithm that shares PPO’s practical benefits while enforcing constraints in a principled CMDP framework. Specifically, we seek to achieve high final reward while approximately satisfying constraints, at least at convergence.

To frame this problem, we distinguish between two commonly conflated settings in constrained RL: (i) *safe exploration*, where constraints must be satisfied throughout training, and (ii) *safe convergence*, where only the final policy is required to satisfy the constraints. Much of the literature has focused on the former, motivated by safety-critical applications in the real world. The dominant approach in this setting is model-based safe RL, which can provide strong safety and stability guarantees [Berkenkamp et al., 2017, As et al., 2025]. In contrast, safe convergence reflects settings like simulation-based training or alignment finetuning [Dai et al., 2024]. Typically, ensuring safety *during* training is considered to decrease the final performance achieved by an algorithm. We show the contrary and present an algorithm that exhibits strict feasibility during training as well as reliable feasibility and high return at convergence.

In nonlinear CMDPs, the constraint surface is typically curved and nonconvex in policy space. Converging prematurely or oscillating near the constraint boundary during training can lead to unreliable constraint satisfaction at convergence. Furthermore, it can lead the iterates to local optima that satisfy the constraints but fail to achieve high reward, see Figure 1. Penalty and barrier methods address this problem by maintaining a feasible trajectory toward the constraint surface, yielding feasible solutions more reliably. However, barrier methods introduce bias [Müller and Cayci, 2024], meaning the optimization problem obtained by adding a barrier penalty does not have the same solution set as the original problem, which can lead to degraded reward in policy optimization [Milosevic et al., 2025]. Barrier methods either require careful tuning or an interior point approach [Liu et al., 2020] to avoid harming reward performance.




Figure 1: Pictorial visualization of the central path argument from the main text. While a wide range of methods technically converge to an optimal feasible solution in the linear programming formulation of finite CMDPs (left), in the function approximation setting (right), approaching the constraint surface too early may result in higher sensitivity to local optima.

The recently proposed C-TRPO [Milosevic et al., 2025] addresses these challenges by combining the strengths of trust-region and barrier methods by deriving a barrier-inspired trust-region formulation using strictly feasible trust regions. This results in an algorithm that acts like a barrier method with an adaptively receding barrier, introducing no regularization bias as a result. This is achieved by producing policies, which are close to the regularization path obtained by altering the regularization strength, which is commonly known as the *central path* [Boyd and Vandenberghe, 2004].

C-TRPO’s scalability is limited due to computational overhead introduced by the TRPO-inspired update, and the update is defined only in the feasible set. To address the need for a simple and scalable CMDP solver, we propose a proximal version of C-TRPO. It also follows the central path, and consequently we call it *Central Path Proximal Policy Optimization (C3PO)*. C3PO is a minibatch-based method that approximates the C-TRPO update using an exact penalty formulation, combining the simplicity and efficiency of PPO-style updates with the feasible geometry of central path methods. At its core, C3PO leverages the central path property of natural policy gradients to gradually guide the policy toward the constraint surface without inducing oscillations or premature convergence.

2 Background

We consider the infinite-horizon discounted constrained Markov decision process (CMDP) and refer the reader to Altman [1999] for a general treatment. The CMDP is given by the tuple $\mathcal{M} \cup \mathcal{C}$, consisting of a finite MDP \mathcal{M} and a set of constraints \mathcal{C} . The finite MDP $\mathcal{M} = \{\mathcal{S}, \mathcal{A}, P, r, \mu, \gamma\}$ is defined by a finite state-space \mathcal{S} , a finite action-space \mathcal{A} , a transition kernel $P: \mathcal{S} \times \mathcal{A} \rightarrow \Delta_{\mathcal{S}}$, an extrinsic reward function $r: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, an initial state distribution $\mu \in \Delta_{\mathcal{S}}$, and a discount factor $\gamma \in [0, 1)$. The space $\Delta_{\mathcal{S}}$ is the set of categorical distributions over \mathcal{S} . Further, $\mathcal{C} = \{(c_i, b_i)\}_{i=1}^m$ is a set of m constraints, where $c_i: \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ are the cost functions and $b_i \in \mathbb{R}$ are the cost thresholds. An agent interacts with the CMDP by selecting a policy $\pi \in \Pi$ and collecting trajectories $\tau = (s_0, a_0, \dots, s_T)$. Let $R(\tau) = \sum_{t=0}^{\infty} \gamma^t r(s_t, a_t)$. Given π , the value function, action-value function, and advantage function associated with the reward function r are defined respectively as

$$V_r^\pi(s) := (1 - \gamma) \mathbb{E}_{\tau \sim \pi} [R(\tau) | s_0 = s], \quad Q_r^\pi(s, a) := (1 - \gamma) \mathbb{E}_{\tau \sim \pi} [R(\tau) | s_0 = s, a_0 = a],$$

and

$$A_r^\pi(s, a) := Q_r^\pi(s, a) - V_r^\pi(s).$$

The expectations are taken over trajectories of the Markov process, meaning with respect to the initial distribution $s_0 \sim \mu$, the policy $a_t \sim \pi(\cdot | s_t)$ and the state transition $s_{t+1} \sim P(\cdot | s_t, a_t)$. $V_{c_i}^\pi(s)$, $Q_{c_i}^\pi(s, a)$ and $A_{c_i}^\pi(s, a)$ are defined analogously for the i -th cost c_i instead of r .

Constrained reinforcement learning addresses the optimization problem

$$\text{maximize}_{\pi \in \Pi} R(\pi) \quad \text{subject to} \quad C_i(\pi) \leq d_i \quad (1)$$

for all $i = 1, \dots, m$, where $R(\pi)$ is the expected value under the initial state distribution $R(\pi) := \mathbb{E}_{s \sim \mu} [V_r^\pi(s)]$ and $C_i(\pi) := \mathbb{E}_{s \sim \mu} [V_{c_i}^\pi(s)]$.

Every stationary policy π induces the occupancy measures $\rho_\pi(s) := (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t \mathbb{P}_\pi(s_t = s)$, and $\rho_\pi(s, a) := \rho_\pi(s) \pi(a | s)$ which indicate the relative frequencies of visiting a state(-action) pair, discounted by how far the event lies in the future. The classical, linear programming (LP) approach to solving finite CMDPs [Altman, 1999], reformulates problem 1 as

$$\text{maximize}_{\rho \in \mathcal{K}} \sum_{s, a} \rho(s, a) r(s, a) \quad \text{subject to} \quad \sum_{s, a} \rho(s, a) c_i(s, a) \leq d_i \quad (2)$$

which can be solved using LP solution methods to obtain an optimal occupancy measure ρ^* . Here, \mathcal{K} is a set of linear constraints that ρ_π must satisfy [Kallenberg, 1994, Mei et al., 2020], sometimes referred to as the Bellman flow equations. Finally, an optimal policy can be extracted by conditioning $\pi^*(a | s) = \rho^*(s, a) / \sum_{a'} \rho^*(s, a')$.

In the function approximation setting, approach 2 is not applicable, which has prompted a large body of research in on-policy policy optimization methods. However, it can be leveraged to derive general constrained RL algorithms [Milosevic et al., 2025]. In the analysis of on-policy methods (including for standard MDPs) the *policy advantage* plays an important role. In CMDPs, the policy advantages are defined as

$$\mathbb{A}_r^{\pi_k}(\pi) = \sum_{s, a} \rho_{\pi_k}(s) \pi(a | s) A_r^{\pi_k}(s, a) \quad \text{and} \quad \mathbb{A}_c^{\pi_k}(\pi) = \sum_{s, a} \rho_{\pi_k}(s) \pi(a | s) A_c^{\pi_k}(s, a). \quad (3)$$

They play an important role in policy optimization, as they approximate the performance difference between two nearby policies with respect to the reward $\mathbb{A}_r^{\pi_k}(\pi) \approx R(\pi) - R(\pi_k)$ if $\pi_k \approx \pi$ and analogously for the cost function.

2.1 Policy Optimization Methods for Constrained Reinforcement Learning

In the following, we review relevant prior constrained policy optimization methods, thereby focusing on a single constraint to reduce notational clutter. However, all mentioned methods are trivial to extend to multiple constraints.

Constrained Policy Optimization (CPO) Constrained policy optimization (CPO) is a modification of trust region policy optimization (TRPO; Schulman et al. [2017a]), where the classic trust region is intersected with the set of safe policies [Achiam et al., 2017]. At each iteration k , the policy of the next iteration π_{k+1} is obtained through the solution of

$$\max_{\pi \in \Pi} \mathbb{A}_r^{\pi_k}(\pi) \text{ s.t. } \bar{D}_{\text{KL}}(\pi, \pi_k) \leq \delta \text{ and } C(\pi_k) + \mathbb{A}_c^{\pi_k}(\pi) \leq d. \quad (4)$$

where $\bar{D}_{\text{KL}}(\pi, \pi_k) = \sum_{s,a} \rho_{\pi_k}(s) D_{\text{KL}}[\pi(\cdot|s)|\pi_k(\cdot|s)]$ and $C(\pi_k) + \mathbb{A}_c^{\pi_k}(\pi)$ is an estimate for $C(\pi)$, see Kakade and Langford [2002], Schulman et al. [2017a], Achiam et al. [2017].

Penalized Proximal Policy Optimization (P3O/P2BPO) Solving the constrained optimization problem equation 4 is difficult to scale up to more challenging tasks and larger model sizes, as it relies on the arguably sample inefficient TRPO update. To circumvent this Zhang et al. [2022] proposed a Constrained RL algorithm derived from the relaxed penalized problem

$$\max_{\pi \in \Pi} \mathbb{A}_r^{\pi_k}(\pi) - \lambda \max\{0, C(\pi_k) + \mathbb{A}_c^{\pi_k}(\pi) - d\}, \text{ s.t. } \bar{D}_{\text{KL}}(\pi, \pi_k) \leq \delta. \quad (5)$$

The appeal of this reformulation is that one can obtain an unconstrained problem that gives the same solution set for λ chosen large enough [Zhang et al., 2022] and by employing a PPO-like loss. A similar approach was taken by Dey et al. [2024], where $\max\{0, \cdot\}$ was replaced with a softplus.

Constrained Trust Region Policy Optimization (C-TRPO) Where equation 4 incorporates constraints by intersecting the trust region with the set of safe policies, an alternative approach was taken by Milosevic et al. [2025] where the geometry was modified such that the resulting trust region automatically consists of safe policies. To this end, C-TRPO proceeds as TRPO but with the usual divergence augmented by a barrier term, meaning

$$\max_{\pi \in \Pi} \mathbb{A}_r^{\pi_k}(\pi) \text{ s.t. } \bar{D}_{\text{KL}}(\pi, \pi_k) + \beta D_{\text{B}}(\pi, \pi_k) \leq \delta. \quad (6)$$

where we'll refer to

$$D_{\text{B}}(\pi, \pi_k) = \frac{b - \mathbb{A}_c^{\pi_k}(\pi)}{b} - \log\left(\frac{b - \mathbb{A}_c^{\pi_k}(\pi)}{b}\right) - 1, \text{ for } b > 0, \text{ else } \infty \quad (7)$$

as the *barrier divergence*, β is a positive safety parameter, and $b = d - C(\pi_k)$ is the *cost budget*.

This update is justified by the general theory of Bregman divergences and the theory of convex programs. It has desirable theoretical properties and results in state-of-the-art performance compared to other on-policy CMDP algorithms. We refer the reader to Milosevic et al. [2025] and Appendix A for detailed discussions.

Other methods So far we have focused on model-free, direct policy optimization methods, specifically trust-region and PPO-penalty based ones. However, it is important to note that model-based approaches, e.g. Berkenkamp et al. [2017], As et al. [2025], are also attractive due to their stability and safety guarantees, but require learning a model, which is not always feasible. Lagrangian methods are a widely adopted approach, where the CMDP is formulated as a primal-dual optimization problem. Optimizing the dual variable with stochastic gradient descent is a popular baseline [Achiam et al., 2017, Ray et al., 2019, Chow et al., 2019, Stooke et al., 2020]. However, a naively optimized dual variable may cause oscillations and overshoot, and analyzing more nuanced update strategies is a subject of current research interest [Sohrabi et al., 2024]. More recently, log-barrier approaches have attracted considerable research interest [Usmanova et al., 2024, Zhang et al., 2024, Dey et al., 2024] due to their algorithmic simplicity through the fixed penalty, but also due to recent rigorous treatments, see e.g. Ni and Kamgarpour [2024]. However, working with an explicit penalty produces suboptimal policies w.r.t the original constrained MDP. Fixed penalties introduce an additional error, which has to be controlled, see for example Geist et al. [2019], Müller and Cayci [2024] for theory, and Liu et al. [2020] for a practical example of regularization bias.

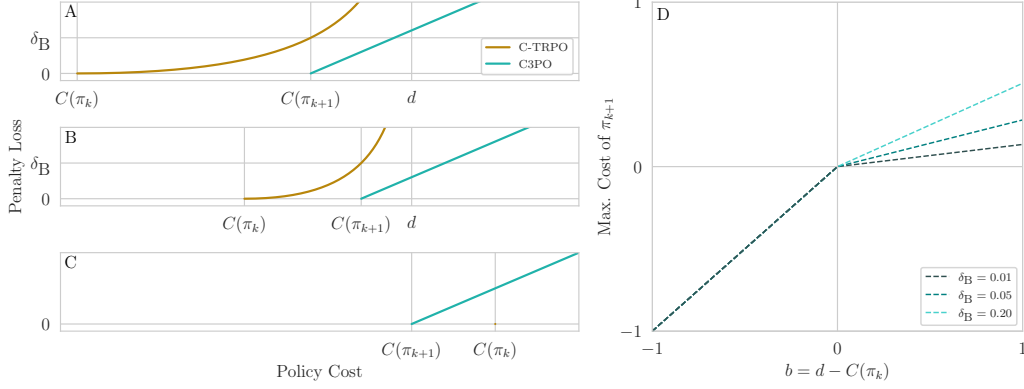


Figure 2: The working principle behind C3PO’s exact penalty approach: As the iterate moves closer towards the constraint (A-C), C3PO’s ReLU-penalty pulls away at a slower rate $0 < w < 1$, mimicking C-TRPO’s barrier divergence. This rate is defined as a function of δ_B (D), see main text. While C-TRPO’s barrier penalty is undefined if either $d \leq C(\pi_k)$ or $d \leq C(\pi_{k+1})$, C3PO’s ReLU-penalty is defined everywhere.

3 Central Path Proximal Policy Optimization

C-TRPO has desirable properties but the practical algorithm 1) scales poorly and is sample-inefficient due to its reliance on the TRPO algorithm and 2) relies on a recovery mechanism (reward-free cost minimization + hysteresis), since the update is not defined if π_k is outside the feasible set.

We propose a first-order approximation of C-TRPO that approximates its central path by solving surrogate optimization problems with the same solution set as C-TRPO’s update on every iteration. In addition, by employing an exact penalty approach, it allows unsafe policies during training, also enabling less strict exploration strategies within the safe convergence regime.

C3PO Update Let us consider a slight modification of C-TRPO’s update, which is constrained with the KL and Barrier constraints separately, since they can be approximated using different methods which result in different precisions, i.e. we consider

$$\max_{\pi \in \Pi} \mathbb{A}_r^{\pi_k}(\pi) \text{ s.t. } D_B(\pi, \pi_k) \leq \delta_B \text{ and } \bar{D}_{KL}(\pi, \pi_k) \leq \delta_{KL}. \quad (8)$$

Note that this is a subtly different problem than that posed by C-TRPO, but δ_{KL} and δ_B can always be chosen to include the feasible set entirely in C-TRPO’s feasible set for a given δ . Instead of solving this constrained problem directly, we consider the penalized problem given by

$$\max_{\pi \in \Pi} \mathbb{A}_r^{\pi_k}(\pi) - \kappa \max\{0, D_B(\pi, \pi_k) - \delta_B\} \text{ s.t. } \bar{D}_{KL}(\pi, \pi_k) \leq \delta_{KL}. \quad (9)$$

Theorem 3.1 (Exactness). *Let λ be the Lagrange multiplier vector for the optimizer of Equation 8. Then for $\kappa \geq |\lambda|$ the solution sets of problem Equation 8 and problem Equation 9 agree.*

Proof. Note that the problem Equation 8 is concave-convex in π . Hence, this is a special case of the general exactness result Theorem B.1. \square

C3PO Algorithm The update Equation 9 is still undefined outside the feasible set of the barrier divergence constraint. Since we use the barrier divergence only to define the feasible solution set of the update, we can replace it with another function, as long as it defines the same feasible set. More precisely, this can be achieved with an equivalent linear constraint that is zero where $D_B(\pi, \pi_k) = \delta_B$ for positive cost advantages. The C3PO algorithm approximates update 9 as

$$\max_{\pi \in \Pi} \mathbb{A}_r^{\pi_k}(\pi) - \kappa \max\{0, \mathbb{A}_c^{\pi_k}(\pi) - \min\{b, w \cdot b\}\} \text{ s.t. } \bar{D}_{KL}(\pi, \pi_k) \leq \delta_{KL}. \quad (10)$$

where $0 < w < 1$.

Proposition 3.2 (Positive Exactness). *For $0 \leq \mathbb{A}_c^{\pi_k}(\pi) < d - C(\pi_k)$, there exist w and δ_B for which the solution sets of problems 8, 9 and 10 agree.*

The new update expresses the same constraint using a linear ReLU-penalty. The rate w is a new hyper-parameter and we refer to Appendix B for a proof of Proposition 3.2. Since the original problem’s penalty function is not defined outside the interior of the feasible set, we must handle the case $C(\pi_k) \geq d$ differently, which is taken care of by the $\min(b, \cdot)$ term: For $b < 0$, problem 10 reduces to the P3O [Zhang et al., 2022] objective Eq. 5. Finally, the additional KL-constraint is approximated as in PPO [Schulman et al., 2017b]. The resulting loss only consists of the PPO loss and an additional loss term which is a function of the policy cost advantage estimate. Let $r(\theta) = \frac{\pi_\theta(a|s)}{\pi_k(a|s)}$ denote the likelihood ratio of the optimized and last behavior policies and let

$$\alpha_{\text{clipped}}(\theta) = \mathbb{E}_{s,a \sim \rho_k} \left[\max \left(r(\theta) \hat{A}_c(s, a), \text{clip}(r(\theta), 1 - \epsilon, 1 + \epsilon) \hat{A}_c(s, a) \right) \right]. \quad (11)$$

The C3PO loss is

$$L^{\text{C3PO}}(\theta) = \text{ReLU}(\alpha_{\text{clipped}}(\theta) - \min\{b, w \cdot b\}). \quad (12)$$

The penalty coefficient remains a hyperparameter, which can be flexibly scheduled to solve CMDPs in the safe convergence regime, as shown in Section 4, where we use a linear schedule to achieve high final performance across multiple tasks. The final method is summarized in Algorithm 1.

Algorithm 1 C3PO (deviation from PPO in green)

Require: Initial policy π_0 and value functions \hat{V}_r, \hat{V}_{c_i} , thresholds d_i , scheduled penalty κ_k , rate w

- 1: **for** $k = 0, 1, 2, \dots$ **do**
- 2: Collect trajectory data $\mathcal{D} = \{s_0, a_0, r_0, c_0, \dots\}$ by running π_k
- 3: Estimate reward advantage \hat{A}_t^r and cost advantages $\hat{A}_t^{c_i}$ using GAE- λ [Schulman et al., 2016]
- 4: Update policy π_{k+1} by minimizing $L^{\text{PPO}} + \kappa_k L^{\text{C3PO}}$ (Equation 12)
- 5: Update value function estimates $\hat{V}_r^{\pi_{k+1}}$ and $\hat{V}_{c_i}^{\pi_{k+1}}$ by regression
- 6: **end for**

Relation to other PPO-Penalty methods C3PO is a superset of P3O [Zhang et al., 2020]. More precisely, if we set $w = 1$ in C3PO, we obtain the P3O loss exactly. Further, C3PO is conceptually similar to P2BPO [Dey et al., 2024], in using a more conservative version of the P3O loss, but C3PO does not use a penalty with a fixed location at the constraint, but a moving penalty which recedes as the iterate gets closer to the constraint. This allows C3PO to approach the optimal feasible solution without regularization bias.

4 Computational Experiments

To evaluate our approach, we conduct experiments aimed at testing the benefits of using central path approximation as a design principle for constrained policy optimization algorithms. We benchmark C3PO against a range of representative constrained reinforcement learning baselines. We include methods from three major algorithmic families: penalty-based methods (P3O, P2BPO), Lagrangian methods (PPO-Lag, CPPO-PID), and trust-region methods (CPO, C-TRPO).

Conceptually, *penalty-based methods*, especially algorithms that augment the PPO loss with a penalty, like P3O [Zhang et al., 2022] and P2BPO [Dey et al., 2024], are closest to our approach. Like C3PO, those penalize constraint violations directly in the policy gradient loss using a ReLU-penalty. *Lagrangian methods* maintain dual variables to enforce constraints adaptively. PPO-Lagrangian [Ray et al., 2019] applies this principle to the PPO algorithm, forming a loss which is similar to C3PO’s. For completeness, we consider CPPO-PID [Stooke et al., 2020] as a more recent Lagrangian baseline. Finally, *trust region methods*, such as CPO [Achiam et al., 2017] and C-TRPO [Milosevic et al., 2025], use trust regions and constrained updates to maintain stable reward improvement and feasibility throughout training. They do not aim for scalability, but form strong baselines on the benchmarks.

We benchmark the algorithms on 4 locomotion tasks and 4 navigation tasks from Safety Gymnasium [Ji et al., 2023], as done by Milosevic et al. [2025]. For the baseline algorithms, we use the hyperparameters reported by Ji et al. [2023], and for P3O and C-TRPO the recommended parameters in




Figure 4: Example of improved performance through approximating the central path: Lagrangian methods tend to converge less reliably towards a safe policy and oscillate around the constraint. This does not yield a higher final reward. Instead, staying feasible from early on in training seems to have a positive effect on final reward.

Zhang et al. [2022] and Milosevic et al. [2025] respectively. For C3PO we use $\kappa = 30.0$ and $w = 0.05$. Each algorithm is trained on each task for 10 million steps with a cost threshold of 25.0. Final iterate performance is measured by aggregating over 5 seeds using `rliable` [Agarwal et al., 2021].

The results provide confirmatory evidence for the usefulness of the central path approach. Policies trained with C3PO exhibit a stable progression toward the constrained optimum, maintaining feasibility for most training iterations, see Figure 4. Furthermore, C3PO consistently outperforms prior PPO-style penalty methods in terms of achieved reward, while also adhering more strictly to the specified constraints, see Figure 3. This improved trade-off between reward and feasibility offers additional support for the effectiveness of the central path approach. While C3PO does not outperform trust-region methods across all tasks in the benchmark, it performs well consistently, resulting in high aggregated performance. The full benchmark results table and more examples like Figure 4 are presented in Appendix C.




Figure 3: Aggregated performance using the inter quartile mean (IQM) across 8 tasks from Safety Gymnasium for a subset of algorithms. The algorithms were chosen as the feasible representatives of their respective group.

5 Conclusion

In this work, we use central path approximation as a guiding principle for designing policy optimization methods for constrained RL. We propose C3PO, an algorithm which is obtained through a simple augmentation of the original PPO-loss inspired by the central path approach. Our experimental results support this design principle: Compared to existing PPO-style penalty and Lagrangian methods, C3PO exhibits improved performance with tighter constraint satisfaction, highlighting the benefits of a central path approach in constrained policy optimization.

While the current results are limited to small-scale simulations and simplified settings, such as a single constraint per task, they suggest that central path approximation is a promising design principle for constrained RL algorithms. We hope this encourages further research, following this paradigm. Future directions include the extensions to high-dimensional tasks, theoretical guarantees, and applications such as safety-critical control and LLM fine-tuning.

Acknowledgments and Disclosure of Funding

N. M. and N.S. are supported by BMBF (Federal Ministry of Education and Research) through ACONITE (01IS22065) and the Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI.) Leipzig and by the European Union and the Free State of Saxony through BIOWIN. N.M. is also supported by the Max Planck IMPRS CoNI Doctoral Program.

References

- Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. Constrained Policy Optimization, 2017.
- Rishabh Agarwal, Max Schwarzer, Pablo Samuel Castro, Aaron C Courville, and Marc Bellemare. Deep reinforcement learning at the edge of the statistical precipice. *Advances in Neural Information Processing Systems*, 34, 2021.
- Eitan Altman. *Constrained Markov Decision Processes*. CRC Press, Taylor & Francis Group, 1999. URL <https://api.semanticscholar.org/CorpusID:14906227>.
- Felipe Alvarez, Jérôme Bolte, and Olivier Brahic. Hessian Riemannian gradient flows in convex programming. *SIAM journal on control and optimization*, 43(2):477–501, 2004.
- Yarden As, Bhavya Sukhija, Lenart Treven, Carmelo Sferrazza, Stelian Coros, and Andreas Krause. ActSafe: Active Exploration with Safety Constraints for Reinforcement Learning. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=aKRADWBJ1I>.
- Felix Berkenkamp, Matteo Turchetta, Angela Schoellig, and Andreas Krause. Safe Model-based Reinforcement Learning with Stability Guarantees. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/766ebcd59621e305170616ba3d3dac32-Paper.pdf.
- Dimitri P Bertsekas. Nonlinear programming. *Journal of the Operational Research Society*, 48(3): 334–334, 1997.
- Stephen P Boyd and Lieven Vandenbergh. *Convex optimization*. Cambridge university press, 2004.
- Yinlam Chow, Ofir Nachum, Aleksandra Faust, Edgar Duenez-Guzman, and Mohammad Ghavamzadeh. Lyapunov-based Safe Policy Optimization for Continuous Control, 2019. URL <https://arxiv.org/abs/1901.10031>.
- Robert M Corless, Gaston H Gonnet, David EG Hare, David J Jeffrey, and Donald E Knuth. On the Lambert W function. *Advances in Computational mathematics*, 5:329–359, 1996.
- Josef Dai, Xuehai Pan, Ruiyang Sun, Jiaming Ji, Xinbo Xu, Mickel Liu, Yizhou Wang, and Yaodong Yang. Safe RLHF: Safe reinforcement learning from human feedback. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=TyFrPOKYXw>.
- Sumanta Dey, Pallab Dasgupta, and Soumyajit Dey. P2BPO: Permeable Penalty Barrier-Based Policy Optimization for Safe RL. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 21029–21036, 2024.
- Matthieu Geist, Bruno Scherrer, and Olivier Pietquin. A theory of regularized markov decision processes. In *International Conference on Machine Learning*, pages 2160–2169. PMLR, 2019.
- 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. URL <https://openreview.net/forum?id=WZmlxIuIGR>.
- Sham Kakade and John Langford. Approximately Optimal Approximate Reinforcement Learning. In *Proceedings of the Nineteenth International Conference on Machine Learning, ICML ’02*, page 267–274, San Francisco, CA, USA, 2002. Morgan Kaufmann Publishers Inc. ISBN 1558608737.

- Lodewijk CM Kallenberg. Survey of linear programming for standard and nonstandard Markovian control problems. Part I: Theory. *Zeitschrift für Operations Research*, 40:1–42, 1994.
- Yongshuai Liu, Jiaxin Ding, and Xin Liu. IPO: Interior-Point Policy Optimization under Constraints. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34:4940–4947, 04 2020. doi: 10.1609/aaai.v34i04.5932.
- Jincheng Mei, Chenjun Xiao, Bo Dai, Lihong Li, Csaba Szepesvári, and Dale Schuurmans. Escaping the gravitational pull of softmax. *Advances in Neural Information Processing Systems*, 33:21130–21140, 2020.
- Nikola Milosevic, Johannes Müller, and Nico Scherf. Embedding Safety into RL: A New Take on Trust Region Methods, 2025. URL <https://arxiv.org/abs/2411.02957>.
- Johannes Müller and Semih Cayci. Optimal Rates of Convergence for Entropy Regularization in Discounted Markov Decision Processes. *arXiv preprint arXiv:2406.04163*, 2024.
- Johannes Müller and Guido Montúfar. Geometry and convergence of natural policy gradient methods. *Information Geometry*, pages 1–39, 2023.
- Gergely Neu, Anders Jonsson, and Vicenç Gómez. A unified view of entropy-regularized Markov decision processes. *arXiv preprint arXiv:1705.07798*, 2017.
- Tingting Ni and Maryam Kamgarpour. A safe exploration approach to constrained Markov decision processes. In *ICML 2024 Workshop: Foundations of Reinforcement Learning and Control—Connections and Perspectives*, 2024.
- Jorge Nocedal and Stephen J Wright. *Numerical optimization*. Springer, 1999.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744, 2022.
- Alex Ray, Joshua Achiam, and Dario Amodei. Benchmarking safe exploration in deep reinforcement learning. *arXiv preprint arXiv:1910.01708*, 7(1):2, 2019.
- Julien Roy, Roger Girgis, Joshua Romoff, Pierre-Luc Bacon, and Christopher Pal. Direct Behavior Specification via Constrained Reinforcement Learning, 2022. URL <https://arxiv.org/abs/2112.12228>.
- John Schulman, Philipp Moritz, Sergey Levine, Michael I. Jordan, and Pieter Abbeel. High-dimensional continuous control using generalized advantage estimation. In *International Conference on Learning Representations (ICLR)*, 2016. URL <https://arxiv.org/abs/1506.02438>.
- John Schulman, Sergey Levine, Philipp Moritz, Michael I. Jordan, and Pieter Abbeel. Trust region policy optimization, 2017a.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal Policy Optimization Algorithms, 2017b. URL <https://arxiv.org/abs/1707.06347>.
- Motahareh Sohrabi, Juan Ramirez, Tianyue H. Zhang, Simon Lacoste-Julien, and Jose Gallego-Posada. On PI Controllers for Updating Lagrange Multipliers in Constrained Optimization, 2024. URL <https://arxiv.org/abs/2406.04558>.
- Adam Stooke, Joshua Achiam, and Pieter Abbeel. Responsive Safety in Reinforcement Learning by PID Lagrangian Methods, 2020. URL <https://arxiv.org/abs/2007.03964>.
- Ilnura Usmanova, Yarden As, Maryam Kamgarpour, and Andreas Krause. Log barriers for safe black-box optimization with application to safe reinforcement learning. *Journal of Machine Learning Research*, 25(171):1–54, 2024.
- Baohe Zhang, Yuan Zhang, Lilli Frison, Thomas Brox, and Joschka Bödecker. Constrained Reinforcement Learning with Smoothed Log Barrier Function. *arXiv preprint arXiv:2403.14508*, 2024.

Linrui Zhang, Li Shen, Long Yang, Shixiang Chen, Xueqian Wang, Bo Yuan, and Dacheng Tao. Penalized Proximal Policy Optimization for Safe Reinforcement Learning. In *Thirty-First International Joint Conference on Artificial Intelligence IJCAI-22*, pages 3719–3725, 07 2022. doi: 10.24963/ijcai.2022/517.

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

A Extended Background

A.1 The Geometry of Policy Optimization

Neu et al. [2017] have shown that the policy divergence used to define the trust-region in TRPO Schuman et al. [2017a] can be derived as the Bregman divergence generated by a mirror function on the state-action polytope. TRPO’s mirror function is the negative conditional entropy

$$\Phi_K(\rho) = \sum_{s,a} \rho(s,a) \log \pi_\rho(a|s) \quad (13)$$

which generates

$$D_K(\pi_k || \pi) = \sum_{s,a} \rho_k(s,a) [\log \pi(a|s) - \log \pi_k(a|s)] \quad (14)$$

via the operator

$$D_\Phi(x || y) := \Phi(x) - \Phi(y) - \nabla \Phi(y)^\top (x - y). \quad (15)$$

In general, a trust region update is defined as

$$\pi_{k+1} \in \arg \max_{\pi \in \Pi} \mathbb{A}_r^{\pi_k}(\pi) \quad \text{sbj. to } D_\Phi(\rho_{\pi_k} || \rho_\pi) \leq \delta, \quad (16)$$

where $D_\Phi: \mathcal{K} \times \mathcal{K} \rightarrow \mathbb{R}$ is the *Bregman divergence* induced by a suitably convex function $\Phi: \text{int}(\mathcal{K}) \rightarrow \mathbb{R}$.

A.2 The Safe Geometry Approach

Milosevic et al. [2025] consider mirror functions of the form

$$\Phi_C(\rho) := \Phi_K(\rho) + \sum_i \beta_i \Phi_B(\rho) \quad (17)$$

$$:= \sum_{s,a} \rho(s,a) \log \pi_\rho(a|s) + \sum_{i=1}^m \beta_i \phi \left(b_i - \sum_{s,a} \rho(s,a) c(s,a) \right), \quad (18)$$

where $\rho \in \mathcal{K}_{\text{safe}}$ is a feasible state-action occupancy, Φ_K is the negative conditional entropy, and ϕ is convex. Further, $\phi: \mathbb{R}_{>0} \rightarrow \mathbb{R}$ with $\phi'(x) \rightarrow +\infty$ for $x \searrow 0$. The log-barrier $\phi(x) = -\log(x)$ considered in this work is a possible candidate. In general, the induced divergence takes the form

$$D_C(\rho_1 || \rho_2) = D_K(\rho_1 || \rho_2) + \sum_{i=1}^m \beta_i D_B(\rho_1 || \rho_2) \quad (19)$$

$$= D_K(\rho_1 || \rho_2) + \sum_{i=1}^m \beta_i [\Phi(b_{1;i}) - \Phi(b_{2;i}) + \Phi'(b_{2;i}) C_i(\pi_1) - \Phi'(b_{2;i}) C_i(\pi_2)], \quad (20)$$

where $b_{\pi;i} = d_i - C_i(\pi)$. The corresponding trust-region scheme is

$$\pi_{k+1} \in \arg \max_{\pi \in \Pi} \mathbb{A}_r^{\pi_k}(\pi) \quad \text{sbj. to } D_C(\rho_{\pi_k} || \rho_\pi) \leq \delta. \quad (21)$$

Analogously to the case of unconstrained TRPO, there is a corresponding natural policy gradient scheme:

$$\theta_{k+1} = \theta_k + \epsilon_k G_C(\theta_k)^+ \nabla R(\theta_k), \quad (22)$$

where $G_C(\theta)^+$ denotes an arbitrary pseudo-inverse of the Gramian

$$G_C(\theta)_{ij} = \partial_{\theta_i} \rho_\theta^\top \nabla^2 \Phi_C(\rho_\theta) \partial_{\theta_j} \rho_\theta.$$

The authors discuss that, under suitable parametrizations of $\theta \mapsto \pi$, this gradient preconditioner is a Riemannian metric on Θ_{safe} and natural policy gradient flows based on $G_C(\theta_k)$ leave Θ_{safe} invariant. Further, $G_C(\theta_k)^+$ is equivalent to the Hessian of D_C :

$$H_C(\theta) = \mathbb{E}_{s \sim \rho_\theta} F(\theta) + \sum_i \beta_i \Phi''(b_i - C_i(\theta)) \nabla_\theta^2 C_i(\theta) \Big|_{\theta=\theta_k}.$$

where F is the fisher information of the policy. Unlike in TRPO, the divergence itself is not easy to estimate, however, the authors demonstrate that another divergence has the same Hessian, i.e. is equivalent up to second order in the policy parameters. It is derived using a ‘‘surrogate advantage trick’’ for C_i and results in the divergence

$$\bar{D}_{KL}(\pi, \pi_k) + \beta \bar{D}_\Phi(\pi, \pi_k) = \bar{D}_{KL}(\pi, \pi_k) + \beta \cdot [\phi(b_k - \mathbb{A}_c^{\pi_k}(\pi)) - \phi(b_k) - \phi'(b_k) \cdot \mathbb{A}_c^{\pi_k}(\pi)], \quad (23)$$

which is ultimately used as a drop-in replacement for the conventional divergence in TRPO.

A.3 Central Paths

In the small step size limit, the trajectories induced by trust region methods converge to the corresponding natural policy gradient (NPG) flow on the state-action polytope \mathcal{K} . The space of state-action occupancies $\rho \in \mathcal{K}$ forms not only a polytope, but a Hessian manifold Müller and Montúfar [2023]. C-TRPO induces such a gradient flow on the LP Equation 2 w.r.t the Hessian geometry induced by the convex function

$$\Phi(\rho) = \sum_{s,a} \rho(s, a) \log \pi(a|s) - \beta \log(\rho - \sum_{s,a} \rho(s, a) c(s, a)). \quad (24)$$

It is well known that Hessian gradient flows (ρ_t) of linear programs follow the central path, meaning that they are characterized as the optimizers of regularized linear programs with regularization strength t^{-1} . In policy space, we obtain for a single constraint

$$\pi_t = \arg \max \{ R(\pi) + t^{-1} D_\Phi(\pi, \pi_0) : \pi, C(\pi) \leq d \}. \quad (25)$$

Since Φ curves infinitely towards the boundary of the feasible set of LP Equation 2, solving the problem posed by C-TRPO corresponds to solving LP Equation 2 using an interior point / barrier method with barrier $D_\Phi(\cdot, \pi_0)$. For a more detailed discussion of Hessian geometries and natural policy gradients see Alvarez et al. [2004], Müller and Montúfar [2023], Müller and Cayci [2024].

B Proofs of Section 3

B.1 Exact Penalty Methods

We provide a general result for the exactness of the penalties considered in this work. For general discussions of exact penalty methods, we refer to standard textbooks in optimization Bertsekas [1997], Nocedal and Wright [1999]. Here, we consider a compact subset $X \subset \mathbb{R}^n$ with non-empty interior, differentiable functions $f, g \in C^1(X)$, and the constrained optimization problem

$$\max f(x) \quad \text{subject to } g(x) \leq b, \quad (26)$$

where we impose Slater’s condition $\{x \in X : g(x) < b\} \neq \emptyset$ to be non-empty and f to be concave and g to be convex. We denote the penalized functions by

$$P_\kappa(x) := f(x) - \kappa \max\{0, g(x) - b\}. \quad (27)$$

Recall the definition of the Lagrangian

$$\mathcal{L}(x, \lambda) = f(x) - \lambda(g(x) - b). \quad (28)$$

Theorem B.1 (Exactness for convex programs). *Assume that there exists a solution $x^* \in X$ of equation 26 and denote the corresponding dual variable by $\lambda^* \geq 0$. For $\kappa > \lambda^*$ we have*

$$\arg \max\{f(x) : x \in X, g(x) \leq b\} = \arg \max\{P_\kappa(x) : x \in X\}. \quad (29)$$

Proof. Consider an infeasible point $\bar{x} \in X$ of P_κ , meaning that $g(\bar{x}) > b$. Note by convexity x^* maximizes the Lagrangian $\mathcal{L}(\cdot, \lambda^*)$. Then

$$P_\kappa(\bar{x}) = f(\bar{x}) - \kappa(g(\bar{x}) - b) < f(\bar{x}) - \lambda^*(g(\bar{x}) - b) = \mathcal{L}(\bar{x}, \lambda^*) \leq \mathcal{L}(x^*, \lambda^*) = P_\kappa(x^*).$$

Hence, every maximizer of P_κ is feasible and thus a solution of the regularized problem, showing the inclusion \supseteq . As P_κ agrees with f for feasible points, we also obtain that every maximizer of f over the feasible set is a maximizer of P_κ . □

B.2 C3PO Exact Penalty

Proposition 3.2 (Positive Exactness). *For $0 \leq \mathbb{A}_c^{\pi_k}(\pi) < d - C(\pi_k)$, there exist w and δ_B for which the solution sets of problems 8, 9 and 10 agree.*

Proof. Let

$$P_{\text{Barrier}} := \{\pi : D_B(\pi, \pi_k) \leq \delta_B, \mathbb{A}_c^{\pi_k}(\pi) \geq 0\}$$

and

$$P_{\text{Lin}} := \{\pi : \mathbb{A}_c^{\pi_k}(\pi) - w \cdot b \leq 0, \mathbb{A}_c^{\pi_k}(\pi) \geq 0\}.$$

Note that

$$D_B(\pi, \pi_k) = \frac{b - \mathbb{A}_c^{\pi_k}(\pi)}{b} - \log\left(\frac{b - \mathbb{A}_c^{\pi_k}(\pi)}{b}\right) - 1 \quad (30)$$

is a strictly convex increasing function of $\mathbb{A}_c^{\pi_k}$ for $\mathbb{A}_c^{\pi_k} \geq 0$ (see Figure 2). This means that there exists a unique $\mathbb{A}_B > 0$ that solves

$$\frac{b - \mathbb{A}_B}{b} - \log\left(\frac{b - \mathbb{A}_B}{b}\right) - 1 = \delta_B \quad (31)$$

and for $\mathbb{A}_B \geq \mathbb{A}_c^{\pi_k}(\pi) > 0$ it holds that $\delta_B \geq D_B(\pi, \pi_k) > 0$. To solve for \mathbb{A}_B , we rewrite 31 as

$$\left(\frac{\mathbb{A}_B - b}{b}\right) \exp\left(\frac{\mathbb{A}_B - b}{b}\right) = -\exp(-\delta_B - 1). \quad (32)$$

and use the definition of Lambert's W-Function Corless et al. [1996] to invert the left hand side as follows

$$\frac{\mathbb{A}_B - b}{b} = W(-\exp(-\delta_B - 1)), \quad (33)$$

where W is the real part of the principle branch of the W-Function. Finally, rearranging yields

$$\mathbb{A}_B = b \cdot (W(-\exp(-\delta_B - 1)) + 1). \quad (34)$$

Note that $b > \mathbb{A}_B > 0$ must still hold. With this result,

$$P_{\text{Barrier}} = \{\pi : \mathbb{A}_c^{\pi_k}(\pi) - \mathbb{A}_B < 0, \mathbb{A}_c^{\pi_k}(\pi) \geq 0\}, \quad (35)$$

$$= \{\pi : \mathbb{A}_c^{\pi_k}(\pi) - b(W(-\exp(-\delta_B - 1)) + 1) < 0, \mathbb{A}_c^{\pi_k}(\pi) \geq 0\}, \quad (36)$$

$$= \{\pi : \mathbb{A}_c^{\pi_k}(\pi) - b w < 0, \mathbb{A}_c^{\pi_k}(\pi) \geq 0\}, \quad (37)$$

showing that $P_{\text{Barrier}} = P_{\text{Lin}}$ for a unique w .

Further, since $\min(b, w \cdot b) = w \cdot b$ for $b > 0$, the solution sets of

$$\max_{\pi \in \Pi} \mathbb{A}_r^{\pi_k}(\pi) \text{ s.t. } \mathbb{A}_c^{\pi_k}(\pi) - \mathbb{A}_B < 0 \quad \text{and } \bar{D}_{\text{KL}}(\pi, \pi_k) < \delta_{\text{KL}} \quad (38)$$

$$\max_{\pi \in \Pi} \mathbb{A}_r^{\pi_k}(\pi) \text{ s.t. } D_B(\pi, \pi_k) < \delta_B \quad \text{and } \bar{D}_{\text{KL}}(\pi, \pi_k) < \delta_{\text{KL}} \quad (39)$$

agree for $\mathbb{A}_c^{\pi_k}(\pi) \geq 0$ and $w = W(-\exp(-\delta_B - 1)) + 1$.

Finally, by theorem B.1, they must also agree with the solutions of

$$\max_{\pi \in \Pi} \mathbb{A}_r^{\pi_k}(\pi) - \kappa_k \max\{0, \mathbb{A}_c^{\pi_k}(\pi) - \min(b, w \cdot b)\} \quad \text{s.t. } \bar{D}_{\text{KL}}(\pi, \pi_k) < \delta_{\text{KL}}, \quad (40)$$

$$\max_{\pi \in \Pi} \mathbb{A}_r^{\pi_k}(\pi) - \kappa_k \max\{0, D_{\text{B}}(\pi, \pi_k) - \delta_B\} \quad \text{s.t. } \bar{D}_{\text{KL}}(\pi, \pi_k) < \delta_{\text{KL}}, \quad (41)$$

under the same conditions and for large enough κ . \square

Note that the cost budget $b = d - C(\pi_k)$ is multiplied with a fixed function of δ_{B} . Hence, we can use w as the hyper-parameter immediately instead of defining it through δ_{B} .




Figure 5: Aggregated performance using the inter quartile mean (IQM) across 8 tasks from Safety Gymnasium for all algorithms (except P2BPO) across 8 tasks. P2BPO has been excluded, since the final cost (right) was off the charts. This may be due to the missing penalty coefficient in the algorithm.

C Experiment Details

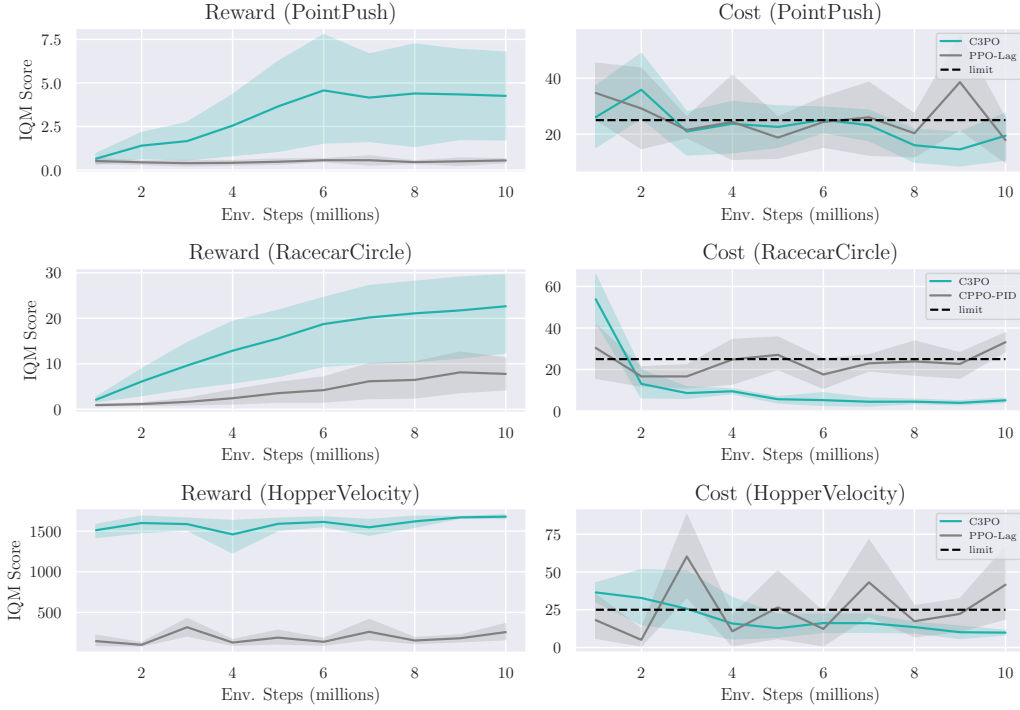


Figure 6: Hand-picked examples where central path approximation improves final reward performance.

Table 1: Performance of 8 representative safe policy optimization algorithms on 8 tasks from Safety Gymnasium for 10 million steps and a cost threshold of 25.0 aggregated over 5 seeds each. Bold marks the algorithm with the highest mean cumulative reward among the admissible ones. An algorithm is admissible, if its average cumulative cost achieved at the end of training is below the threshold.

		Ant	HalfCheetah	Humanoid	Hopper	CarButton	PointGoal	RacecarCircle	PointPush
C3PO	<i>R</i>	3043 ± 44	2458 ± 436	5389 ± 93	1674 ± 35	2.3 ± 0.7	23.8 ± 0.9	25.9 ± 5.1	4.5 ± 2.6
	<i>C</i>	15.0 ± 4.7	13.3 ± 6.4	1.2 ± 0.9	9.9 ± 1.7	53.4 ± 22.3	37.9 ± 1.7	5.0 ± 1.7	20.2 ± 10.0
C-TRPO	<i>R</i>	3019 ± 149	2841 ± 41	5746 ± 248	1621 ± 82	1.1 ± 0.2	19.3 ± 0.9	29.5 ± 3.1	1.0 ± 6.6
	<i>C</i>	13.2 ± 9.2	12.1 ± 7.6	12.2 ± 5.9	17.7 ± 8.0	34.0 ± 10.2	23.3 ± 3.6	20.2 ± 4.0	25.3 ± 7.0
CPO	<i>R</i>	3106 ± 21	2824 ± 104	5569 ± 349	1696 ± 19	1.1 ± 0.2	20.4 ± 2.0	29.8 ± 1.9	0.7 ± 2.9
	<i>C</i>	25.1 ± 11.3	23.1 ± 8.0	16.2 ± 8.6	25.7 ± 4.4	33.5 ± 8.7	28.2 ± 4.1	23.1 ± 4.5	28.9 ± 20.0
PPO-LAG	<i>R</i>	3210 ± 85	3033 ± 1	5814 ± 122	240 ± 159	0.3 ± 0.8	9.4 ± 1.8	30.9 ± 1.8	0.6 ± 0.0
	<i>C</i>	28.9 ± 8.7	23.2 ± 1.9	12.7 ± 31.0	38.8 ± 36.4	39.2 ± 41.1	22.5 ± 10.1	31.7 ± 2.7	18.2 ± 9.5
CPPO-PID	<i>R</i>	3205 ± 76	3036 ± 10	5877 ± 84	1657 ± 61	-1.2 ± 0.6	6.1 ± 4.8	8.1 ± 4.3	1.0 ± 1.1
	<i>C</i>	26.2 ± 4.4	26.5 ± 7.2	20.3 ± 6.0	18.6 ± 8.1	23.8 ± 6.0	21.8 ± 6.8	33.3 ± 5.9	22.8 ± 9.9
P2BPO	<i>R</i>	3269 ± 18	2928 ± 46	5293 ± 171	1573 ± 85	6.1 ± 0.9	25.9 ± 0.2	15.7 ± 7.5	1.1 ± 0.5
	<i>C</i>	32.3 ± 8.9	26.0 ± 19.7	1.5 ± 1.1	13.2 ± 11.7	125 ± 14	39.6 ± 5.7	5.5 ± 8.0	43.8 ± 28.9
P3O	<i>R</i>	3122 ± 24	3020 ± 12	5492 ± 118	1633 ± 49	0.2 ± 0.3	5.7 ± 0.3	0.9 ± 0.1	0.7 ± 0.6
	<i>C</i>	21.2 ± 2.5	27.0 ± 1.1	4.2 ± 2.2	14.6 ± 1.6	40.9 ± 18.2	17.1 ± 6.2	13.1 ± 4.6	14.1 ± 9.4
PPO	<i>R</i>	5402 ± 274	6583 ± 954	6138 ± 699	1810 ± 390	18.2 ± 1.2	26.6 ± 0.2	40.8 ± 0.5	0.9 ± 0.7
	<i>C</i>	887 ± 27	976 ± 1	783 ± 60	435 ± 85	378 ± 18	50.7 ± 3.3	200 ± 4	42.9 ± 24.0