

ABSOLUTE POLICY OPTIMIZATION

Anonymous authors

Paper under double-blind review

ABSTRACT

In recent years, trust region on-policy reinforcement learning has achieved impressive results in addressing complex control tasks and gaming scenarios. However, contemporary state-of-the-art algorithms within this category primarily emphasize improvement in expected performance, lacking the ability to control over the worst-case performance outcomes. To address this limitation, we introduce a novel objective function; optimizing which leads to guaranteed monotonic improvement in the lower bound of near-total performance samples. We call it improvement of absolute performance. Building upon this groundbreaking theoretical advancement, we further introduce a practical solution called Absolute Policy Optimization (APO). Our experiments demonstrate the effectiveness of our approach across challenging continuous control benchmark tasks and extend its applicability to mastering Atari games. Our findings reveal that APO significantly outperforms state-of-the-art policy gradient algorithms, resulting in substantial improvements in worst-case performance, as well as expected performance.

1 INTRODUCTION

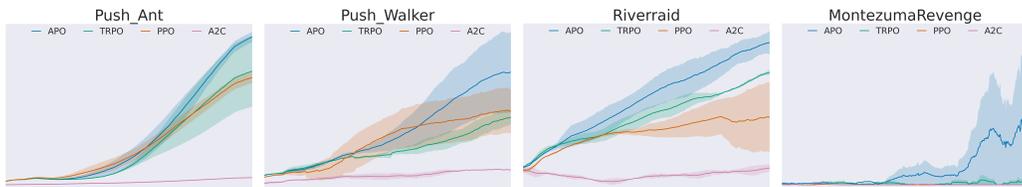


Figure 1: Illustration of performance improvement.

Existing reinforcement learning algorithms have taken the expectation of improving cumulative rewards (referred to as **performance**) as their core optimization objective. Within this framework, trust region-based on-policy reinforcement learning algorithms have achieved the most outstanding results. However, the representative trust-region policy optimization (TRPO) Schulman et al. (2015) only ensures the monotonic improvement of the expectation of this performance distribution, it fails to exert control over the worst-case performance sample originating from the same distribution. In this paper, we introduce a novel theoretical breakthrough that ensures the monotonic improvement of the lower bound of near-total performance samples (**absolute performance**) from the distribution. Subsequently, we implement a series of approximations to transform this theoretically-grounded algorithm into a practical solution, which we refer to as **Absolute Policy Optimization (APO)**. Remarkably, APO exhibits scalability and can efficiently optimize nonlinear policies characterized by tens of thousands of parameters. Our experimental results underscore the effectiveness of APO, demonstrating substantial performance improvements in terms of both absolute performance and expected performance compared to state-of-the-art policy gradient algorithms. These improvements are evident across challenging continuous control benchmark tasks and extend to the realm of playing Atari games. Figure 1 provides a visual representation of our approach’s superiority, where APO can effectively handle tasks that other algorithms inherently struggle to optimize, spanning both continuous and discrete control domains. This work serves as the next generation base reinforcement learning algorithm, and represents a significant step towards developing practical RL algorithms that can be robustly applied to many real-world problems.

2 RELATED WORKS

Model-Free Deep Reinforcement Learning Model-free deep reinforcement learning (RL) algorithms have found applications from the realm of games (Mnih et al., 2013b; Silver et al., 2016) to the intricate domain of robotic control (Schulman et al., 2015). The leading contenders of the model free reinforcement learning algorithms include (i) deep Q-learning (Mnih et al., 2013a; Hausknecht & Stone, 2015; van Hasselt et al., 2015; Hessel et al., 2018), (ii) off-policy policy gradient methods (Silver et al., 2014; Lillicrap et al., 2015; Gu et al., 2016; Fujimoto et al., 2018; Haarnoja et al., 2018), and (iii) trust region on-policy policy gradient methods (Schulman et al., 2015; 2017).

Among those categories, Q-learning-based techniques, augmented with function approximation, have exhibited remarkable prowess over tasks with discrete action spaces, e.g. Atari game playing (Bellemare et al., 2013). However, these methods performs poorly in the realm of continuous control benchmarks, notably exemplified in OpenAI Gym (Brockman et al., 2016a; Duan et al., 2016).

In contrast, off-policy policy gradient methods extend Q-learning-based strategies via introducing an independent actor network to handle continuous control tasks, as exemplified by the Deep Deterministic Policy Gradient (DDPG)(Lillicrap et al., 2015). However, off-policy methods suffer from stability issues and susceptibility to hyperparameter tuning nuances(Haarnoja et al., 2018). Recently, enhancements are made to incorporate entropy to foster exploration (Haarnoja et al., 2018) and mitigate the overestimation bias through target networks (Fujimoto et al., 2018). Despite these advancements, the convergence characteristics of off-policy policy gradient methods remain incompletely understood, primarily explored under stringent assumptions such as infinite sampling and Q-updates (Fujimoto et al., 2018). Moreover, off-policy policy gradient methods are primarily tailored for continuous action spaces.

Conversely, trust region on-policy policy gradient methods harmoniously accommodate both continuous and discrete action spaces while showcasing superior stability and dependable convergence properties. Notably, the representative Trust Region Policy Optimization (TRPO)(Schulman et al. (2015), complemented by its pragmatic counterpart, Proximal Policy Optimization (PPO)(Schulman et al. (2017), have consistently delivered impressive performance across an array of demanding benchmark tasks. Furthermore, those methods have largely helped training of groundbreaking artificial intelligence applications, including ChatGPT (Schulman et al., 2022), the automated Rubik’s Cube solver with a robotic hand (Akkaya et al., 2019), and the championship-level drone racing (Kaufmann et al., 2023), thereby reaffirming their profound impact on advancing the frontiers of AI technology.

Attempts to Improve Trust Region Methods Recently, many efforts are made to improve trust region on-policy methods, including (i) *improve computation efficiency*. TREFree (Sun et al., 2023) introduced a novel surrogate objective that eliminates trust region constraints. (ii) *encourage exploration*. COPOS (Pajarinen et al., 2019) applied compatible value function approximation to effectively control entropy during policy updates. (iii) *improve training stability and data-efficiency*. Truly PPO (TR-PPO) (Wang et al., 2020) introduced a new clipping function and trust region-based triggering condition. Generalized PPO (GePPO) (Queeney et al., 2021) extended PPO to an off-policy variant, thereby enhancing sampling efficiency through data reuse. AlphaPPO (Xu et al., 2023) introduced alpha divergence, a parametric metric that offers a more effective description of policy differences, resulting more stable training performance.

There are also improvements considering variance control, including (i) *variance reduction of policy gradient*. Xu et al. and Papini et al. applied the stochastic variance reduced gradient descent (SVRG) technique for getting stochastic variance-reduced version of policy gradient (SVRPO) to improve the sample efficiency. (Yuan et al.) incorporates the StochAstic Recursive grAdient algoritHm (SARAH) into the TRPO framework to get more stable variance. (ii) *variance reduction of performance update*. (Tomczak et al., 2019) introduced a surrogate objective with approximate importance sampling to strike a balance between performance update bias and variance. (iii) *variance reduction of importance sampling*. (Lin et al., 2023) introduced sample dropout to bound the variance of importance sampling estimate by dropping out samples when their ratio deviation is too high.

Although trust region-based methods have achieved notable success, there remains substantial potential for improvement. A critical gap in existing approaches lies in their inability to exert control over the worst-case individual performance samples stemming from the policy. Unforeseen instances of poor performance can result in training instability, thereby jeopardizing the reliability of solutions

in real-world applications. In our research, we bridge this gap by introducing novel theoretical results that ensure a monotonic improvement of the lower bound of near-total performance samples.

3 PROBLEM FORMULATION

3.1 NOTATIONS

Consider an infinite-horizon discounted Markov decision process (MDP) defined by the tuple $(\mathcal{S}, \mathcal{A}, \gamma, \mathcal{R}, P, \mu)$, where \mathcal{S} is the state space, and \mathcal{A} is the control space, $R : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$ is a bounded reward function, $0 \leq \gamma < 1$ is the discount factor, $\mu : \mathcal{S} \mapsto \mathbb{R}$ is the bounded initial state distribution, and $P : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \mapsto \mathbb{R}$ is the transition probability. $P(s'|s, a)$ is the probability of transitioning to state s' when the agent takes action a at state s . A stationary policy $\pi : \mathcal{S} \mapsto \mathcal{P}(\mathcal{A})$ is a mapping from states to a probability distribution over actions, with $\pi(a|s)$ denoting the probability of selecting action a in state s . We denote the set of all stationary policies by Π . Subsequently, we denote π_θ as the policy that is parameterized by the parameter θ .

The standard goal for MDP is to learn a policy π that maximizes a performance measure $\mathcal{J}(\pi)$ which is computed via the discounted sum of reward:

$$\mathcal{J}(\pi) = \mathbb{E}_{\tau \sim \pi} [\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t, s_{t+1})], \quad (1)$$

where $\tau = [s_0, a_0, s_1, \dots]$, and $\tau \sim \pi$ is shorthand for that the distribution over trajectories depends on $\pi : s_0 \sim \mu, a_t \sim \pi(\cdot|s_t), s_{t+1} \sim P(\cdot|s_t, a_t)$. And the objective is to select a policy π that maximizes the performance measure: $\max_{\pi \in \Pi} \mathcal{J}(\pi)$.

3.2 ABSOLUTE PERFORMANCE BOUND

Notice that the above considers maximizing the expected reward performance, which, unfortunately, does not provide control over each individual performance sample derived from the policy π . To clarify, a performance sample is defined here as $R_\pi(s_0) \doteq \sum_{t=0}^{\infty} \gamma^t R(s_t, a_t, s_{t+1})$, where the state action sequence $\hat{\tau} = [a_0, s_1, \dots] \sim \pi$ starts with an initial state s_0 , which follows initial state distribution μ . In practical reinforcement learning setting, unexpected poor performance samples can lead to training instability, compromising the reliability of solutions in real-world applications. To tackle this issue, our fundamental insight is that policy optimization should not be solely fixated on enhancing expected performance, but also on improving the worst-case performance samples originating from the distribution of the variable $R_\pi(s_0)$.

However, it's important to acknowledge that within the Markov Decision Process framework, any possible state visitation is inherently assigned a non-zero probability. In essence, any policy performance distribution inherently accommodates the statistical possibility of all conceivable $R_\pi(s_0)$. Therefore, our ultimate goal is to improve the lower bound of near-total performance samples derived from the policy. We denote this lower bound as the absolute performance bound for the policy:

Definition 1 (Absolute Performance Bound). $\mathcal{B}_k(\pi)$ is called the absolute performance bound with p_k confidence if it satisfies the following condition:

$$Pr(R_\pi(s_0) \geq \mathcal{B}_k(\pi)) \geq p_k, \quad (2)$$

where $p_k \doteq 1 - \frac{1}{k^2} \in (0, 1)$ and k is the probability factor ($k > 1$ and $k \in \mathbb{R}$), which can be set according to the demand for probability magnitude.

Remark 1. Definition 1 shows that more than p_k of the samples from the distribution of $R_\pi(s_0)$ will be larger than the bound $\mathcal{B}_k(\pi)$. By setting $p_k \rightarrow 1$, $\mathcal{B}_k(\pi)$ represents the lower bound of near-total performance samples of policy π .

3.3 ABSOLUTE MARKOV DECISION PROCESS

In this paper, our focus is on a special class of Markov Decision Processes (MDP) characterized by improvement of the absolute absolute performance. This unique class is termed **Absolute Markov Decision Process (AMDP)**. Much like a standard MDP, an AMDP is defined by a tuple $(\mathcal{S}, \mathcal{A}, \gamma, \mathcal{R}, P, \mu, k)$, with the inclusion of an extra probabilistic factor denoted as k , which

serves to modulate the degree of conservatism in the absolute performance bound. In accordance with Definition 1, the overarching objective within the AMDP framework is to identify a policy π that effectively improves \mathcal{B}_k . For an unknown performance distribution, we first define

$$\mathcal{V}(\pi) \doteq \mathbb{E}_{\tau \sim \pi} \left[\left(\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t, s_{t+1}) - \mathcal{J}(\pi) \right)^2 \right]$$
 as the variance of the performance distribution.

Then, we can leverage the Chebyshev's inequality theory Saw et al. (1984) to obtain an absolute performance bound as $\mathcal{B}_k(\pi) \doteq \mathcal{J}(\pi) - k\mathcal{V}(\pi)$, which is guaranteed to satisfy Definition 1 (proved in Proposition 1). Thus, AMDP addresses the following optimization

$$\max_{\pi \in \Pi} \mathcal{J}(\pi) - k\mathcal{V}(\pi). \quad (3)$$

Here we define the on-policy value function as $V_\pi(s) \doteq \mathbb{E}_{\tau \sim \pi} [R_\pi(s) | s_0 = s]$, the on-policy action-value function as $Q_\pi(s, a) = Q_\pi(s, a, s') \doteq \mathbb{E}_{\tau \sim \pi} [R_\pi(s) | s_0 = s, a_0 = a]$, and the advantage function as $A_\pi(s, a) = A_\pi(s, a, s') \doteq Q_\pi(s, a) - V_\pi(s)$. We also define $\bar{A}_\pi(s)$ as the expected advantage of π' over π at state s : $\bar{A}_\pi(s) \doteq \mathbb{E}_{a \sim \pi'} [A_\pi(s, a)]$.

4 ABSOLUTE POLICY OPTIMIZATION

To optimize equation 3, we need to evaluate the objective with respect to an unknown π . Our main intuition is to find a surrogate function for the objective, such that (i) it represents a tight lower bound of the objective; and (ii) it can be easily estimated from the samples on the most recent policy. To solve large and continuous AMDPs, policy search algorithms look for the optimal policy within a set $\Pi_\theta \subset \Pi$ of parametrized policies. Mathematically, APO updates solve the following optimization:

$$\pi_{j+1} = \mathbf{argmax}_{\pi \in \Pi_\theta} \mathcal{J}_{\pi, \pi_j}^l - k(MV_{\pi, \pi_j} + VM_{\pi, \pi_j}) \quad , \quad (4)$$

where $\mathcal{J}_{\pi, \pi_j}^l$ represents the lower bound surrogate function for $\mathcal{J}(\pi)$ and $(MV_{\pi, \pi_j} + VM_{\pi, \pi_j})$ represents the upper bound surrogate function for $\mathcal{V}(\pi)$.

Remark 2. *Since the performance samples from the same start state belong to a one-dimensional distribution, performance samples from different start states belong to a mixture of one-dimensional distributions. MV_{π, π_j} reflects the upper bound of expected variance of the return over different start states. VM_{π, π_j} reflects the upper bound of variance of the expected return of different start states. The detailed interpretations are discussed in Equation (18), Lemma 1, and Lemma 2.*

Here $\mathcal{J}_{\pi, \pi_j}^l, MV_{\pi, \pi_j}, VM_{\pi, \pi_j}$ are defined as:

$$\mathcal{J}_{\pi, \pi_j}^l \doteq \mathcal{J}(\pi_j) + \frac{1}{1-\gamma} \mathbb{E}_{s \sim d^{\pi_j}, a \sim \pi} \left[A_{\pi_j}(s, a) - \frac{2\gamma\epsilon^\pi}{1-\gamma} \sqrt{\frac{1}{2} \mathcal{D}_{KL}(\pi || \pi_j)[s]} \right] \quad (5)$$

$$MV_{\pi, \pi_j} \doteq \frac{\|\mu^T\|_\infty}{1-\gamma^2} \mathbf{max}_s \left| \mathbb{E}_{\substack{a \sim \pi \\ s' \sim P}} [A_{\pi_j}(s, a, s')^2] - \mathbb{E}_{\substack{a \sim \pi_j \\ s' \sim P}} [A_{\pi_j}(s, a, s')^2] + |H(s, a, s')|_{max}^2 \right. \quad (6)$$

$$\left. + 2 \mathbb{E}_{\substack{a \sim \pi \\ s' \sim P}} [A_{\pi_j}(s, a, s')] \cdot |H(s, a, s')|_{max} \right| + MV_{\pi_j} + \frac{2\gamma^2 \|\mu^T\|_\infty}{1-\gamma^2} \sqrt{\frac{1}{2} \mathcal{D}_{KL}(\pi || \pi_j)[s]} \cdot \|\Omega_{\pi_j}\|_\infty$$

$$VM_{\pi, \pi_j} \doteq \|\mu^T\|_\infty \mathbf{max}_s \left| |\eta(s)|_{max}^2 + 2|V_{\pi_j}(s)| \cdot |\eta(s)|_{max} \right| - \mathbf{min}_{s_0 \sim \mu} (\mathcal{J}(\pi))^2 + \mathbb{E}_{s_0 \sim \mu} [V_{\pi_j}^2(s_0)] \quad (7)$$

where $\mathcal{D}_{KL}(\pi || \pi_j)[s]$ is KL divergence between two policy (π, π_j) at state s , $\epsilon^\pi \doteq$

$$\mathbf{max}_s |\mathbb{E}_{a \sim \pi} [A_{\pi_j}(s, a)]|, d^{\pi_j} \doteq (1-\gamma) \sum_{t=0}^H \gamma^t P(s_t = s | \pi_j), \Omega_{\pi_j} \doteq \begin{bmatrix} \omega_{\pi_j}(s^1) \\ \omega_{\pi_j}(s^2) \\ \vdots \end{bmatrix}, \omega_{\pi_j}(s) \doteq$$

$$\mathbb{E}_{a \sim \pi_j, s' \sim P} [Q_{\pi_j}(s, a, s')^2] - V_{\pi_j}(s)^2, MV_{\pi_j} \doteq \mathbb{E}_{s_0 \sim \mu} [\text{Var}[R_{\pi_j}(s_0)]] \text{ and } \mathbf{min} (\mathcal{J}(\pi))^2 \doteq$$

$$\mathbf{min}_{\mathcal{J}(\pi) \in [\mathcal{J}_{\pi, \pi_j}^l, \mathcal{J}_{\pi, \pi_j}^u]} (\mathcal{J}(\pi))^2 \text{ with } \mathcal{J}_{\pi, \pi_j}^u \doteq \mathcal{J}(\pi_j) + \frac{1}{1-\gamma} \mathbb{E}_{\substack{s \sim d^{\pi_j} \\ a \sim \pi}} \left[A_{\pi_j}(s, a) + \frac{2\gamma\epsilon^\pi}{1-\gamma} \sqrt{\frac{1}{2} \mathcal{D}_{KL}(\pi || \pi_j)[s]} \right].$$

Additionally,

$$|H(s, a, s')|_{max} \doteq \left| \gamma \mathbb{E}_{\substack{s_0=s', \\ \hat{\tau} \sim \pi_j}} \left[\sum_{t=0}^{\infty} \gamma^t \bar{A}_{\pi_j}(s_t) \right] - \mathbb{E}_{\substack{s_0=s, \\ \hat{\tau} \sim \pi_j}} \left[\sum_{t=0}^{\infty} \gamma^t \bar{A}_{\pi_j}(s_t) \right] \right| + \frac{2\gamma(1+\gamma)\epsilon}{(1-\gamma)^2} \mathcal{D}_{KL}(\pi \parallel \pi_j)[s] \quad (8)$$

$$|\eta(s)|_{max} \doteq \left| \mathbb{E}_{\substack{s_0=s, \\ \hat{\tau} \sim \pi_j}} \left[\sum_{t=0}^{\infty} \gamma^t \bar{A}_{\pi_j}(s_t) \right] \right| + \frac{2\gamma\epsilon}{(1-\gamma)^2} \mathcal{D}_{KL}(\pi \parallel \pi_j)[s] \quad (9)$$

$$\epsilon \doteq \max_{s,a} |A_{\pi_j}(s, a)|. \quad (10)$$

Theoretical Guarantees for APO

Theorem 1 (Monotonic Improvement of Absolute Performance). *Suppose π, π' are related by equation 4, then absolute performance bound $\mathcal{B}_k(\pi) = \mathcal{J}(\pi) - k\mathcal{V}(\pi)$ satisfies $\mathcal{B}_k(\pi') \geq \mathcal{B}_k(\pi)$.*

The proof for Theorem 1 is summarized in Appendix A.

5 OPTIMIZED APO

In addition to the guaranteed enhancement of the absolute performance bound, it is also highly desirable for policy optimization to improve expected performance. By leveraging the concept of policy improvement bounds as introduced in trust region methods (Schulman et al., 2015; Achiam et al., 2017), we can enhance APO by imposing additional constraint on the feasible solutions:

$$\pi_{j+1} = \underset{\pi \in \Pi_{\theta}}{\operatorname{argmax}} \mathcal{J}_{\pi, \pi_j}^l - k(MV_{\pi, \pi_j} + VM_{\pi, \pi_j}) \quad \text{s.t.} \quad \mathcal{J}_{\pi, \pi_j}^l \geq \mathcal{J}(\pi_j) \quad (11)$$

Remark 3. *Any feasible solution to equation 11 ensures (i) monotonic improvement of absolute performance [Theorem 1], and (ii) monotonic improvement of expected performance [Corollary 1, Corollary 3, (Achiam et al., 2017)]. Additionally, π_j is always a feasible solution towards equation 11.*

In local policy search (Peters & Schaal, 2008), the policy is iteratively updated by maximizing objective within a local neighborhood of the most recent policy π_j . Inspired by trust region methods, we approximate equation 11 with a trust region constraint instead of penalties on policy divergence:

$$\begin{aligned} \pi_{j+1} = \underset{\pi \in \Pi_{\theta}}{\operatorname{argmax}} \quad & \frac{1}{1 - \gamma} \mathbb{E}_{\substack{s \sim d^{\pi_j}, \\ a \sim \pi}} [A_{\pi_j}(s, a)] - k(\overline{MV}_{\pi, \pi_j} + \overline{VM}_{\pi, \pi_j}) \quad (12) \\ \text{s.t.} \quad & \frac{1}{1 - \gamma} \mathbb{E}_{\substack{s \sim d^{\pi_j}, \\ a \sim \pi}} [A_{\pi_j}(s, a)] \geq 0 \quad \text{and} \quad \bar{\mathcal{D}}_{KL}(\pi \parallel \pi_j) \leq \delta \end{aligned}$$

where δ is the step size, $\overline{MV}_{\pi, \pi_j} \doteq MV_{\pi, \pi_j} - MV_{\pi_j} - \frac{2\gamma^2 \|\mu^T\|_{\infty}}{1-\gamma^2} \sqrt{\frac{1}{2} \mathcal{D}_{KL}(\pi \parallel \pi_j)[s] \cdot \|\Omega_{\pi_j}\|_{\infty}}$ and $\overline{VM}_{\pi, \pi_j} = VM_{\pi, \pi_j} - \mathbb{E}_{s_0 \sim \mu} [V_{\pi_j}^2(s_0)]$. The set $\{\pi \in \Pi_{\theta} : \bar{\mathcal{D}}_{KL}(\pi \parallel \pi_j) = \mathbb{E}_{s \sim \pi_j} [\mathcal{D}_{KL}(\pi \parallel \pi_j)[s]] \leq \delta\}$ is called *trust region*. Notice that MV_{π_j} and $\mathbb{E}_{s_0 \sim \mu} [V_{\pi_j}^2(s_0)]$ are computable constant.

Equation 12 represents a constrained policy optimization problem, the solution to which has been discussed in prior work (Achiam et al., 2017). In practical implementations, due to various approximation errors, equation 12 may often lead to infeasible solutions. Hence, the line search trick (Achiam et al., 2017) results in very small step sizes and makes it hard to update the policy effectively. To overcome that challenge, we optimize the problem in the form of weighted sum:

$$\begin{aligned} \pi_{j+1} = \underset{\pi \in \Pi_{\theta}}{\operatorname{argmax}} \quad & \left(w_1 \left(\frac{1}{1 - \gamma} \mathbb{E}_{\substack{s \sim d^{\pi_j}, \\ a \sim \pi}} [A_{\pi_j}(s, a)] \right) + \right. \quad (13) \\ & \left. w_2 \left(\frac{1}{1 - \gamma} \mathbb{E}_{\substack{s \sim d^{\pi_j}, \\ a \sim \pi}} [A_{\pi_j}(s, a)] - k(\overline{MV}_{\pi, \pi_j} + \overline{VM}_{\pi, \pi_j}) \right) \right) \\ \text{s.t.} \quad & \bar{\mathcal{D}}_{KL}(\pi \parallel \pi_j) \leq \delta \end{aligned}$$

where w_1, w_2 are weights to control the importance of optimization objectives.

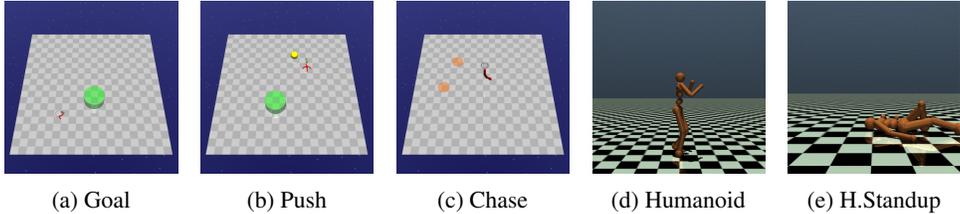


Figure 2: Tasks of continuous experiments

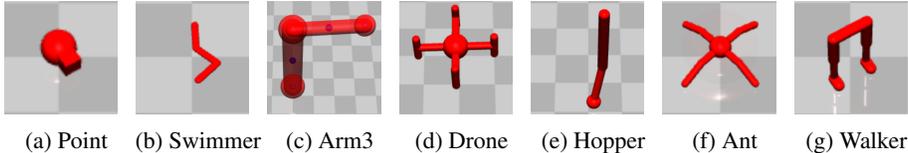


Figure 3: Robots of continuous tasks benchmark GUARD.

6 PRACTICAL IMPLEMENTATION

In this section, we show how to (i) simplify complex computations (ii) implement an efficient approximation to Equation (13) and (iii) encourage learning even when Equation (13) becomes infeasible. The full APO pseudocode is provided as Algorithm 1 in Appendix B.

Special Hyperparameters for Practical Implementation For practical implementation, we treat two items when implementing Equation (13) as hyperparameters. (i) $\|\mu^T\|_\infty$: we treat the infinity norm of μ (initial distribution of the system) as a constant parameter due to its inaccessibility during implementation. (ii) $|H(s, a, s')|_{max}$: We can either compute $|H(s, a, s')|_{max}$ from the most recent policy with equation 8 or treat it as a hyperparameter since $|H(s, a, s')|_{max}$ is bounded for any system with a bounded reward function. In practice, we found that the hyperparameter option helps increasing the performance. This setting will be discussed more detailedly in Section 7.5.

Efficient Approximation of APO We first estimate the expected reward advantage in Equation (13) via importance sampling with a sampling distribution π_k (Schulman et al., 2015) as

$$\mathbb{E}_{a \sim \pi, s' \sim P} [A_{\pi_j}(s, a, s')^2] = \mathbb{E}_{a \sim \pi_j, s' \sim P} [(\pi(a|s)/\pi_j(a|s)) A_{\pi_j}(s, a, s')^2] \quad (14)$$

Equation (14) allows us to replace $\mathbb{E}_{a \sim \pi, s' \sim P} [A_{\pi_j}(s, a, s')^2]$ with empirical estimates at each state-action pair (s, a) from rollouts by the previous policy π_j . The empirical estimate of reward advantage is given by $R(s, a, s') + \gamma V_{\pi_j}(s') - V_{\pi_j}(s)$. $V_{\pi_j}(s)$ can be computed at each state by taking the discounted future return. To proceed, we convexify Equation (13) by approximating the objective via first-order expansions, and the trust region constraint via second-order expansions. Then, Equation (13) can be efficiently solved using duality (Schulman et al., 2015).

Infeasible Solution An update to θ is computed according to the techniques described in Schulman et al. (2015) every time Equation (13) is solved. However, due to approximation errors, sometimes Equation (13) can become infeasible. In that case, we use a line search to ensure improvement of the surrogate objective and satisfaction of the KL divergence constraint. Starting with the maximal value of the step length, we shrink it exponentially until the objective improves. Without this line search, the algorithm occasionally computes large steps that cause a catastrophic degradation of performance.

7 EXPERIMENT

In our experiments, we want to answer the following questions:

- Q1:** How does APO compare with state-of-the-art on-policy RL algorithms?
- Q2:** What benefits are demonstrated by directly optimizing the absolute performance?
- Q3:** Is treating H_{max} as a hyperparameter necessary?
- Q4:** What are the impacts of different probability factor k choices?

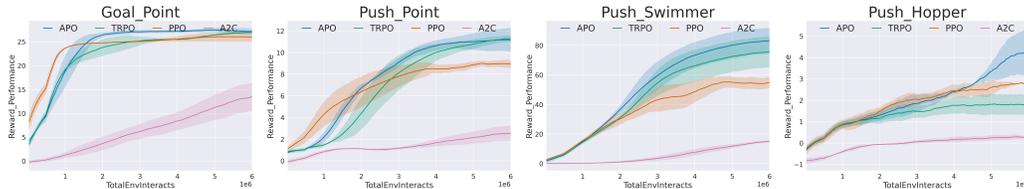


Figure 4: Comparison of results from four representative test suites in low dimensional continuous systems

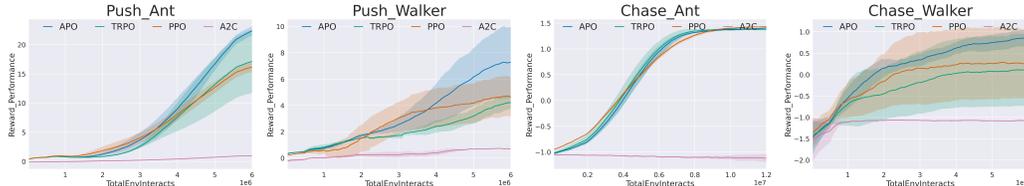


Figure 5: Comparison of results from four representative test suites in high dimensional continuous systems

7.1 EXPERIMENT SETUP

To answer the above, we run experiments on both continuous domain and the discrete domain.

Continuous Tasks Our continuous experiments are conducted on GUARD (Zhao et al., 2023), a challenging robot locomotion benchmark build upon Mujoco (Todorov et al., 2012) and Gym. Seven different robots are included: (i) **Point**: (Figure 3a) A point-mass robot ($\mathcal{A} \subseteq \mathbb{R}^2$) that can move on the ground. (ii) **Swimmer**: (Figure 3b) A three-link robot ($\mathcal{A} \subseteq \mathbb{R}^2$) that can move on the ground. (iii) **Arm3**: (Figure 3c) A fixed three-joint robot arm ($\mathcal{A} \subseteq \mathbb{R}^3$) that can move its end effector around with high flexibility. (iv) **Drone**: (Figure 3d) A quadrotor robot ($\mathcal{A} \subseteq \mathbb{R}^4$) that can move in the air. (v) **Hopper**: (Figure 3e) A one-legged robot ($\mathcal{A} \subseteq \mathbb{R}^5$) that can move on the ground. (vi) **Ant**: (Figure 3f) A quadrupedal robot ($\mathcal{A} \subseteq \mathbb{R}^8$) that can move on the ground. (vii) **Walker**: (Figure 3g) A bipedal robot ($\mathcal{A} \subseteq \mathbb{R}^{10}$) that can move on the ground. Furthermore, three different types of tasks are considered, including (i) **Goal**: (Figure 2a) robot navigates towards a series of 2D or 3D goal positions. (ii) **Push**: (Figure 2b) robot pushes a ball toward different goal positions. (iii) **Chase**: (Figure 2c) robot tracks multiple dynamic targets. Considering these different robots and tasks, we design 8 low-dim test suites and 4 high-dim test suits with 7 types of robots and 3 types of tasks, which are summarized in Table 3 in Appendix. We name these test suites as $\{\text{Task Type}\}_{\text{Robot}}$. Further details are listed in Appendix C.1.

Additionally, we conduct continuous control experiments on Mujoco Openai Gym (Brockman et al., 2016b). Two tasks are considered: (i) **Humanoid**: (Figure 2d) The 3D bipedal robot ($\mathcal{A} \subseteq \mathbb{R}^{17}$) is designed to simulate a human. And the goal of the environment is to walk forward as fast as possible without falling over. (ii) **Humanoid Standup**: (Figure 2e) The robot ($\mathcal{A} \subseteq \mathbb{R}^{17}$) is same with task **Humanoid**, but the goal is to make the humanoid standup and then keep it standing by applying torques on the various hinges. These two tasks are also summarized in Table 3.

Discrete Tasks We also test APO in all 62 Atari environments of (Brockman et al., 2016b) which are simulated on the Arcade Learning Environment benchmark (Bellemare et al., 2018). All experiments are based on ‘v5’ environments and ‘ram’ observation space.

Comparison Group We compare APO to the state-of-the-art base on-policy RL algorithms: (i) TRPO (Schulman et al., 2015) (ii) Advantage Actor Critic (A2C) (Mnih et al., 2016) (iii) PPO (Schulman et al., 2017). For all experiments, we take the best specific parameters mentioned in the original paper and keep the common parameters as the same. The policy π , the value V^π are all encoded in feedforward neural networks using two hidden layers of size (64,64) with tanh activations. The full list of parameters of all methods and tasks compared can be found in Appendix C.2.

7.2 COMPARISON TO OTHER ALGORITHMS IN THE CONTINUOUS DOMAIN

7.2.1 GUARD BENCHMARK

Low dimension Figure 4 shows representative comparison results on low dimensional system (See Appendix D for all results). APO is successful at getting more steady and higher final reward. We notice that PPO only gains faster convergence in part of the simplest task owing to its exploration abilities, the advantage decreases rapidly with more complex tasks such as PUSH. In difficult tasks, APO can perform best at the combined level of convergence speed and final performance.

High dimension Figure 5 reports the comparison results on challenging high-dimensional Ant_{PUSH, CHASE} and Walker_{PUSH, CHASE} tasks, where APO outperforms other baselines in getting higher reward and convergence speed.

7.2.2 MUJOCO BENCHMARK

To showcase the performance of APO on other high-dimensional continuous benchmark, we conduct additional experiments involving 3D humanoid robot (details have been introduced in Section 7.1) and compare the results with baseline method TRPO. See Table 5 for detailed hyperparameters and Figure 6 for learning curves.

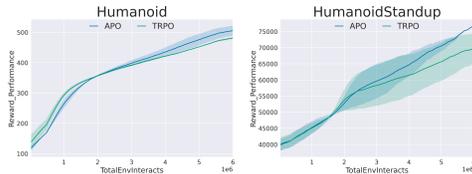


Figure 6: Additional experiments on MuJoCo

7.3 COMPARISON TO OTHER ALGORITHMS IN THE ATARI DOMAIN

The hyperparameters for Atari Domain are also provided in Appendix C.2. For the other three algorithms, we used hyperparameters that were tuned to maximize performance on this benchmark. Then we follow the metrics of (Schulman et al., 2017) to quantitatively evaluate the strengths of APO: (i) average expected reward per episode over **all epochs of training** (which favors fast learning), and (ii) average expected reward per episode over **last 100 epochs of training** (which favors final performance). Table 1 records the number of highest evaluation scores obtained by each algorithm across all games. The learning curves for all testing atari games are provided in Appendix D. To compare the performance of all testing algorithm to the TRPO baseline across games, we slightly change the normalization algorithm proposed by (van Hasselt et al., 2015) to obtain more reasonable score (See Appendix C.3 for further explanation of rationality) in percent. The score we used is average reward per episode over last 100 epochs of training:

$$\Delta_1 \doteq score_{agent} - score_{random}, \Delta_2 \doteq score_{TRPO} - score_{random} \quad (15)$$

$$score_{normalized} = \frac{\Delta_2}{\Delta_1} \text{ if } \Delta_1 < 0 \text{ and } \Delta_2 < 0 \text{ else } \frac{\Delta_1}{\Delta_2}$$

Then we use stacked bar chart in Figure 7 to visualize APO’s capabilities. Figure 7 show that APO has a superior combination of capabilities compared to other algorithms. So far the above experimental comparison answers **Q1**.

	APO	PPO	TRPO	A2C	Tie
(1) average expected reward over all epochs	26	22	10	3	1
(2) average expected reward over last 100 epochs	29	17	12	3	1

Table 1: The number of highest evaluation scores obtained by each algorithm across all games

7.4 ABSOLUTE PERFORMANCE COMPARISON

We use large probability factor k in practical implementation, which means we are close to optimizing the lower bound for all samples. Thus we use another two similar metrics to evaluate the effectiveness of algorithms for lower bound lifting: (iii) average worst reward per episode over **all epochs of training**, and (iv) average worst reward per episode over **last 20 epochs of training**. We summarize the absolute performance of APO in Atari games and GUARD in Table 2, which answers **Q2**.

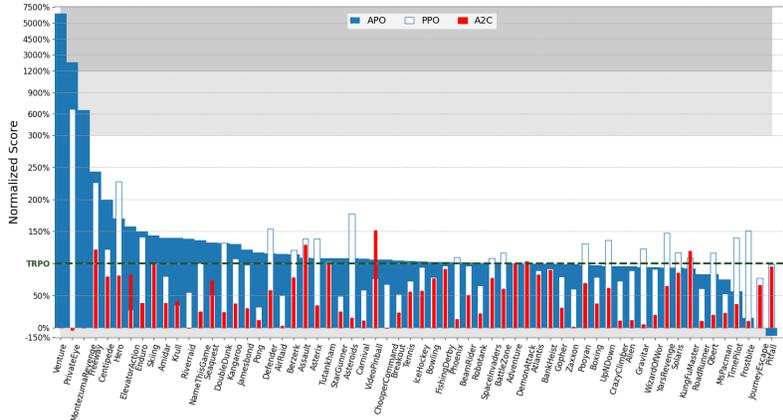


Figure 7: Stacked bar chart for all 62 atari games

	APO	PPO	TRPO	A2C	Tie
(1) average worst reward over all epochs(Atari)	26	24	7	3	2
(2) average worst reward over last 100 epochs(Atari)	27	20	10	3	2
(3) average worst reward over all epochs(GUARD)	7	3	1	1	0
(4) average worst reward over last 20 epochs(GUARD)	8	2	2	0	0

Table 2: The number of highest evaluation scores obtained by each algorithm across all games

7.5 ABLATION ON H_{max} HYPERPARAMETER TRICK

We chose Riverraid of discrete tasks and PUSH_Ant of continuous tasks to perform ablation experiments against the $|H(s, a, s')|_{max}$ implementation. Figure 8 shows that although both boosts are similar in the early stages of tasks, hypeterparameter method can more consistently converge to a higher reward value. Thus, the figures and description answer **Q3**.

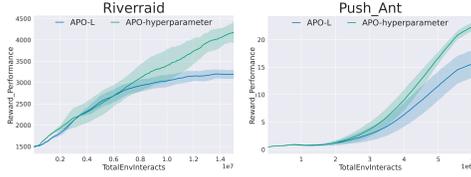


Figure 8: Ablation on H_{max} hyperparameter trick

7.6 ABLATION ON PROBABILITY FACTOR k

For ablation, we selected Riverraid to investigate the impact of different choices for the probability factor k . As illustrated in Figure 9, when k takes on a very small value, indicating optimization of only a limited portion of performance samples, the effectiveness diminishes. This is attributed to the loss of control over the lower bound of near-total performance samples. Conversely, when k becomes excessively large, the optimization shifts its focus towards the most extreme worst-case performance scenarios. This ultra-conservative approach tends to render the overall optimization less effective. Therefore, a moderate choice of k will be favorable to the overall improvement of the effect, which answers **Q4**.

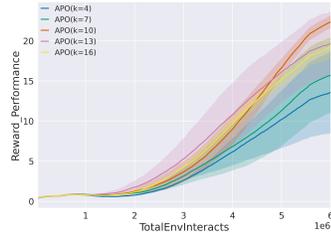


Figure 9: Ablation on probability factor k

8 CONCLUSION

This paper proposed APO, the first general-purpose policy search algorithm that improve both expected performance and absolute performance. Our approach is grounded by a pioneering theoretical advancement, where maximization of a specific objective function ensures monotonic improvement of expected performance and lower bound of near-total performance samples. We demonstrate APO’s effectiveness on challenging continuous control benchmark tasks and playing Atari games, showing its significant performance improvement compared to existing methods and ability to enhance both expected performance and worst-case performance.

REFERENCES

- Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. Constrained policy optimization. In *International conference on machine learning*, pp. 22–31. PMLR, 2017.
- Ilge Akkaya, Marcin Andrychowicz, Maciek Chociej, Mateusz Litwin, Bob McGrew, Arthur Petron, Alex Paino, Matthias Plappert, Glenn Powell, Raphael Ribas, et al. Solving rubik’s cube with a robot hand. *arXiv preprint arXiv:1910.07113*, 2019.
- M. G. Bellemare, Y. Naddaf, J. Veness, and M. Bowling. The arcade learning environment: an evaluation platform for general agents. *Journal of Artificial Intelligence Research*, pp. 253–279, Jul 2018. doi: 10.1613/jair.3912. URL <http://dx.doi.org/10.1613/jair.3912>.
- Marc G Bellemare, Yavar Naddaf, Joel Veness, and Michael Bowling. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47: 253–279, 2013.
- David R. Brillinger. Information and Information Stability of Random Variables and Processes. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 13(2):134–135, 12 2018. ISSN 0035-9254. doi: 10.2307/2985711. URL <https://doi.org/10.2307/2985711>.
- Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016a.
- Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv: Learning, arXiv: Learning*, Jun 2016b.
- Yan Duan, Xi Chen, Rein Houthoofd, John Schulman, and Pieter Abbeel. Benchmarking deep reinforcement learning for continuous control. In *International conference on machine learning*, pp. 1329–1338. PMLR, 2016.
- Scott Fujimoto, Herke Hoof, and David Meger. Addressing function approximation error in actor-critic methods. In *International conference on machine learning*, pp. 1587–1596. PMLR, 2018.
- Shixiang Gu, Timothy Lillicrap, Ilya Sutskever, and Sergey Levine. Continuous deep q-learning with model-based acceleration. In *International conference on machine learning*, pp. 2829–2838. PMLR, 2016.
- Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, pp. 1861–1870. PMLR, 2018.
- Matthew Hausknecht and Peter Stone. Deep recurrent q-learning for partially observable mdps. In *2015 aaai fall symposium series*, 2015.
- Matteo Hessel, Joseph Modayil, Hado Van Hasselt, Tom Schaul, Georg Ostrovski, Will Dabney, Dan Horgan, Bilal Piot, Mohammad Azar, and David Silver. Rainbow: Combining improvements in deep reinforcement learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.
- Elia Kaufmann, Leonard Bauersfeld, Antonio Loquercio, Matthias Müller, Vladlen Koltun, and Davide Scaramuzza. Champion-level drone racing using deep reinforcement learning. *Nature*, 620(7976):982–987, 2023.
- Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- Zichuan Lin, Xiapeng Wu, Mingfei Sun, Deheng Ye, Qiang Fu, Wei Yang, and Wei Liu. Sample dropout: A simple yet effective variance reduction technique in deep policy optimization. *arXiv preprint arXiv:2302.02299*, 2023.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013a.

- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013b.
- Volodymyr Mnih, Adrià Puigdomènech Badia, Mehdi Mirza, Alex Graves, Timothy P. Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. *arXiv: Learning, arXiv: Learning*, Feb 2016.
- Joni Pajarinen, Hong Linh Thai, Riad Akrou, Jan Peters, and Gerhard Neumann. Compatible natural gradient policy search. *Machine Learning*, 108:1443–1466, 2019.
- Matteo Papini, Damiano Binaghi, Giuseppe Canonaco, Matteo Pirotta, and Marcello Restelli. Stochastic variance-reduced policy gradient. In *International conference on machine learning*, pp. 4026–4035. PMLR, 2018.
- Jan Peters and Stefan Schaal. Reinforcement learning of motor skills with policy gradients. *Neural networks*, 21(4):682–697, 2008.
- James Queeney, Yannis Paschalidis, and Christos G. Cassandras. Generalized proximal policy optimization with sample reuse. *Advances in Neural Information Processing Systems*, 34:11909–11919, 2021.
- John G. Saw, Mark CK Yang, and Tse Chin Mo. Chebyshev inequality with estimated mean and variance. *The American Statistician*, 38(2):130–132, 1984.
- John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International conference on machine learning*, pp. 1889–1897. PMLR, 2015.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- John Schulman, Barret Zoph, Christina Kim, Jacob Hilton, Jacob Menick, Jiayi Weng, Juan Felipe Ceron Uribe, Liam Fedus, Luke Metz, Michael Pokorny, et al. Chatgpt: Optimizing language models for dialogue. *OpenAI blog*, 2022.
- David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. Deterministic policy gradient algorithms. In *International conference on machine learning*, pp. 387–395. Pmlr, 2014.
- David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.
- Matthew J. Sobel. The variance of discounted markov decision processes. *Journal of Applied Probability*, 19(4):794–802, 1982.
- Mingfei Sun, Benjamin Ellis, Anuj Mahajan, Sam Devlin, Katja Hofmann, and Shimon Whiteson. Trust-region-free policy optimization for stochastic policies. *arXiv preprint arXiv:2302.07985*, 2023.
- Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Sep 2012. doi: 10.1109/iros.2012.6386109. URL <http://dx.doi.org/10.1109/iros.2012.6386109>.
- Marcin B. Tomczak, Dongho Kim, Peter Vrancx, and Kee-Eung Kim. Policy optimization through approximate importance sampling. *arXiv preprint arXiv:1910.03857*, 2019.
- Hado van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning, 2015.
- Yuhui Wang, Hao He, and Xiaoyang Tan. Truly proximal policy optimization. In *Uncertainty in Artificial Intelligence*, pp. 113–122. PMLR, 2020.

Haotian Xu, Zheng Yan, Junyu Xuan, Guangquan Zhang, and Jie Lu. Improving proximal policy optimization with alpha divergence. *Neurocomputing*, 534:94–105, 2023.

Tianbing Xu, Qiang Liu, and Jian Peng. Stochastic variance reduction for policy gradient estimation. *arXiv preprint arXiv:1710.06034*, 2017.

Huizhuo Yuan, Chris Junchi Li, Yuhao Tang, and Yuren Zhou. Policy optimization via stochastic recursive gradient algorithm, 2019. In URL <https://openreview.net/forum>.

Weiye Zhao, Rui Chen, Yifan Sun, Ruixuan Liu, Tianhao Wei, and Changliu Liu. Guard: A safe reinforcement learning benchmark. *arXiv preprint arXiv:2305.13681*, 2023.

A MONOTONIC IMPROVEMENT OF ABSOLUTE PERFORMANCE

A.1 PRELIMINARIES

First we define $\mathcal{R}_\pi(s) = \sum_{t=0}^{\infty} \gamma^t R(s_t, a_t, s_{t+1})$ as infinite-horizon discounted return starts at state s and define expected return $V_\pi(s) = \mathbb{E}_{\hat{\tau} \sim \pi}(\mathcal{R}_\pi(s))$ as the value of state s . Then for all trajectories $\hat{\tau} \sim \pi$ start from state $s_0 \sim \mu$, the expectation and variance of $\mathcal{R}_\pi(s_0)$ can be respectively defined as $\mathcal{J}(\pi)$ and $\mathcal{V}(\pi)$. Formally:

$$\mathcal{J}(\pi) = \mathbb{E}_{\substack{s_0 \sim \mu \\ \hat{\tau} \sim \pi}}[\mathcal{R}_\pi(s_0)] = \mathbb{E}_{s_0 \sim \mu}[V_\pi(s_0)] \quad (16)$$

$$\begin{aligned} \mathcal{V}(\pi) &= \mathbb{E}_{\substack{s_0 \sim \mu \\ \hat{\tau} \sim \pi}}[(\mathcal{R}_\pi(s_0) - \mathcal{J}(\pi))^2] \quad (17) \\ &= \mathbb{E}_{s_0 \sim \mu} \left[\mathbb{E}_{\hat{\tau} \sim \pi} [\text{Var}[\mathcal{R}_\pi(s_0)]] + \left[\mathbb{E}_{\hat{\tau} \sim \pi} [\mathcal{R}_\pi(s_0)] \right]^2 \right] - \mathcal{J}(\pi)^2 \\ &= \mathbb{E}_{s_0 \sim \mu} \left[\mathbb{E}_{\hat{\tau} \sim \pi} [\text{Var}[\mathcal{R}_\pi(s_0)]] + V_\pi(s_0)^2 \right] - \mathcal{J}(\pi)^2 \\ &= \mathbb{E}_{s_0 \sim \mu} \left[\mathbb{E}_{\hat{\tau} \sim \pi} [\text{Var}[\mathcal{R}_\pi(s_0)]] \right] + \mathbb{E}_{s_0 \sim \mu} [V_\pi(s_0)^2] - \mathcal{J}(\pi)^2 \\ &= \underbrace{\mathbb{E}_{s_0 \sim \mu} \left[\mathbb{E}_{\hat{\tau} \sim \pi} [\text{Var}[\mathcal{R}_\pi(s_0)]] \right]}_{\text{MeanVariance}} + \underbrace{\mathbb{E}_{s_0 \sim \mu} [V_\pi(s_0)^2]}_{\text{VarianceMean}} - \mathcal{J}(\pi)^2 \end{aligned}$$

Note that for the derivation of $\mathcal{V}(\pi)$ we treat the return of all trajectories as a mixture of one-dimensional distributions. Each distribution consists of the returns of trajectories from the same start state. The variance can then be divided into two parts:

1. **MeanVariance** reflects the expected variance of the return over different start states.
2. **VarianceMean** reflects the variance of the average return of different start states.

Proposition 1. $\mathcal{B}_k(\pi) = \mathcal{J}(\pi) - k\mathcal{V}(\pi)$ is guaranteed to be an absolute performance bound defined by Definition 1.

Proof. According to Chebyshev’s inequality theory Saw et al. (1984), if random variable $\mathcal{R}_\pi(s_0)$ has finite non-zero variance $\mathcal{V}(\pi)$ and finite expected value $\mathcal{J}(\pi)$. Then for any real number $k > 0$ (but only the case $k > 1$ is useful), following inequality holds:

$$Pr(\mathcal{R}_\pi(s_0) < \mathcal{J}(\pi) - k\mathcal{V}(\pi)) < Pr(|\mathcal{R}_\pi(s_0) - \mathcal{J}(\pi)| \geq k\mathcal{V}(\pi)) \leq \frac{1}{k^2}, \quad (18)$$

Then we have:

$$Pr(\mathcal{R}_\pi(s_0) \geq \mathcal{B}_k(\pi)) \geq 1 - \frac{1}{k^2} \quad (19)$$

□

Now to prove Theorem 1, we need to prove the optimization object in equation 4 serves as the lower bound of $\mathcal{B}_k(\pi')$. To do this, we first try to obtain the following terms:

1. Upper bound of **MeanVariance** with π' . (Appendix A.2)
2. Upper bound of **VarianceMean** with π' . (Appendix A.3)
3. Lower bound of $\mathcal{J}(\pi')$. (Appendix A.4)

And then we can proof our theorem by leveraging this lower bound in Appendix A.5

A.2 MEANVARIANCE BOUND

Lemma 1 (Bound of MeanVariance). *Denote **MeanVariance** of policy π as $MV_\pi = \mathbb{E}_{s_0 \sim \mu} [\text{Var}[\mathcal{R}_\pi(s_0)]]$. Given two policies π', π that are α -coupled, the following bound holds:*

$$\begin{aligned} |MV_{\pi'} - MV_\pi| &\leq \|\mu^T\|_\infty \left(\frac{1}{1-\gamma^2} \max_s \left| \mathbb{E}_{\substack{a \sim \pi \\ s' \sim P}} \left[\left(\frac{\pi'(a|s)}{\pi(a|s)} - 1 \right) A_\pi(s, a, s')^2 \right] \right. \right. \\ &\quad \left. \left. + 2 \mathbb{E}_{\substack{a \sim \pi \\ s' \sim P}} \left[\left(\frac{\pi'(a|s)}{\pi(a|s)} \right) A_\pi(s, a, s') \right] |H(s, a, s')|_{max} + |H(s, a, s')|_{max}^2 \right| \right. \\ &\quad \left. + \frac{2\gamma^2}{1-\gamma^2} \sqrt{\frac{1}{2} \mathcal{D}_{KL}(\pi' \|\pi)[s] \cdot \|\Omega_\pi\|_\infty} \right) \end{aligned} \quad (20)$$

where $\Omega_\pi = \begin{bmatrix} \omega_\pi(s^1) \\ \omega_\pi(s^2) \\ \vdots \end{bmatrix}$ and $\omega_\pi(s) = \mathbb{E}_{\substack{a \sim \pi \\ s' \sim P}} [Q_\pi(s, a, s')^2] - V_\pi(s)^2$ is defined as the variance of the state-action value function Q_π at state s .

$$|H(s, a, s')|_{max} = |L(s, a, s')| + \frac{2(1+\gamma)\gamma\epsilon}{(1-\gamma)^2} \mathcal{D}_{KL}(\pi' \|\pi)[s] \quad (21)$$

$$\begin{aligned} L(s, a, s') &= \gamma \mathbb{E}_{\substack{s_0=s' \\ \hat{\tau} \sim \pi}} \left[\sum_{t=0}^{\infty} \gamma^t \bar{A}_\pi(s_t) \right] - \mathbb{E}_{\substack{s_0=s \\ \hat{\tau} \sim \pi}} \left[\sum_{t=0}^{\infty} \gamma^t \bar{A}_\pi(s_t) \right] \\ \epsilon &= \max_{s,a} |A_\pi(s, a)| \end{aligned}$$

Proof. According to [Theorem 1, (Sobel, 1982)], the following proposition holds:

Proposition 2. Define $\mathbf{V}_\pi = \begin{bmatrix} \text{Var}[\mathcal{R}_\pi(s^1)] \\ \text{Var}[\mathcal{R}_\pi(s^2)] \\ \vdots \end{bmatrix}$, where $\mathbb{E}_{s_0 \sim \mu} [\text{Var}[\mathcal{R}_\pi(s_0)]] = \mu^\top \mathbf{V}_\pi$, we have

$$\mathbf{V}_\pi = (I - \gamma^2 P_\pi)^{-1} \Omega_\pi. \quad (22)$$

Therefore, we have

$$\begin{aligned} \mathbb{E}_{s_0 \sim \mu} [\text{Var}[\mathcal{R}_\pi(s_0)]] &= \mu^\top \mathbf{V}_\pi \\ &= \mu^\top ((I - \gamma^2 P_\pi)^{-1} \Omega_\pi) \end{aligned} \quad (23)$$

The divergence of **MeanVariance** we want to bound can be written as:

$$\begin{aligned} |MV_{\pi'} - MV_\pi| &= \left| \mathbb{E}_{s_0 \sim \mu} [\text{Var}[\mathcal{R}_{\pi'}(s_0)]] - \mathbb{E}_{s_0 \sim \mu} [\text{Var}[\mathcal{R}_\pi(s_0)]] \right| \\ &= \|\mu^\top (\mathbf{V}_{\pi'} - \mathbf{V}_\pi)\|_1 \\ &\leq \|\mu^\top\|_\infty \|\mathbf{V}_{\pi'} - \mathbf{V}_\pi\|_1 \end{aligned} \quad (24)$$

To bound $\|\mathbf{V}_{\pi'} - \mathbf{V}_\pi\|_1$, consider the following conditions:

$$\begin{aligned} \mathbf{V}_\pi - \gamma^2 P_\pi \mathbf{V}_\pi &= \Omega_\pi \\ \mathbf{V}_{\pi'} - \gamma^2 P_{\pi'} \mathbf{V}_{\pi'} &= \Omega_{\pi'} \end{aligned} \quad (25)$$

Let $G_\pi = (I - \gamma^2 P_\pi)^{-1}$, then \mathbf{V}_π can be written as:

$$\mathbf{V}_\pi = G_\pi \Omega_\pi \quad (26)$$

Then we have:

$$\begin{aligned}
\mathbf{V}_{\pi'} - \mathbf{V}_{\pi} &= G_{\pi'} \boldsymbol{\Omega}_{\pi'} - G_{\pi} \boldsymbol{\Omega}_{\pi} \\
&= G_{\pi'} \boldsymbol{\Omega}_{\pi'} - G_{\pi'} \boldsymbol{\Omega}_{\pi} + G_{\pi'} \boldsymbol{\Omega}_{\pi} - G_{\pi} \boldsymbol{\Omega}_{\pi} \\
&= G_{\pi'} (\boldsymbol{\Omega}_{\pi'} - \boldsymbol{\Omega}_{\pi}) + (G_{\pi'} - G_{\pi}) \boldsymbol{\Omega}_{\pi}
\end{aligned} \tag{27}$$

Now $\|\mathbf{V}_{\pi'} - \mathbf{V}_{\pi}\|_1$ can be bounded by:

$$\begin{aligned}
\|\mathbf{V}_{\pi'} - \mathbf{V}_{\pi}\|_1 &\leq \|G_{\pi'}\|_1 \|\boldsymbol{\Omega}_{\pi'} - \boldsymbol{\Omega}_{\pi}\|_{\infty} + \|G_{\pi'} - G_{\pi}\|_1 \|\boldsymbol{\Omega}_{\pi}\|_{\infty} \\
&\leq \|G_{\pi'}\|_1 \|\boldsymbol{\Omega}_{\pi'} - \boldsymbol{\Omega}_{\pi}\|_{\infty} + \|G_{\pi'} - G_{\pi}\|_1 \|\boldsymbol{\Omega}_{\pi}\|_{\infty}
\end{aligned} \tag{28}$$

Since we already know $\|G_{\pi}\|_1 = \|G_{\pi'}\|_1 = (1 - \gamma^2)^{-1}$ and $\|\boldsymbol{\Omega}_{\pi}\|_{\infty}$ can be obtained with $\boldsymbol{\Omega}_{\pi}$. We only need to tackle with $\|G_{\pi'} - G_{\pi}\|_1$ and $\|\boldsymbol{\Omega}_{\pi'} - \boldsymbol{\Omega}_{\pi}\|_{\infty}$.

To get $\|G_{\pi'} - G_{\pi}\|_1$, we have

$$\begin{aligned}
G_{\pi}^{-1} - G_{\pi'}^{-1} &= (I - \gamma^2 P_{\pi}) - (I - \gamma^2 P'_{\pi}) \\
&= \gamma^2 (P'_{\pi} - P_{\pi}) \\
&= \gamma^2 \Delta
\end{aligned} \tag{29}$$

where $\Delta = P'_{\pi} - P_{\pi}$. Then we have

$$G_{\pi'} - G_{\pi} = \gamma^2 G_{\pi'} \Delta G_{\pi} \tag{30}$$

With $\|\Delta\|_1 = 2\mathcal{D}_{TV}(\pi' \|\pi)[s]$, we have

$$\begin{aligned}
\|\gamma^2 G_{\pi'} \Delta G_{\pi}\|_1 &\leq \gamma^2 \|G_{\pi'}\|_1 \|\Delta\|_1 \|G_{\pi}\|_1 \\
&= \gamma^2 \cdot \frac{1}{1 - \gamma^2} \cdot 2\mathcal{D}_{TV}(\pi' \|\pi)[s] \cdot \frac{1}{1 - \gamma^2} \\
&= \frac{2\gamma^2}{1 - \gamma^2} \mathcal{D}_{TV}(\pi' \|\pi)[s]
\end{aligned} \tag{31}$$

And according to Brillinger (2018) $\mathcal{D}_{TV}(\pi' \|\pi)[s] \leq \sqrt{\frac{1}{2} \mathcal{D}_{KL}(\pi' \|\pi)[s]}$, so far we have

$$\begin{aligned}
&\left| \mathbb{E}_{s_0 \sim \mu} [\text{Var}[\mathcal{R}_{\pi'}(s_0)]] - \mathbb{E}_{s_0 \sim \mu} [\text{Var}[\mathcal{R}_{\pi}(s_0)]] \right| \\
&= \|\boldsymbol{\mu}^{\top} (\mathbf{V}_{\pi'} - \mathbf{V}_{\pi})\|_1 \\
&\leq \|\boldsymbol{\mu}^{\top}\|_{\infty} \|\mathbf{V}_{\pi'} - \mathbf{V}_{\pi}\|_1 \\
&\leq \|\boldsymbol{\mu}^{\top}\|_{\infty} \left(\frac{1}{1 - \gamma^2} \|\boldsymbol{\Omega}_{\pi'} - \boldsymbol{\Omega}_{\pi}\|_{\infty} + \frac{2\gamma^2}{1 - \gamma^2} \mathcal{D}_{TV}(\pi' \|\pi)[s] \|\boldsymbol{\Omega}_{\pi}\|_{\infty} \right) \\
&\leq \|\boldsymbol{\mu}^{\top}\|_{\infty} \left(\frac{1}{1 - \gamma^2} \|\boldsymbol{\Omega}_{\pi'} - \boldsymbol{\Omega}_{\pi}\|_{\infty} + \frac{2\gamma^2}{1 - \gamma^2} \sqrt{\frac{1}{2} \mathcal{D}_{KL}(\pi' \|\pi)[s]} \cdot \|\boldsymbol{\Omega}_{\pi}\|_{\infty} \right)
\end{aligned} \tag{32}$$

To address $\|\boldsymbol{\Omega}_{\pi'} - \boldsymbol{\Omega}_{\pi}\|_{\infty}$, we notice that $\omega_{\pi}(s) = \mathbb{V}_{a \sim \pi} [Q_{\pi}(s, a, s')] = \mathbb{V}_{a \sim \pi} [A_{\pi}(s, a, s')]$, which means:

$$\|\boldsymbol{\Omega}_{\pi'} - \boldsymbol{\Omega}_{\pi}\|_{\infty} = \max_s \left| \mathbb{V}_{\substack{a \sim \pi' \\ s' \sim P}} [A_{\pi'}(s, a, s')] - \mathbb{V}_{\substack{a \sim \pi \\ s' \sim P}} [A_{\pi}(s, a, s')] \right| \tag{33}$$

Where

$$\begin{aligned} A_\pi(s, a, s') &= R(s, a, s') + \gamma V_\pi(s') - V_\pi(s) \\ A_{\pi'}(s, a, s') &= R(s, a, s') + \gamma V_{\pi'}(s') - V_{\pi'}(s) \end{aligned} \quad (34)$$

Define $H(s, a, s') = A_{\pi'}(s, a, s') - A_\pi(s, a, s')$, we have:

$$H(s, a, s') = \gamma(V_{\pi'}(s') - V_\pi(s')) - (V_{\pi'}(s) - V_\pi(s)) \quad (35)$$

Similar to TRPO Schulman et al. (2015):

$$\begin{aligned} & \mathbb{E}_{\substack{s_0=s \\ \hat{\tau} \sim \pi'}} \left[\sum_{t=0}^{\infty} \gamma^t A_\pi(s_t, a_t, s_{t+1}) \right] \\ &= \mathbb{E}_{\substack{s_0=s \\ \hat{\tau} \sim \pi'}} \left[\sum_{t=0}^{\infty} \gamma^t (R_\pi(s_t, a_t, s_{t+1}) + \gamma V_\pi(s_{t+1}) - V_\pi(s_t)) \right] \\ &= \mathbb{E}_{\substack{s_0=s \\ \hat{\tau} \sim \pi'}} \left[-V_\pi(s_0) + \sum_{t=0}^{\infty} \gamma^t R_\pi(s_t, a_t, s_{t+1}) \right] \\ &= \mathbb{E}_{s_0=s} \left[-V_\pi(s_0) \right] + \mathbb{E}_{\substack{s_0=s \\ \hat{\tau} \sim \pi'}} \left[\sum_{t=0}^{\infty} \gamma^t R_\pi(s_t, a_t, s_{t+1}) \right] \\ &= -V_\pi(s) + V_{\pi'}(s) \end{aligned} \quad (36)$$

Then $H(s, a, s')$ can be written as:

$$H(s, a, s') = \gamma \mathbb{E}_{\substack{s_0=s' \\ \hat{\tau} \sim \pi'}} \left[\sum_{t=0}^{\infty} \gamma^t A_\pi(s_t, a_t, s_{t+1}) \right] - \mathbb{E}_{\substack{s_0=s \\ \hat{\tau} \sim \pi'}} \left[\sum_{t=0}^{\infty} \gamma^t A_\pi(s_t, a_t, s_{t+1}) \right] \quad (37)$$

Define $\bar{A}(s)$ to be the expected advantage of π' over π at state s :

$$\bar{A}(s) = \mathbb{E}_{a \sim \pi'} \left[A_\pi(s, a) \right] \quad (38)$$

Now $H(s, a, s')$ can be written as:

$$H(s, a, s') = \gamma \mathbb{E}_{\substack{s_0=s' \\ \hat{\tau} \sim \pi'}} \left[\sum_{t=0}^{\infty} \gamma^t \bar{A}_\pi(s_t) \right] - \mathbb{E}_{\substack{s_0=s \\ \hat{\tau} \sim \pi'}} \left[\sum_{t=0}^{\infty} \gamma^t \bar{A}_\pi(s_t) \right] \quad (39)$$

Define $L(s, a, s')$ as:

$$L(s, a, s') = \gamma \mathbb{E}_{\substack{s_0=s' \\ \hat{\tau} \sim \pi}} \left[\sum_{t=0}^{\infty} \gamma^t \bar{A}_\pi(s_t) \right] - \mathbb{E}_{\substack{s_0=s \\ \hat{\tau} \sim \pi}} \left[\sum_{t=0}^{\infty} \gamma^t \bar{A}_\pi(s_t) \right] \quad (40)$$

With $\epsilon = \max_{s,a} |A_\pi(s, a)|$ we have:

$$\begin{aligned}
& |H(s, a, s') - L(s, a, s')| \tag{41} \\
&= \left| \gamma \left(\mathbb{E}_{\substack{s_0=s', \\ \hat{\tau} \sim \pi'}} \left[\sum_{t=0}^{\infty} \gamma^t \bar{A}_\pi(s_t) \right] - \mathbb{E}_{\substack{s_0=s', \\ \hat{\tau} \sim \pi}} \left[\sum_{t=0}^{\infty} \gamma^t \bar{A}_\pi(s_t) \right] \right) \right. \\
&\quad \left. - \left(\mathbb{E}_{\substack{s_0=s, \\ \hat{\tau} \sim \pi'}} \left[\sum_{t=0}^{\infty} \gamma^t \bar{A}_\pi(s_t) \right] - \mathbb{E}_{\substack{s_0=s, \\ \hat{\tau} \sim \pi}} \left[\sum_{t=0}^{\infty} \gamma^t \bar{A}_\pi(s_t) \right] \right) \right| \\
&\leq \gamma \left| \mathbb{E}_{\substack{s_0=s', \\ \hat{\tau} \sim \pi'}} \left[\sum_{t=0}^{\infty} \gamma^t \bar{A}_\pi(s_t) \right] - \mathbb{E}_{\substack{s_0=s', \\ \hat{\tau} \sim \pi}} \left[\sum_{t=0}^{\infty} \gamma^t \bar{A}_\pi(s_t) \right] \right| \\
&\quad + \left| \mathbb{E}_{\substack{s_0=s, \\ \hat{\tau} \sim \pi'}} \left[\sum_{t=0}^{\infty} \gamma^t \bar{A}_\pi(s_t) \right] - \mathbb{E}_{\substack{s_0=s, \\ \hat{\tau} \sim \pi}} \left[\sum_{t=0}^{\infty} \gamma^t \bar{A}_\pi(s_t) \right] \right| \\
&\leq \frac{4\gamma(1+\gamma)\epsilon}{(1-\gamma)^2} (\mathcal{D}_{TV}(\pi' \parallel \pi)[s])^2
\end{aligned}$$

Then according to Brillinger (2018) $\mathcal{D}_{TV}(\pi' \parallel \pi)[s] \leq \sqrt{\frac{1}{2} \mathcal{D}_{KL}(\pi' \parallel \pi)[s]}$, we can then bound $|H(s, a, s')|$ with:

$$\begin{aligned}
|H(s, a, s')| &\leq |L(s, a, s')| + \frac{4\gamma(1+\gamma)\epsilon}{(1-\gamma)^2} (\mathcal{D}_{TV}(\pi' \parallel \pi)[s])^2 \tag{42} \\
&\leq |L(s, a, s')| + \frac{2\gamma(1+\gamma)\epsilon}{(1-\gamma)^2} \mathcal{D}_{KL}(\pi' \parallel \pi)[s] \doteq |H(s, a, s')|_{max}
\end{aligned}$$

With $A_{\pi'}(s, a, s') = A_\pi(s, a, s') + H(s, a, s')$ we have:

$$\begin{aligned}
& \mathbb{V}ar_{\substack{a \sim \pi', \\ s' \sim P}} [A_{\pi'}(s, a, s')] - \mathbb{V}ar_{\substack{a \sim \pi, \\ s' \sim P}} [A_\pi(s, a, s')] \tag{43} \\
&= \mathbb{E}_{\substack{a \sim \pi', \\ s' \sim P}} [A_{\pi'}(s, a, s')^2] - \mathbb{E}_{\substack{a \sim \pi, \\ s' \sim P}} [A_\pi(s, a, s')^2] \\
&= \mathbb{E}_{\substack{a \sim \pi', \\ s' \sim P}} [(A_\pi(s, a, s') + H(s, a, s'))^2] - \mathbb{E}_{\substack{a \sim \pi, \\ s' \sim P}} [A_\pi(s, a, s')^2] \\
&= \mathbb{E}_{\substack{a \sim \pi', \\ s' \sim P}} [A_\pi(s, a, s')^2] - \mathbb{E}_{\substack{a \sim \pi, \\ s' \sim P}} [A_\pi(s, a, s')^2] + 2 \mathbb{E}_{\substack{a \sim \pi', \\ s' \sim P}} [A_\pi(s, a, s')H(s, a, s')] + \mathbb{E}_{\substack{a \sim \pi', \\ s' \sim P}} [H(s, a, s')^2] \\
&= \mathbb{E}_{\substack{a \sim \pi, \\ s' \sim P}} \left[\left(\frac{\pi'(a|s)}{\pi(a|s)} - 1 \right) A_\pi(s, a, s')^2 \right] + 2 \mathbb{E}_{\substack{a \sim \pi', \\ s' \sim P}} [A_\pi(s, a, s')H(s, a, s')] + \mathbb{E}_{\substack{a \sim \pi', \\ s' \sim P}} [H(s, a, s')^2] \\
&\leq \mathbb{E}_{\substack{a \sim \pi, \\ s' \sim P}} \left[\left(\frac{\pi'(a|s)}{\pi(a|s)} - 1 \right) A_\pi(s, a, s')^2 \right] + 2 \mathbb{E}_{\substack{a \sim \pi, \\ s' \sim P}} \left[\left(\frac{\pi'(a|s)}{\pi(a|s)} \right) A_\pi(s, a, s') \right] |H(s, a, s')|_{max} + |H(s, a, s')|_{max}^2
\end{aligned}$$

Then we can bound $\|\Omega_{\pi'} - \Omega_\pi\|_\infty$ with:

$$\begin{aligned}
& \|\Omega_{\pi'} - \Omega_\pi\|_\infty \tag{44} \\
&\leq \max_s \left| \mathbb{E}_{\substack{a \sim \pi, \\ s' \sim P}} \left[\left(\frac{\pi'(a|s)}{\pi(a|s)} - 1 \right) A_\pi(s, a, s')^2 \right] + 2 \mathbb{E}_{\substack{a \sim \pi, \\ s' \sim P}} \left[\left(\frac{\pi'(a|s)}{\pi(a|s)} \right) A_\pi(s, a, s') \right] |H(s, a, s')|_{max} + |H(s, a, s')|_{max}^2 \right|
\end{aligned}$$

By substituting Equation (45) into Equation (32), Lemma 1 is proved. \square

A.3 VARIANCEMEAN BOUND

Lemma 2 (Bound of VarianceMean). *Denote **VarianceMean** of policy π as $VM_\pi = \mathbb{V}_{s_0 \sim \mu} [V_\pi(s_0)]$. Given two policies π', π that are α -coupled, the **VarianceMean** of π' can be bounded by:*

$$VM_{\pi'} \leq \mathbb{E}_{s_0 \sim \mu} [V_\pi^2(s_0)] + \|\mu^T\|_\infty \max_s \left| |\eta(s)|_{max}^2 + 2|V_\pi(s)| \cdot |\eta(s)|_{max} \right| - \left(\min \{ \max \{ 0, \mathcal{J}_{\pi', \pi}^l \}, \mathcal{J}_{\pi', \pi}^u \} \right)^2 \quad (45)$$

Proof.

$$VM_{\pi'} = \mathbb{E}_{s_0 \sim \mu} [V_{\pi'}^2(s_0)] - \mathcal{J}(\pi')^2 \quad (46)$$

Since both terms on the right of Equation (46) are non-negative, we can bound $VM_{\pi'}$ with the upper bound of $\mathbb{E}_{s_0 \sim \mu} [V_{\pi'}^2(s_0)]$ and the lower bound of $\mathcal{J}(\pi')^2$.

Define $\mathbf{V}_\pi^2 = \begin{bmatrix} V_\pi^2(s^1) \\ V_\pi^2(s^2) \\ \vdots \end{bmatrix}$, where $\mathbb{E}_{s_0 \sim \mu} [V_\pi^2(s_0)] = \mu^\top \mathbf{V}_\pi^2$. Then we have

$$\begin{aligned} & \left| \mathbb{E}_{s_0 \sim \mu} [V_{\pi'}^2(s_0)] - \mathbb{E}_{s_0 \sim \mu} [V_\pi^2(s_0)] \right| \\ &= \|\mu^\top (\mathbf{V}_{\pi'}^2 - \mathbf{V}_\pi^2)\|_1 \\ &\leq \|\mu^\top\|_\infty \|\mathbf{V}_{\pi'}^2 - \mathbf{V}_\pi^2\|_1 \end{aligned} \quad (47)$$

To address $\|\mathbf{V}_{\pi'}^2 - \mathbf{V}_\pi^2\|_1$, we have:

$$V_{\pi'}^2(s) - V_\pi^2(s) = \left(V_{\pi'}(s) - V_\pi(s) \right) \left(V_{\pi'}(s) + V_\pi(s) \right) \quad (48)$$

According to equation 36:

$$\begin{aligned} V_{\pi'}(s) - V_\pi(s) &= \mathbb{E}_{\substack{s_0=s \\ \hat{\tau} \sim \pi}} \left[\sum_{t=0}^{\infty} \gamma^t A_\pi(s_t, a_t, s_{t+1}) \right] \\ &= \mathbb{E}_{\substack{s_0=s \\ \hat{\tau} \sim \pi'}} \left[\sum_{t=0}^{\infty} \gamma^t \bar{A}_\pi(s_t) \right] \\ &\doteq \eta(s) \end{aligned} \quad (49)$$

Define $L(s) = \mathbb{E}_{\substack{s_0=s \\ \hat{\tau} \sim \pi}} \left[\sum_{t=0}^{\infty} \gamma^t \bar{A}_\pi(s_t) \right]$, then we have:

$$|\eta(s) - L(s)| = \left| \mathbb{E}_{\substack{s_0=s \\ \hat{\tau} \sim \pi'}} \left[\sum_{t=0}^{\infty} \gamma^t \bar{A}_\pi(s_t) \right] - \mathbb{E}_{\substack{s_0=s \\ \hat{\tau} \sim \pi}} \left[\sum_{t=0}^{\infty} \gamma^t \bar{A}_\pi(s_t) \right] \right| \leq \frac{4\gamma\epsilon}{(1-\gamma)^2} (\mathcal{D}_{TV}(\pi' \parallel \pi)[s])^2 \quad (50)$$

And according to Brillinger (2018), we can bound $|\eta(s)|$ with:

$$\begin{aligned} |\eta(s)| &\leq |L(s)| + \frac{4\gamma\epsilon}{(1-\gamma)^2} (\mathcal{D}_{TV}(\pi' \parallel \pi)[s])^2 \\ &\leq |L(s)| + \frac{2\gamma\epsilon}{(1-\gamma)^2} \mathcal{D}_{KL}(\pi' \parallel \pi)[s] \doteq |\eta(s)|_{max} \end{aligned} \quad (51)$$

Further, we can obtain:

$$\begin{aligned} |V_{\pi'}(s) + V_\pi(s)| &\leq |V_{\pi'}(s)| + |V_\pi(s)| \\ &= |V_{\pi'}(s) - V_\pi(s)| + 2|V_\pi(s)| \\ &\leq |V_{\pi'}(s) - V_\pi(s)| + 2|V_\pi(s)| \\ &\leq |\eta(s)|_{max} + 2|V_\pi(s)| \end{aligned} \quad (52)$$

Thus the following inequality holds:

$$\begin{aligned}
& \|V_{\pi'}^2 - V_{\pi}^2\|_1 & (53) \\
& \leq \max_s \left| |V_{\pi'}(s) - V_{\pi}(s)| \cdot |V_{\pi'}(s) + V_{\pi}(s)| \right| \\
& \leq \max_s \left| |\eta(s)|_{max} \cdot (|\eta(s)|_{max} + 2|V_{\pi}(s)|) \right| \\
& = \max_s \left| |\eta(s)|_{max}^2 + 2|V_{\pi}(s)| \cdot |\eta(s)|_{max} \right|
\end{aligned}$$

Substitute Equation (53) into Equation (47) the upper bound of $\mathbb{E}_{s_0 \sim \mu} [V_{\pi'}^2(s_0)]$ is obtained:

$$\mathbb{E}_{s_0 \sim \mu} [V_{\pi'}^2(s_0)] \leq \mathbb{E}_{s_0 \sim \mu} [V_{\pi}^2(s_0)] + \|\mu^T\|_{\infty} \max_s \left| |\eta(s)|_{max}^2 + 2|V_{\pi}(s)| \cdot |\eta(s)|_{max} \right| \quad (54)$$

The lower bound of $\mathcal{J}(\pi')^2$ can then be obtained according to (Achiam et al., 2017):

$$\mathcal{J}(\pi')^2 \geq (\min \{ \max \{ 0, \mathcal{J}_{\pi', \pi}^l \}, \mathcal{J}_{\pi', \pi}^u \})^2 \quad (55)$$

where

$$\begin{aligned}
\mathcal{J}_{\pi', \pi}^l &= \mathcal{J}(\pi) + \frac{1}{1 - \gamma} \mathbb{E}_{\substack{s \sim d_{\pi}^{\pi} \\ a \sim \pi'}} \left[A_{\pi}(s, a) - \frac{2\gamma\epsilon^{\pi'}}{1 - \gamma} \sqrt{\frac{1}{2} \mathcal{D}_{KL}(\pi' \|\pi)[s]} \right] \\
\mathcal{J}_{\pi', \pi}^u &= \mathcal{J}(\pi) + \frac{1}{1 - \gamma} \mathbb{E}_{\substack{s \sim d_{\pi}^{\pi} \\ a \sim \pi'}} \left[A_{\pi}(s, a) + \frac{2\gamma\epsilon^{\pi'}}{1 - \gamma} \sqrt{\frac{1}{2} \mathcal{D}_{KL}(\pi' \|\pi)[s]} \right]
\end{aligned}$$

By substituting Equation (54) and Equation (55) into Equation (46) Lemma 2 is proved. \square

A.4 EXPECTATION BOUND

Lemma 3. For any policies π', π , with $\epsilon^{\pi'} \doteq \max_s |E_{a \sim \pi'} [A^{\pi}(s, a)]|$, and define $d^{\pi} = (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t P(s_t = s | \pi)$ as the discounted augmented state distribution using π , then the following bound holds:

$$\mathcal{J}(\pi') - \mathcal{J}(\pi) \geq \frac{1}{1 - \gamma} \mathbb{E}_{\substack{s \sim d_{\pi}^{\pi} \\ a \sim \pi'}} \left[A^{\pi}(s, a) - \frac{2\gamma\epsilon^{\pi'}}{1 - \gamma} \mathcal{D}_{TV}(\pi' \|\pi)[s] \right] \quad (56)$$

Proof. d^{π} we used is defined as

$$d^{\pi}(\hat{s}) = (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t P(s_t = s | \pi). \quad (57)$$

Then it allows us to express the expected discounted total reward compactly as:

$$\mathcal{J}(\pi) = \frac{1}{1 - \gamma} \mathbb{E}_{\substack{s \sim d_{\pi}^{\pi} \\ a \sim \pi \\ s' \sim P}} [R(s, a, s')], \quad (58)$$

where by $a \sim \pi$, we mean $a \sim \pi(\cdot | s)$, and by $s' \sim P$, we mean $s' \sim P(\cdot | s, a)$. We drop the explicit notation for the sake of reducing clutter, but it should be clear from context that a and s' depend on s .

Define $P(s' | s, a)$ is the probability of transitioning to state s' given that the previous state was s and the agent took action a at state s , and $\mu : \mathcal{S} \mapsto [0, 1]$ is the initial augmented state distribution. Let $p_{\pi}^t \in \mathbb{R}^{|\mathcal{S}|}$ denote the vector with components $p_{\pi}^t(s) = P(s_t = s | \pi)$, and let $P_{\pi} \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{S}|}$ denote

the transition matrix with components $P_\pi(s'|s) = \int P(s'|s, a)\pi(a|s)da$; then $p_\pi^t = P_\pi p_\pi^{t-1} = P_\pi^t \mu$ and

$$\begin{aligned} d^\pi &= (1 - \gamma) \sum_{t=0}^{\infty} (\gamma P_\pi)^t \mu \\ &= (1 - \gamma)(I - \gamma P_\pi)^{-1} \mu \end{aligned} \quad (59)$$

This formulation helps us easily obtain the following lemma.

Lemma 4. For any function $f : \mathcal{S} \mapsto \mathbb{R}$ and any policy π ,

$$(1 - \gamma) \mathbb{E}_{s \sim \mu} [f(s)] + \mathbb{E}_{\substack{s \sim d^\pi \\ a \sim \pi \\ s' \sim P}} [\gamma f(s')] - \mathbb{E}_{s \sim d^\pi} [f(s)] = 0. \quad (60)$$

Proof. Multiply both sides of equation 59 by $(I - \gamma P_\pi)$ and take the inner product with the vector $f \in \mathbb{R}^{|\mathcal{S}|}$. \square

Combining Lemma 4 with equation 58, we obtain the following, for any function f and any policy π :

$$\mathcal{J}(\pi) = \mathbb{E}_{s \sim \mu} [f(s)] + \frac{1}{1 - \gamma} \mathbb{E}_{\substack{s \sim d^\pi \\ a \sim \pi \\ s' \sim P}} [R(s, a, s') + \gamma f(s') - f(s)] \quad (61)$$

Then we will derive and present the new policy improvement bound. We will begin with a lemma:

Lemma 5. For any function $f \mapsto \mathbb{R}$ and any policies π' and π , define

$$L_{\pi, f}(\pi') \doteq \mathbb{E}_{\substack{s \sim d^{\pi'} \\ a \sim \pi' \\ s' \sim P}} \left[\left(\frac{\pi'(a|s)}{\pi(a|s)} - 1 \right) (R(s, a, s') + \gamma f(s') - f(s)) \right], \quad (62)$$

and $\epsilon_f^{\pi'} \doteq \max_s |\mathbb{E}_{a \sim \pi', s' \sim P} [R(s, a, s') + \gamma f(s') - f(s)]|$. Then the following bounds hold:

$$\mathcal{J}(\pi') - \mathcal{J}(\pi) \geq \frac{1}{1 - \gamma} \left(L_{\pi, f}(\pi') - 2\epsilon_f^{\pi'} D_{TV}(d^{\pi'} \| d^\pi) \right), \quad (63)$$

$$\mathcal{J}(\pi') - \mathcal{J}(\pi) \leq \frac{1}{1 - \gamma} \left(L_{\pi, f}(\pi') + 2\epsilon_f^{\pi'} D_{TV}(d^{\pi'} \| d^\pi) \right), \quad (64)$$

where D_{TV} is the total variational divergence. Furthermore, the bounds are tight (when $\pi' = \pi$, the LHS and RHS are identically zero).

Proof. First, for notational convenience, let $\delta_f(s, a, s') \doteq R(s, a, s') + \gamma f(s') - f(s)$. By equation 61, we obtain the identity

$$\mathcal{J}(\pi') - \mathcal{J}(\pi) = \frac{1}{1 - \gamma} \left(\mathbb{E}_{\substack{s \sim d^{\pi'} \\ a \sim \pi' \\ s' \sim P}} [\delta_f(s, a, s')] - \mathbb{E}_{\substack{s \sim d^\pi \\ a \sim \pi \\ s' \sim P}} [\delta_f(s, a, s')] \right) \quad (65)$$

Now, we restrict our attention to the first term in equation 65. Let $\dagger \delta_f^{\pi'} \in \mathbb{R}^{|\mathcal{S}|}$ denote the vector of components, where $\dagger \delta_f^{\pi'}(s) = \mathbb{E}_{a \sim \pi', s' \sim P} [\delta_f(s, a, s') | s]$. Observe that

$$\begin{aligned} \mathbb{E}_{\substack{s \sim d^{\pi'} \\ a \sim \pi' \\ s' \sim P}} [\delta_f(s, a, s')] &= \langle d^{\pi'}, \dagger \delta_f^{\pi'} \rangle \\ &= \langle d^\pi, \dagger \delta_f^{\pi'} \rangle + \langle d^{\pi'} - d^\pi, \dagger \delta_f^{\pi'} \rangle \end{aligned}$$

With the Hölder's inequality; for any $p, q \in [1, \infty]$ such that $\frac{1}{p} + \frac{1}{q} = 1$, we have

$$\langle d^\pi, \dagger \delta_f^{\pi'} \rangle + \|d^{\pi'} - d^\pi\|_p \|\dagger \delta_f^{\pi'}\|_q \geq \mathbb{E}_{\substack{s \sim d^{\pi'} \\ a \sim \pi' \\ s' \sim P}} [\delta_f(s, a, s')] \geq \langle d^\pi, \dagger \delta_f^{\pi'} \rangle - \|d^{\pi'} - d^\pi\|_p \|\dagger \delta_f^{\pi'}\|_q \quad (66)$$

We choose $p = 1$ and $q = \infty$; With $\|d^{\pi'} - d^\pi\|_1 = 2D_{TV}(d^{\pi'} \| d^\pi)$ and $\|\dagger \delta_f^{\pi'}\|_\infty = \epsilon_f^{\pi'}$, and by the importance sampling identity, we have

$$\begin{aligned} \langle d^\pi, \dagger \delta_f^{\pi'} \rangle &= \mathbb{E}_{\substack{s \sim d^\pi \\ a \sim \pi' \\ s' \sim P}} [\delta_f(s, a, s')] \\ &= \mathbb{E}_{\substack{s \sim d^\pi \\ a \sim \pi \\ s' \sim P}} \left[\left(\frac{\pi'(a|s)}{\pi(a|s)} \right) \delta_f(s, a, s') \right] \end{aligned} \quad (67)$$

After bringing equation 67, $\|d^{\pi'} - d^\pi\|_1$, $\|\dagger \delta_f^{\pi'}\|_\infty$ into equation 66, then subtract $\mathbb{E}_{\substack{s \sim d^\pi \\ a \sim \pi \\ s' \sim P}} [\delta_f(s, a, s')]$, the bounds are obtained. The lower bound leads to equation 63, and the upper bound leads to equation 64. \square

Then we will bound the divergence term, $\|d^{\pi'} - d^\pi\|_1$, i.e. $2D_{TV}(d^{\pi'} \| d^\pi)$.

Lemma 6. *The divergence between discounted future state visitation distributions, $\|d^{\pi'} - d^\pi\|_1$, is bounded by an average divergence of the policies π' and π :*

$$\|d^{\pi'} - d^\pi\|_1 \leq \frac{2\gamma}{1 - \gamma} \mathbb{E}_{s \sim d^\pi} [D_{TV}(\pi' \| \pi)[s]], \quad (68)$$

where $D_{TV}(\pi' \| \pi)[s] = \frac{1}{2} \sum_a |\pi'(a|s) - \pi(a|s)|$.

Proof. Firstly, we introduce an identity for the vector difference of the discounted future state visitation distributions on two different policies, π' and π . Define the matrices $G \doteq (I - \gamma P_\pi)^{-1}$, $\bar{G} \doteq (I - \gamma P_{\pi'})^{-1}$, and $\Delta = P_{\pi'} - P_\pi$. Then:

$$\begin{aligned} G^{-1} - \bar{G}^{-1} &= (I - \gamma P_\pi) - (I - \gamma P_{\pi'}) \\ &= \gamma \Delta, \end{aligned} \quad (69)$$

left-multiplying by G and right-multiplying by \bar{G} , we obtain

$$\bar{G} - G = \gamma \bar{G} \Delta G. \quad (70)$$

Thus, the following equality holds:

$$\begin{aligned} d^{\pi'} - d^\pi &= (1 - \gamma) (\bar{G} - G) \mu \\ &= \gamma (1 - \gamma) \bar{G} \Delta G \mu \\ &= \gamma \bar{G} \Delta d^\pi. \end{aligned} \quad (71)$$

Using equation 71, we obtain

$$\begin{aligned} \|d^{\pi'} - d^\pi\|_1 &= \gamma \|\bar{G} \Delta d^\pi\|_1 \\ &\leq \gamma \|\bar{G}\|_1 \|\Delta d^\pi\|_1, \end{aligned} \quad (72)$$

where $\|\bar{G}\|_1$ is bounded by:

$$\|\bar{G}\|_1 = \|(I - \gamma P_{\pi'})^{-1}\|_1 \leq \sum_{t=0}^{\infty} \gamma^t \|P_{\pi'}\|_1^t = (1 - \gamma)^{-1}. \quad (73)$$

Next, we bound $\|\Delta d_1^\pi\|$ as following:

$$\begin{aligned}
\|\Delta d_1^\pi\|_1 &= \sum_{s'} \left| \sum_s \Delta(s'|s) d^\pi(s) \right| \\
&\leq \sum_{s,s'} |\Delta(s'|s)| d^\pi(s) \\
&= \sum_{s,s'} \left| \sum_a P(s'|s,a) (\pi'(a|s) - \pi(a|s)) \right| d^\pi(s) \\
&\leq \sum_{s,a,s'} P(s'|s,a) |\pi'(a|s) - \pi(a|s)| d^\pi(s) \\
&= \sum_{s,a} |\pi'(a|s) - \pi(a|s)| d^\pi(s) \\
&= 2 \mathbb{E}_{s \sim d^\pi} [D_{TV}(\pi' || \pi)[s]].
\end{aligned} \tag{74}$$

By taking equation 74 and equation 73 into equation 72, this lemma is proved. \square

The new policy improvement bound follows immediately.

Lemma 7. For any function $f : \mathcal{S} \mapsto \mathbb{R}$ and any policies π' and π , define $\delta_f(s, a, s') \doteq R(s, a, s') + \gamma f(s') - f(s)$,

$$\begin{aligned}
\epsilon_f^{\pi'} &\doteq \max_s |\mathbb{E}_{a \sim \pi', s' \sim P} [\delta_f(s, a, s')]|, \\
L_{\pi, f}(\pi') &\doteq \mathbb{E}_{\substack{s \sim d^\pi \\ a \sim \pi \\ s' \sim P}} \left[\left(\frac{\pi'(a|s)}{\pi(a|s)} - 1 \right) \delta_f(s, a, s') \right], \text{ and} \\
D_{\pi, f}^\pm(\pi') &\doteq \frac{L_{\pi, f}(\pi')}{1 - \gamma} \pm \frac{2\gamma \epsilon_f^{\pi'}}{(1 - \gamma)^2} \mathbb{E}_{s \sim d^\pi} [D_{TV}(\pi' || \pi)[s]],
\end{aligned}$$

where $D_{TV}(\pi' || \pi)[s] = \frac{1}{2} \sum_a |\pi'(a|s) - \pi(a|s)|$ is the total variational divergence between action distributions at s . The following bounds hold:

$$D_{\pi, f}^+(\pi') \geq \mathcal{J}(\pi') - \mathcal{J}(\pi) \geq D_{\pi, f}^-(\pi').$$

Furthermore, the bounds are tight (when $\pi' = \pi$, all three expressions are identically zero)

Proof. Begin with the bounds from lemma 5 and bound the divergence $D_{TV}(d^{\pi'} || d^\pi)$ by lemma 6. \square

The choice of $f = V_\pi$ in lemma 7 leads to following inequality:

For any policies π', π , with $\epsilon^{\pi'} \doteq \max_s |E_{a \sim \pi'} [A_\pi(s, a)]|$, the following bound holds:

$$\mathcal{J}(\pi') - \mathcal{J}(\pi) \geq \frac{1}{1 - \gamma} \mathbb{E}_{\substack{s \sim d^\pi \\ a \sim \pi'}} \left[A_\pi(s, a) - \frac{2\gamma \epsilon^{\pi'}}{1 - \gamma} D_{TV}(\pi' || \pi)[s] \right]$$

At this point, the lemma 3 is proved. \square

A.5 PROOF OF THEOREM 1

With Lemma 2, Lemma 1 and Lemma 3, we have the following lower bound of absolute performance bound $\mathcal{B}_k(\pi')$:

$$\mathcal{B}_k(\pi') \geq \mathcal{J}_{\pi', \pi}^l - k (MV_{\pi', \pi} + VM_{\pi', \pi}) \tag{75}$$

where

$$\begin{aligned}
MV_{\pi',\pi} &= \frac{\|\mu^T\|_\infty}{1-\gamma^2} \mathbf{max}_s \left| \mathbb{E}_{\substack{a \sim \pi' \\ s' \sim P}} [A_\pi(s, a, s')^2] - \mathbb{E}_{\substack{a \sim \pi \\ s' \sim P}} [A_\pi(s, a, s')^2] + |H(s, a, s')|_{max}^2 \right. \\
&\quad \left. + 2 \mathbb{E}_{\substack{a \sim \pi' \\ s' \sim P}} [A_\pi(s, a, s')] \cdot |H(s, a, s')|_{max} \right| + MV_\pi + \frac{2\gamma^2 \|\mu^T\|_\infty}{1-\gamma^2} \sqrt{\frac{1}{2} \mathcal{D}_{KL}(\pi' \|\pi)[s]} \cdot \|\Omega_\pi\|_\infty \\
VM_{\pi',\pi} &= \|\mu^T\|_\infty \mathbf{max}_s \left| |\eta(s)|_{max}^2 + 2|V_\pi(s)| \cdot |\eta(s)|_{max} \right| + \mathbb{E}_{s_0 \sim \mu} [V_\pi^2(s_0)] - (\min \{ \max \{0, \mathcal{J}_{\pi',\pi}^l\}, \mathcal{J}_{\pi',\pi}^u \})^2 \\
\mathcal{J}_{\pi',\pi}^l &= \mathcal{J}(\pi) + \frac{1}{1-\gamma} \mathbb{E}_{\substack{s \sim d^\pi \\ a \sim \pi'}} \left[A_\pi(s, a) - \frac{2\gamma\epsilon^\pi}{1-\gamma} \sqrt{\frac{1}{2} \mathcal{D}_{KL}(\pi' \|\pi)[s]} \right]
\end{aligned}$$

We define $\mathcal{M}_j(\pi) = \mathcal{J}_{\pi,\pi_j}^l - k(MV_{\pi,\pi_j} + VM_{\pi,\pi_j})$, and it can be found that $\mathcal{B}_k(\pi_j) = \mathcal{M}_j(\pi_j)$. Then by Equation (75), we have $\mathcal{B}_k(\pi_{j+1}) \geq \mathcal{M}_j(\pi_{j+1})$ and the following holds:

$$\mathcal{B}_k(\pi_{j+1}) - \mathcal{B}_k(\pi_j) \geq \mathcal{M}_j(\pi_{j+1}) - \mathcal{M}_j(\pi_j) \quad (76)$$

Thus, by maximizing \mathcal{M}_j at each iteration, we guarantee that the true absolute performance \mathcal{B}_k is non-decreasing. So far Theorem 1 has been proved.

B APO PSEUDOCODE

Algorithm 1 Absolute Policy Optimization

Input: Initial policy $\pi_0 \in \Pi_\theta$.
for $j = 0, 1, 2, \dots$ **do**
 Sample trajectory $\tau \sim \pi_j = \pi_{\theta_j}$
 Estimate gradient $g \leftarrow \nabla_{\theta} O_{\pi,\pi_j} |_{\theta=\theta_j}$
 ▷ Define $O_{\pi,\pi_j} = \left(\frac{w_1+w_2}{1-\gamma} \mathbb{E}_{\substack{s \sim d^{\pi_j} \\ a \sim \pi}} [A_{\pi_j}(s, a)] - kw_2 (\overline{MV}_{\pi,\pi_j} + \overline{VM}_{\pi,\pi_j}) \right)$
 Estimate Hessian $H \leftarrow \nabla_{\theta}^2 \mathbb{E}_{s \sim \pi_j} [\mathcal{D}_{KL}(\pi \|\pi_j)[s]] |_{\theta=\theta_j}$
 Solve convex programming
 $\theta_{j+1}^* = \arg \max_{\theta} g^\top (\theta - \theta_j)$
 s.t. $\frac{1}{2} (\theta - \theta_j)^\top H (\theta - \theta_j) \leq \delta$
 Get search direction $\Delta\theta^* \leftarrow \theta_{j+1}^* - \theta_j$
 for $k = 0, 1, 2, \dots$ **do** ▷ Line search
 $\theta' \leftarrow \theta_j + \xi^k \Delta\theta^*$ ▷ $\xi \in (0, 1)$ is the backtracking coefficient
 if $\mathbb{E}_{s \sim \pi_j} [\mathcal{D}_{KL}(\pi_{\theta'} \|\pi_j)[s]] \leq \delta$ **and** ▷ Trust region
 $O_{\pi_{\theta'},\pi_j} \geq O_{\pi_j,\pi_j}$ **then** ▷ Objective
 $\theta_{j+1} \leftarrow \theta'$ ▷ Update policy
 break
 end if
 end for
end for

C EXPERIMENT DETAILS

C.1 GUARD ENVIRONMENT SETTINGS

Goal Task In the Goal task environments, the reward function is:

$$r(x_t) = d_{t-1}^g - d_t^g + \mathbf{1}[d_t^g < R^g],$$

where d_t^g is the distance from the robot to its closest goal and R^g is the size (radius) of the goal. When a goal is achieved, the goal location is randomly reset to someplace new while keeping the rest of the layout the same.

Push Task In the Push task environments, the reward function is

$$r(x_t) = d_{t-1}^r - d_t^r + d_{t-1}^b - d_t^b + \mathbf{1}[d_t^g < R^g],$$

where d^r and d^b are the distance from the robot to its closest goal and the distance from the box to its closest goal, and R^g is the size (radius) of the goal. The box size is 0.2 for all the Push task environments. Like the goal task, a new goal location is drawn each time a goal is achieved.

Chase Task In the Chase task environments, the reward function is

$$r(x_t) = d_{t-1}^r - d_t^r + \mathbf{1}[d_t^g < R^g],$$

where d^r is the distance from the robot to its closest goal and R^g is the size (radius) of the goal. Those targets continuously move away from the robot at a slow speed. The dense reward component provides a bonus for minimizing the distance between the robot and the targets. The targets are constrained to a circular area.

The test suites of our continuous experiments are summarized in Table 3.

Table 3: The test suites environments of our continuous experiments

Task Settings		Moving Area			Task			Dimension	
		Ground	Aerial	Fixed	Goal	Push	Chase	Low	High
GUARD Robot	Arm3 (\mathbb{R}^3)			✓	✓			✓	
	Drone (\mathbb{R}^4)		✓		✓			✓	
	Point (\mathbb{R}^2)	✓			✓	✓		✓	
	Swimmer (\mathbb{R}^2)	✓			✓	✓		✓	
	Hopper (\mathbb{R}^5)	✓			✓	✓		✓	
	Ant (\mathbb{R}^8)	✓				✓	✓		✓
	Walker (\mathbb{R}^{10})	✓				✓	✓		✓
Mujoco	Humanoid (\mathbb{R}^{17})	✓							✓
Robot	HumanoidStandup (\mathbb{R}^{17})	✓							✓

State Space The internal state spaces describe the state of the robots, which can be obtained from standard robot sensors (accelerometer, gyroscope, magnetometer, velocimeter, joint position sensor, joint velocity sensor and touch sensor). The details of the internal state spaces of the robots in our test suites are summarized in Table 4.

Control Space For all the experiments, the control space of all robots are continuous, and linearly scaled to $[-1, +1]$.

C.2 POLICY SETTINGS

The hyper-parameters used in our experiments are listed in Table 5 as default.

Table 4: The internal state space components of different test suites environments.

Internal State Space	Point	Swimmer	Walker	Ant	Drone	Hopper	Arm3
Accelerometer (\mathbb{R}^3)	✓	✓	✓	✓	✓	✓	✓
Gyroscope (\mathbb{R}^3)	✓	✓	✓	✓	✓	✓	✓
Magnetometer (\mathbb{R}^3)	✓	✓	✓	✓	✓	✓	✓
Velocimeter (\mathbb{R}^3)	✓	✓	✓	✓	✓	✓	✓
Joint position sensor (\mathbb{R}^n)	$n = 0$	$n = 2$	$n = 10$	$n = 8$	$n = 0$	$n = 6$	$n = 3$
Joint velocity sensor (\mathbb{R}^n)	$n = 0$	$n = 2$	$n = 10$	$n = 8$	$n = 0$	$n = 6$	$n = 3$
Touch sensor (\mathbb{R}^n)	$n = 0$	$n = 4$	$n = 2$	$n = 8$	$n = 0$	$n = 1$	$n = 1$

Our experiments use separate multi-layer perception with *tanh* activations for the policy network and value network. Each network consists of two hidden layers of size (64,64). Policy networks and value networks are trained using *Adam* optimizer. Policy networks are trained with learning rate of 1e-3 while value networks are trained with 3e-4.

We apply an on-policy framework in our experiments. During each epoch the agent interact B times with the environment and then perform a policy update based on the experience collected from the current epoch. The maximum length of the trajectory is set to 1000. The steps in each epoch is set to 30000. The total epoch number N is set to 200 in continuous tasks and 500 in atari tasks as default.

The policy update step is based on the scheme of TRPO, which performs up to 100 steps of back-tracking with a coefficient of 0.8 for line searching.

For all experiments, we use a discount factor of $\gamma = 0.99$, an advantage discount factor $\lambda = 0.97$, and a KL-divergence step size of $\delta_{KL} = 0.02$.

Other unique hyper-parameters for each algorithms follow the original paper to attain best performance.

Each model is trained on a server with a 48-core Intel(R) Xeon(R) Silver 4214 CPU @ 2.2.GHz, Nvidia RTX A4000 GPU with 16GB memory, and Ubuntu 20.04.

C.3 NORMALIZED SCORE SETTINGS

We intend to use 100% as the TRPO baseline measure. Thus the score we used is slightly changed from the normalization algorithm proposed by van Hasselt et al. (2015) with the form as follow:

$$\begin{aligned} \Delta_1 &\doteq score_{agent} - score_{random} \\ \Delta_2 &\doteq score_{TRPO} - score_{random} \\ score_{normalized} &= \frac{\Delta_2}{\Delta_1} \text{ if } \Delta_1 < 0 \text{ and } \Delta_2 < 0 \text{ else } \frac{\Delta_1}{\Delta_2} \end{aligned}$$

The difference is that we take the inverse of the metrics used in van Hasselt et al. (2015) when $\Delta_1 < 0$ and $\Delta_2 < 0$. For further explanation, we need to address the positive and negative cases of Δ_1 and Δ_2 :

- $\Delta_1 > 0$ and $\Delta_2 > 0$ In this case, $\frac{\Delta_1}{\Delta_2}$ can effectively demonstrate the ability of algorithms.
- $\Delta_1 < 0$ and $\Delta_2 < 0$ Practically in this case, assuming $score_{agent} > score_{TRPO}$, $score_{normalized}$ should be greater than 1 to demonstrate that agent is more capable than baseline TRPO. However, we will get a decimal if we obey the original score algorithms in van Hasselt et al. (2015) which is incorrect. Thus we take the inverse of it in this situation.
- $\Delta_1 < 0$ and $\Delta_2 > 0$ In this case we will get a negative number which is reasonable to show the negative effects in terms of reward enhancement.
- $\Delta_1 > 0$ and $\Delta_2 < 0$ The changed normalization algorithm is still incorrect in this situation. However, we have not encountered such cases in all of our atari game statistics.

Table 5: Important hyper-parameters of different algorithms in our experiments

Policy Parameter		A2C	TRPO	PPO	APO
Epochs in continuous tasks	N_1	200	200	200	200
Epochs in discrete tasks	N_2	500	500	500	500
Steps per epoch	B	30000	30000	30000	30000
Maximum length of trajectory	L	1000	1000	1000	1000
Policy network hidden layers		(64, 64)	(64, 64)	(64, 64)	(64, 64)
Policy network iteration		80	80	80	80
Policy network optimizer		Adam	Adam	Adam	Adam
Policy learning rate		3e-4	3e-4	3e-4	3e-4
Discount factor	γ	0.99	0.99	0.99	0.99
Advantage discount factor	λ	0.97	0.97	0.97	0.97
backtracking steps		-	100	-	100
backtracking coefficient		-	0.8	-	0.8
Target KL	δ_{KL}	-	0.02	0.02	0.02
Clip Ratio	ϵ	-	-	0.2	-
Probability factor	k	-	-	-	10
Weight of $\mathcal{J}(\pi)$	w_1	-	-	-	1/3
Weight of $\mathcal{B}_k(\pi)$	w_2	-	-	-	2/3
Value network hidden layers		(64, 64)	(64, 64)	(64, 64)	(64, 64)
Value network iteration		80	80	80	80
Value network optimizer		Adam	Adam	Adam	Adam
Value learning rate		1e-3	1e-3	1e-3	1e-3

D TOTAL EXPERIMENT RESULTS

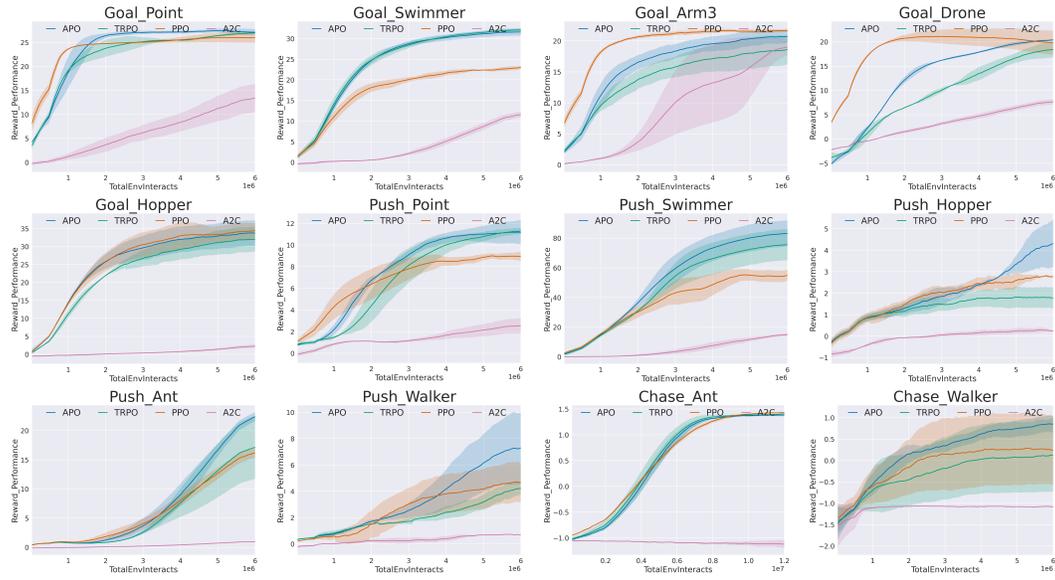


Figure 10: All continuous GUARD tasks learning curves.

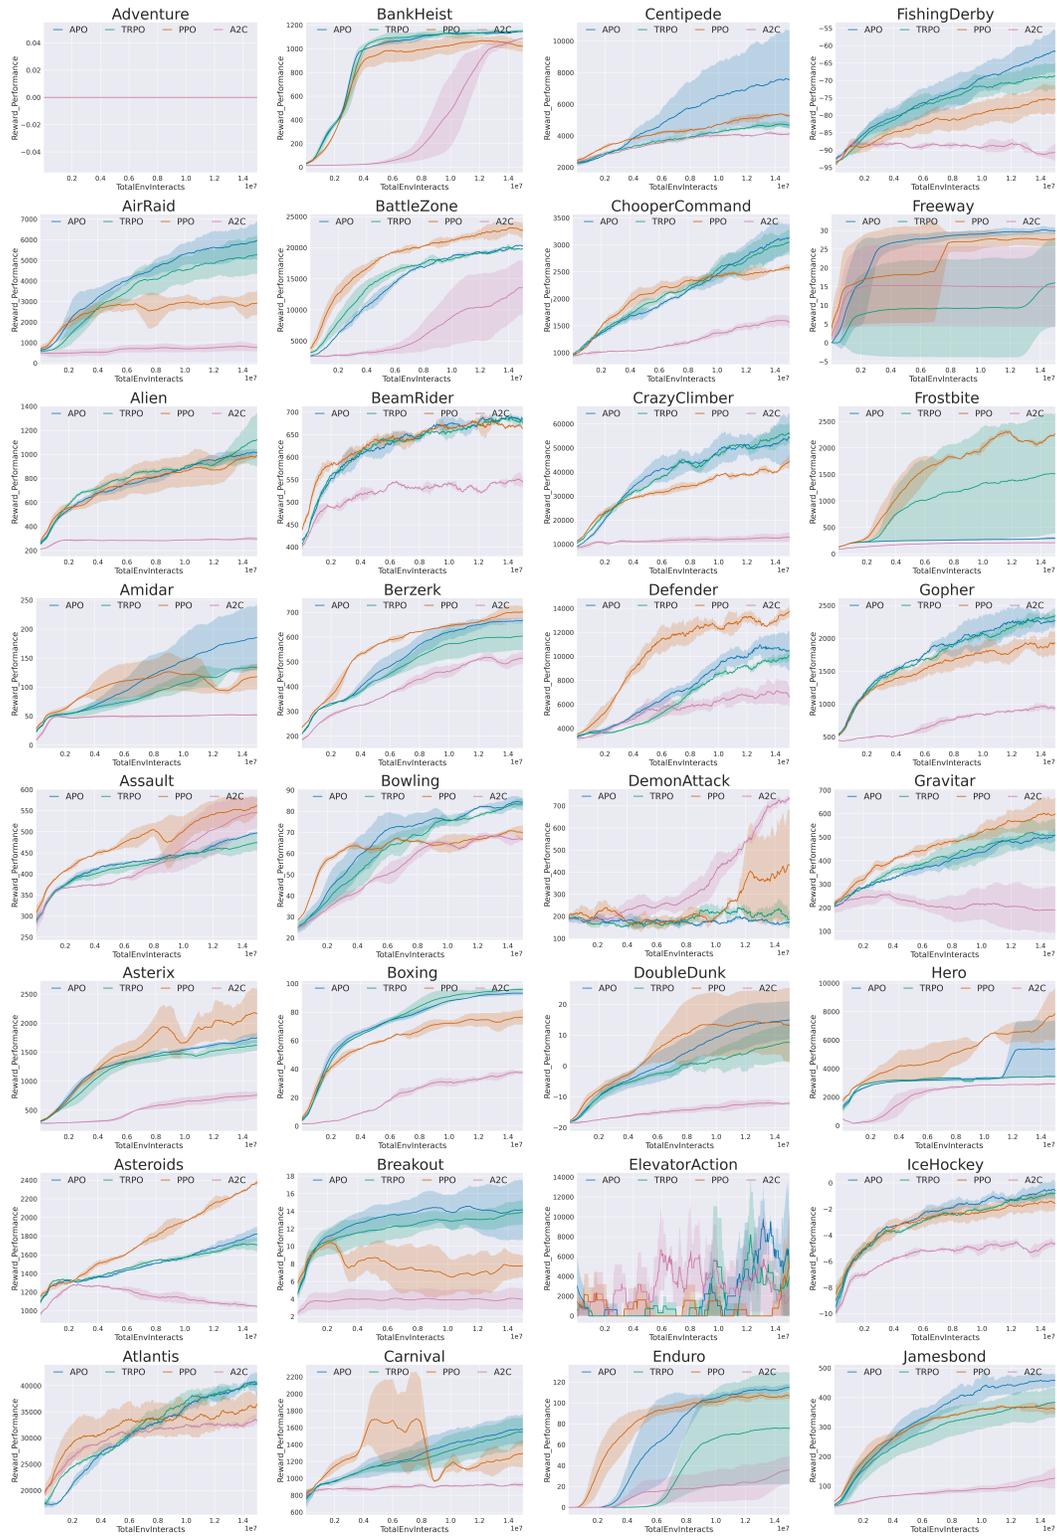


Figure 11: Comparison of results on Atari Game No.1 - No. 32

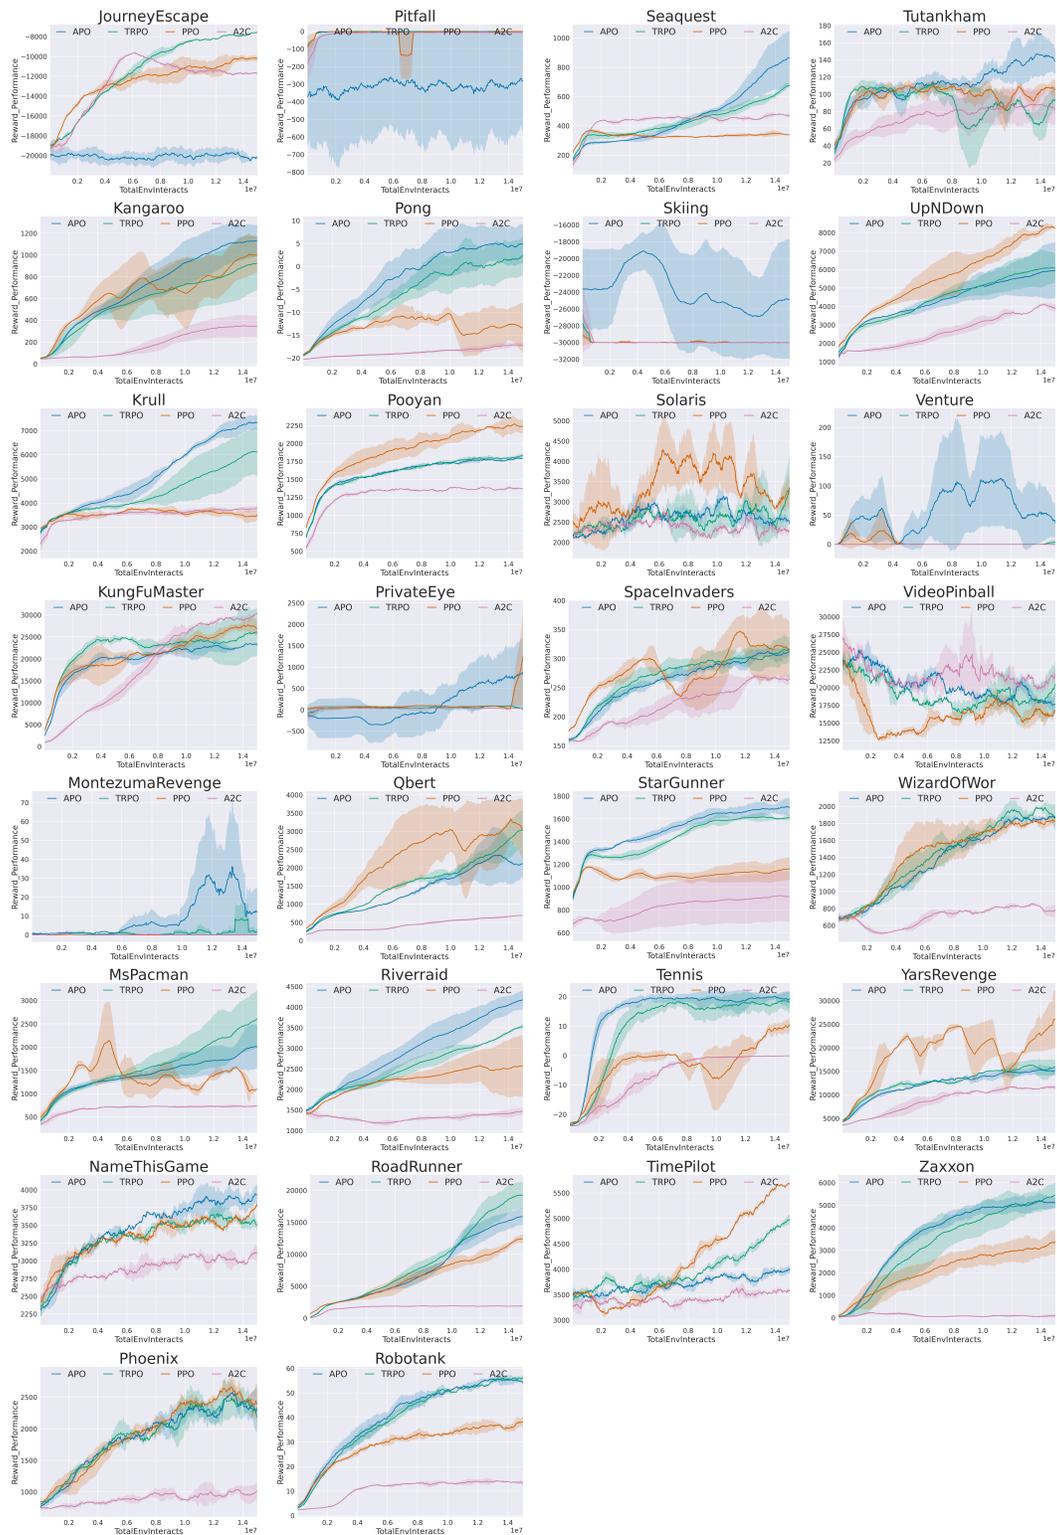


Figure 12: Comparison of results on Atari Game No.33 - No. 62