

Stabilizing Policy Gradient Methods via Reward Profiling

Shihab Ahmed¹, El Houcine Bergou⁴, Yue Wang^{1,2}, Aritra Dutta^{2,3}

¹Department of Electrical and Computer Engineering, University of Central Florida, FL, USA

²Department of Computer Science, University of Central Florida, FL, USA

³Department of Mathematics, University of Central Florida, FL, USA

⁴College of Computing, Mohammed VI Polytechnic University, Ben Guerir, Morocco
{Shihab.Ahmed, Yue.Wang, Aritra.Dutta}@ucf.edu, elhoucine.bergou@um6p.ma

Abstract

Policy gradient methods, which have been extensively studied in the last decade, offer an effective and efficient framework for reinforcement learning problems. However, their performances can often be unsatisfactory, suffering from unreliable reward improvements and slow convergence, due to high variance in gradient estimations. In this paper, we propose a universal reward profiling framework that can be seamlessly integrated with any policy gradient algorithm, where we selectively update the policy based on *high-confidence performance estimations*. We theoretically justify that our technique will not slow down the convergence of the baseline policy gradient methods, but with high probability, will result in stable and monotonic improvements of their performance. Empirically, on eight continuous-control benchmarks (Box2D and MuJoCo/PyBullet), our profiling yields up to $1.5\times$ faster convergence to near-optimal returns, up to $1.75\times$ reduction in return variance on some setups. Our profiling approach offers a general, theoretically grounded path to more reliable and efficient policy learning in complex environments.

Code — <https://github.com/shlhab/reward-profiling>

Extended version — <https://arxiv.org/abs/2511.16629>

1 Introduction

Reinforcement learning (RL) optimizes an agent’s performance in a stochastic, sequential decision-making task. Policy-gradient (PG) methods, which directly optimize parameterized policies from sampled trajectories rather than relying on value-function bootstrapping, form one of the core paradigms in RL (Sutton, Barto et al. 1998; Schulman et al. 2015). Its direct formulation hence enables PG’s effective implementations in high-dimensional and continuous-control settings, including robotic manipulation (Peters and Schaal 2006), locomotion (Todorov, Erez, and Tassa 2012), and autonomous driving (Schulman et al. 2017), where value-based approaches are inefficient.

However, PG methods generally suffer from *high variance* in gradient estimations. Small stochastic fluctuations in early-episode rewards of the trajectories can propagate through the Monte Carlo estimator (e.g. REINFORCE (Williams 1992)), leading to erratic updates, slow convergence, and occasional collapse of performance (Ilyas et al. 2018; Lehmann 2024).

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

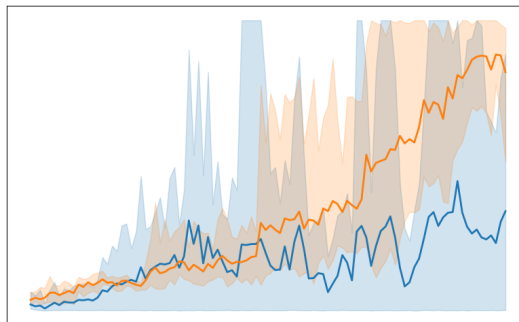


Figure 1: Training performance of REINFORCE. The agent converges with stable behavior with our reward profiling. — REINFORCE — REINFORCE+Lookback.

Classic variance-reduction techniques, including baseline/advantage normalization and trust-region constraints, can improve the stability but also bring drawbacks: problem-specific tuning (Chung et al. 2021), second-order solvers (Schulman et al. 2015), or extra computational overhead (Wu et al. 2017). A unified framework combining sample efficiency, stability, and simplicity that does not rely on problem-specific baselines can be useful. Therefore, a natural question arises:

Is it possible to reduce PG variance and stabilize learning across arbitrary policy-gradient methods, without relying on specialized tuning or heavy second-order machinery?

In this paper, we provide an affirmative answer to this question by proposing our reward profiling framework. Our framework is based on a straightforward, natural idea: we only update the policy if the updated one implies a better performance. Under ideal circumstances, when our justifications are exact, our reward profiling will imply a monotonically increasing performance, addressing the issue of unstable performance in vanilla PG methods. We perform a simple experiment to entertain this idea. We apply our reward profiling to a simple REINFORCE algorithm, and implement two algorithms under the CartPole environment (Brockman et al. 2016). Figure 1 shows vanilla REINFORCE can suffer from unstable learning and severe performance fluctuations, whereas the reward profiling provides a much more stable, nearly monotonic improvement. This observation shows that the profiling framework can address the fundamental challenge in PG without acquiring any complicated techniques.

In this paper, we further investigate this idea and develop

our reward profiling framework to overcome the common instability issues of PG methods. Our contributions are summarized as follows.

Design of a universal reward-profiling framework (§4). We first develop our universal reward-profiling framework in Algorithm 1, where we selectively accept, reject, or blend the policy updates based on the high-confidence value estimations of their value functions. The framework requires a small number of additional rollouts per iteration, without incurring significant additional computation. First, we introduce our `Lookback` technique, which updates the policy only when the estimated cumulative reward of the new policy surpasses that of the current policy. We further develop two variants of our profiling technique (`Mix-up` and `Three-Points`) to address the potential issue of being stuck at a local optimum. They are designed to accelerate convergence while maintaining stability. This can be applied to any PG methods, and is expected to stabilize the learning process and improve their performance.

Theoretical guarantees (§5). We establish theoretical guarantees on the convergence and global optimality of the presented framework. We showed that, in the justification phase, with additional samples of order $\mathcal{O}(\epsilon^{-2} \ln(T/\delta))$, we can compare the performances of the current and updated policies with high probability, thus ensuring monotonic improvements. We also show that the profiling framework will not slow down convergence theoretically, and enjoys an $\mathcal{O}(T^{-1/4})$ sub-optimality gap on the last iterate, providing a solid foundation for our proposed framework.

Extensive empirical evaluation (§6). We adapt our framework with three representative algorithms: DDPG (Lillicrap et al. 2015), TRPO (Schulman et al. 2015), and PPO (Schulman et al. 2017), and design corresponding profiling algorithms. We then evaluate their performance on continuous-control benchmarks. Results show consistent gains in final performance, sample efficiency, and training stability, providing validation for the algorithmic framework. We also port the algorithm to a Unity-ML “Reacher” DDPG agent with a separate simulation backend to verify that our framework is broadly applicable while yielding faster convergence and lower variance.

2 Related Work

There are two lines of research aiming to tame high variance and improve stability in PG methods: (i) *algorithm-centric* variants that bake variance-reduction or trust-region ideas into the core update, and (ii) *wrapper-style* frameworks that layer on generic stability checks without re-engineering the underlying optimizer.

(i) **Algorithm-centric approaches.** Early variance reduction in REINFORCE introduced *baselines*, which are often a learned value function, to center Monte-Carlo returns, yielding unbiased and lower-variance updates (Sutton et al. 1999). Actor-critic architectures extend this idea by bootstrapping via temporal-difference learning, trading bias for further variance reduction (Konda and Tsitsiklis 1999; Sutton 1984); however, a misspecified critic step can introduce harmful bias (Chen, Sun, and Yin 2021; Olshevsky and Ghahesifard 2023). As an alternative approach, trust-region methods such as TRPO enforce a KL-constraint to guarantee monotonic policy improvement under certain regularity conditions (Schul-

man et al. 2015), but rely on second-order solvers, which are expensive. PPO replaces TRPO’s conjugate-gradient and Hessian computations with a clipped surrogate objective, retaining many of the stability benefits while using only first-order updates (Schulman et al. 2017); nevertheless, it can still exhibit overly conservative steps or exploration failures in complex, contact-rich tasks (Wang et al. 2019b). Beyond on-policy methods, off-policy actor-critic algorithms like DDPG (Lillicrap et al. 2015) and TD3 (Fujimoto, van Hoof, and Meger 2018) aim for high sample efficiency via replay buffers, but suffer from spiky Q-value estimates and divergent updates without careful regularization (Haarnoja et al. 2018; Islam et al. 2017). Variants such as SAC inject entropy bonuses for exploration and stability (Haarnoja et al. 2018), yet their gains come at the cost of extra hyperparameters and temperature tuning.

(ii) **Framework-based approaches.** An orthogonal strand of work wraps a generic PG optimizer with lightweight checks or corrections. SVRG-style control variates (e.g., SVRPG) freeze a reference policy to reduce gradient variance, at the expense of extra memory and on-policy rollouts (Xu, Liu, and Peng 2018). Momentum injections, such as Nesterov and heavy-ball variants, have been proposed to accelerate PG under smoothness assumptions (Xiao 2022; Chen et al. 2024); although they sometimes amplify early spikes. In acceptance-based schemes, updates are only applied if a held-out performance estimate improves—trace back to optimistic or hysteretic updates in multi-agent settings (Omidshafiei et al. 2017; Palma et al. 2018); however, repeatedly validating a “true” return can be prohibitively costly.

Our reward-profiling wrapper unifies and extends these ideas with three simple, hyperparameter-minimal schemes: `Lookback` acts like a backtracking line search in policy space, rejecting any update whose empirical returns are lower than the previous policy. `Mixup` is a convex combination of the old and new parameters to smooth the transition and escape rejection deadlocks. `Three-Points` evaluates an additional “midpoint” to choose the best of the old, new, and mixed policies. Requiring only $\mathcal{O}(\epsilon^{-2} \ln(T/\delta))$ extra rollouts, it can be plugged into any PG method (DDPG, TRPO, PPO, etc.) without per-environment tuning. This design delivers theoretical improvement guarantees and consistent empirical gains in final performance, sample efficiency, and training stability. The framework can be related to line-search and momentum (Moré, Garbow, and Hillstom 1994; Sutskever et al. 2013; Muehlebach and Jordan 2021), but is designed as a virtually zero-hyperparameter, plug-in wrapper.

3 Preliminaries

The foundational framework for Reinforcement Learning is the Markov Decision Process (MDP), roughly $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, r, \gamma, \rho)$, where \mathcal{S} is the state space, \mathcal{A} the action set, $P(s' | s, a)$ the transition kernel, $r(s, a) \in [0, R_{\max}]$ the bounded one-step reward, $\gamma \in [0, 1)$ the discount factor, and ρ the initial state distribution. At each step, the agent observes the current state s_t and takes an action a_t . The environment then transits the next state s_{t+1} according to the transition kernel $P(\cdot | s_t, a_t)$, and receiving an immediate reward $r(s_t, a_t)$. A parameterized Markov policy is a mapping $\pi_\theta : \mathcal{S} \rightarrow \Delta(\mathcal{A})$ ¹ that captures the probability of taking

¹ $\Delta(\cdot)$ is the probability simplex over the space \cdot .

actions under each state, for some parameter $\theta \in \Theta$. A policy π_θ can randomly induce a trajectory $\tau = (s_0, a_0, s_1, \dots)$ under the MDP, whose return is defined as the accumulated reward along the trajectory: $G(\tau) = \sum_{t=0}^{\infty} \gamma^t r(s_t, a_t)$. The value function of a policy π_θ is then defined as the expectation of returns: $V^\pi(s) \triangleq \mathbb{E}[G(\tau) \mid s_0 = s]$. The goal is to find a policy π , following which the agent can get the highest cumulative reward:

$$\theta^* = \arg \max_{\theta \in \Theta} J(\theta), \text{ where } J(\theta) \triangleq \mathbb{E}[V^{\pi_\theta}(s) \mid s \sim \rho], \quad (1)$$

for some initial distribution ρ .

Policy gradient algorithms optimize $J(\theta)$ through gradient ascent, based on the Policy Gradient Theorem (Sutton et al. 1999): $\nabla_\theta J(\theta) = \mathbb{E}_{\tau \sim \pi_\theta} \left[\sum_{t=0}^{\infty} \nabla_\theta \log \pi_\theta(a_t \mid s_t) Q^{\pi_\theta}(s_t, a_t) \right]$, where the action-value function is defined as $Q^\pi(s, a) \triangleq \mathbb{E}[G(\tau) \mid s_0 = s, a_0 = a]$.

In practice, as the value functions are unknown, one needs to replace them with the (Monte-Carlo) estimation, resulting in the straightforward REINFORCE algorithm (Sutton, Barto et al. 1998). However, REINFORCE suffers from variances in gradient estimation and unstable performance; hence, modern approaches employ distinct stabilization mechanisms through constrained optimization and off-policy learning. Among on-policy methods, TRPO maximizes a surrogate-advantage objective subject to a hard KL-constraint (Schulman et al. 2015):

$$\begin{aligned} \max_{\theta} \quad & \mathcal{L}^{\text{TRPO}}(\theta) = \mathbb{E}_{\tau \sim \pi_{\theta_{\text{old}}}} \left[r_t(\theta) \hat{A}^{\pi_{\theta_{\text{old}}}}(s_t, a_t) \right], \\ \text{s.t.} \quad & \mathbb{E}_{\tau \sim \pi_{\theta_{\text{old}}}} \left[\mathbb{D}_{KL}(\pi_{\theta_{\text{old}}}(\cdot \mid s_t) \parallel \pi_\theta(\cdot \mid s_t)) \right] \leq \delta, \end{aligned}$$

where $r_t(\theta) = \frac{\pi_\theta(a_t \mid s_t)}{\pi_{\theta_{\text{old}}}(a_t \mid s_t)}$, and \hat{A} being an advantage estimator. TRPO enjoys a theoretical guarantee of (approximate) monotonic policy improvement with the KL constraint. PPO replaces TRPO’s hard constraint with a clipped surrogate that is cheaper to compute, still discourages large updates (Schulman et al. 2017). Unlike TRPO’s KL-based trust region, PPO does not admit the same theoretical assurance of monotonic improvement, yet a practical heuristic that works well in large-scale implementations. Meanwhile, DDPG has remained a standard off-policy PG algorithm, training a deterministic actor $\mu_\theta : \mathcal{S} \rightarrow \mathcal{A}$ and a critic Q_ϕ off-policy using a replay buffer \mathcal{D} and slowly-updated target networks (Lillicrap et al. 2015) ($\mu_{\theta'}, Q_{\phi'}$):

$$\begin{aligned} \mathcal{L}(\phi) &= \mathbb{E}_{(s, a, r, s') \sim \mathcal{D}} \left[r + \gamma Q_{\phi'}(s', \mu_{\theta'}(s')) - Q_\phi(s, a) \right]^2, \\ \nabla_\theta J &\approx \mathbb{E}_{s \sim \mathcal{D}} \left[\nabla_a Q_\phi(s, a) \Big|_{a=\mu_\theta(s)} \nabla_\theta \mu_\theta(s) \right]. \end{aligned}$$

Twin Delayed DDPG (TD3) (Fujimoto, van Hoof, and Meger 2018) further stabilizes DDPG by (i) using two critics and taking the minimum to reduce overestimation bias, (ii) delaying policy and target network updates, and (iii) adding clipped noise to target actions for policy smoothing.

4 Reward Profiling Framework

In its simplicity, our reward profiling framework compares the performance of two policies $\pi_{\theta_1}, \pi_{\theta_2}$. Within any PG method step, a potential update on the policy is made:

$\theta_1 \rightarrow \theta_2$. This updated, as mentioned, is based on the inaccurate gradient estimation, which is the key source of high variance, leading to occasional updates that *decrease* the true performance $J(\pi)$, slow overall convergence, and provide no safeguard on the quality of the policy selection. Under the reward profiling framework, this update contributes only if it improves the performance. In the ideal case, we exactly know the value functions of $J(\theta_1)$ and $J(\theta_2)$, then we will only accept the update $\theta_1 \rightarrow \theta_2$ if $J(\theta_2) \geq J(\theta_1)$, which results in a monotonically increasing performance and stabilizes training. In practice, however, we do not know $J(\theta)$ and need to make the comparison based on estimation. It is clear that the performance and stability highly depend on the accuracy of the comparison. To handle this, we provide our **high-confidence comparison** scheme.

In general, for two candidate parameterized policies π_j with $j \in \{1, 2\}$, a number (E) of *i.i.d.* trajectories are sampled from each policy forming evaluation sample sets $\mathcal{D}_j = \{\tau_i^{(j)}\}_{i=1}^E \sim \pi_j$. The empirical return estimate for π_j is computed as

$$\hat{J}(\pi_j) = \frac{1}{|\mathcal{D}_j|} \sum_{\tau \in \mathcal{D}_j} G(\tau), \quad (2)$$

where $G(\tau)$ denotes the cumulative (discounted) return of trajectory τ . Based on the comparison of these return estimates, the corresponding parameters are selected for the policy. We refer to this scheme as the “lookback” framework, where we compare the performance of the policies from two successive steps:

$$\theta_{t+1} = \begin{cases} \theta_{\text{new}}, & \hat{J}(\pi_{\text{new}}) \geq \hat{J}(\pi_{\text{old}}), \\ \theta_{\text{old}}, & \text{otherwise.} \end{cases} \quad (3)$$

If the noise masks genuine improvements in the noisy estimates, the agent might get stuck. To address this, we draw inspiration from the method *mixup*, a classic data augmentation technique used in supervised learning (Zhang et al. 2017), where inputs and labels are convexly combined to reduce inductive bias, widely used in computer vision applications (Dutta et al. 2024). Specifically, we consider an intermediate policy defined by

$$\theta_{\text{mix}} = \lambda \theta_{\text{new}} + (1 - \lambda) \theta_{\text{old}}, \quad (4)$$

where $\lambda \in [0, 1]$ is a mixing parameter. Similarly, we compare their performance and accept the update if it results from a higher reward:

$$\theta_{t+1} = \arg \max_{\theta \in \{\theta_t, \theta_{\text{mix}}\}} \hat{J}(\pi_\theta). \quad (5)$$

We also highlight that, as θ_{mix} lies closer to θ_t in parameter space (when λ is small), this step functions as a cheap trust region, yet without any Hessian computation or KL constraint as in TRPO (Schulman et al. 2015, 2017). This step is expected to result in a better performance. Finally, the last modification unifies the best of the first two cases. The *Three-Points* variant considers all candidates $\{\theta_t, \theta', \theta_{\text{mix}}\}$ and selects the best:

$$\theta_{t+1} = \arg \max_{\theta \in \{\theta_{\text{old}}, \theta_{\text{new}}, \theta_{\text{mix}}\}} \hat{J}(\pi_\theta). \quad (6)$$

By design, this technique ensures that π mostly improves over the training process and provides enough room for exploration. We summarize the entire procedure in Algorithm 1, noting that it can be wrapped around any first-order PG method without altering its core update logic.

Algorithm 1: Reward Profiling Framework

Input: initial θ_0 ; iterations T ; rollouts E ; mix-weight λ ; variant $\in \{\text{LB, MU, TP}\}$
for $t = 0, \dots, T - 1$ **do**
 Update the policy following any PG method to the new parameter θ'
 $\theta_{\text{mix}} \leftarrow \lambda \theta' + (1 - \lambda) \theta_t$ (Eq. (4))
 Estimate $\hat{J}_{\text{old}}, \hat{J}_{\text{new}}, \hat{J}_{\text{mix}}$ with $\theta_{t+1} \leftarrow$
 $\begin{cases} \arg \max_{\{\theta_t, \theta'\}} \hat{J}, & \text{\textbackslash\textbackslash Lookback \textbackslash\textbackslash} \\ \arg \max_{\{\theta_t, \theta_{\text{mix}}\}} \hat{J}, & \text{\textbackslash\textbackslash Mix-up \textbackslash\textbackslash} \\ \arg \max_{\{\theta_t, \theta', \theta_{\text{mix}}\}} \hat{J}, & \text{\textbackslash\textbackslash Three-points \textbackslash\textbackslash} \end{cases}$
 (Eqs. (5),(6))
end

5 Convergence Analysis

To provide theoretical validation of the framework, we study the choice of E , which serves as a sort of evaluation budget, in our method to ensure accuracy, and then we propose a detailed convergence analysis. We start with adapting the following concentration inequality.

Lemma 1 (Concentration) *Let the policy π have per-step rewards lie in $[0, R_{\text{max}}]$, so that any trajectory return satisfies $0 \leq G(\tau) \leq B$ with $B = \frac{R_{\text{max}}}{1-\gamma}$. Then for any $\epsilon > 0$,*

$$\mathbb{P}(|\hat{J}(\pi) - J(\pi)| \geq \epsilon) \leq 2 \exp\left(-\frac{2E\epsilon^2}{B^2}\right).$$

Applying Lemma 1 with a per-step failure probability of $\frac{\delta}{T}$ and then union-bounding over T consecutive updates, we get for every $t \in [T]$ empirical estimates satisfies $|\hat{J}(\pi_t) - J(\pi_t)| \leq \epsilon$, provided $E \geq \frac{B^2}{2\epsilon^2} \ln\left(\frac{2T}{\delta}\right)$, with probability at least $1 - \delta$.²

Lemma 2 (Monotonicity) *Let $E \geq \frac{B^2}{2\epsilon^2} \ln\left(\frac{2T}{\delta}\right)$. Then with probability at least $1 - \delta$, for every update $t = 1, 2, \dots, T$, whenever the lookback rule accepts (i.e. $\hat{J}(\pi_{\theta_{t+1}}) \geq \hat{J}(\pi_{\theta_t})$), the true returns satisfy*

$$J(\pi_{\theta_{t+1}}) \geq J(\pi_{\theta_t}) - 2\epsilon.$$

This implies that, with this choice of E , the comparison ensures the accepted updates yield performance improvements, and hence the learning can be monotonic and more stable.

We then show that when instantiated with the *Lookback* decision rule and REINFORCE algorithm, Algorithm 1 converges to a near-optimal policy at rate $\mathcal{O}(T^{-1/4})$ with high probability. The analysis for the mix-up and three-point strategies follows similar lines. We make the following standard assumption.

Assumption 1 (Smoothness) *The policy class π_θ is smooth in θ , and there exists $\sigma \geq 0$ such that*

$$\|\mathbb{E}_{d^{\pi_\theta} \times \pi_\theta}[(\nabla \log \pi_\theta(A|S))(\nabla \log \pi_\theta(A|S))^T]\| \leq \sigma,$$

²This requirement is a worst-case bound on the return range. In our experiments, the variance observed is much lower with $E = 5 - 10$, only.

where $d^{\pi_\theta}(s)$ is the discounted-state-visitation distribution under π_θ .

Assumption 1 is standard in policy gradient studies, e.g., (Xu, Liu, and Peng 2018; Wang et al. 2024; Ganesh and Aggarwal 2024). There is a rich family of policies that satisfy the conditions, including the Softmax policy and policies defined through neural networks with smooth and Lipschitz activation functions (Tian, Olshevsky, and Paschalidis 2023). Finally, we are all set to quote the convergence of our algorithm.

Theorem 1 (Convergence) *Under Assumption 1, setting $\eta = \mathcal{O}(T^{-1/2})$ and $E = \frac{B^2}{2\epsilon^2} \ln\left(\frac{2T}{\delta}\right)$, results in $J(\pi^*) - J(\pi_{\theta_T}) = \mathcal{O}(T^{-1/4})$, with probability at least $1 - \delta$.*

The deterministic PG method (Agarwal et al. 2021; Xiao 2022) achieves a faster convergence rate than ours because it has access to the true policy gradient without error. In contrast, our algorithm achieves a similar convergence rate when the algorithm is stochastic and the update is based on policy gradient estimation, as in actor-critic algorithms (Xu, Liu, and Peng 2018; Suttle et al. 2023; Olshevsky and Ghahserifard 2023; Chen, Sun, and Yin 2021; Wang et al. 2024). This demonstrates that the convergence rate of our approach is not slowed down much by our profiling framework. More importantly, unlike previous works, Theorem 1 provides a high-probability guarantee on the *last* iteration, whereas prior results typically only guarantee the expectation or the best policy found during the learning process. This framework requires additional trajectories of order $\mathcal{O}(\epsilon^{-2} \ln(T/\delta))$, which is still acceptable when the training samples are redundant.

Remark 1 *The same convergence guarantee of Theorem 1 holds for the other two cases. We note that the Mixup is a special case of Lookback. Because $\theta_{\text{mix}} = \lambda \theta' + (1 - \lambda) \theta_t = \theta_t + \lambda \eta \hat{\nabla}_\theta J(\pi_{\theta_t})$. Three points strategy adds another point in Lookback and generates a better reward than Lookback in every iteration.*

Biased Case. In the above discussions, we consider the case where we compare performances based on Monte-Carlo estimation. In practice, however, it is often the case that PG methods are utilized in *actor-critic* algorithms, where the value functions are estimated based on the *critic*. Since the *critic* part is generally inaccurate, the resulting gradient estimations can have bias. Thus, it is important to extend our framework to the biased case. Given θ , we assume that the *critic* part provides a biased estimation of the value function \hat{Q}^{π_θ} , such that $\mathbb{E}[\hat{Q}^{\pi_\theta}] = Q^{\pi_\theta} + \epsilon_\theta$, where ϵ_θ is the bias introduced by the errors of the *critic* part. Such results can be obtained, for instance, when a deep neural network is used to estimate value functions, e.g., (Du et al. 2019; Neyshabur 2017; Miyato et al. 2018), and is widely adapted in actor-critic analysis, e.g., (Wang et al. 2019a; Zhou and Lu 2023; Chen et al. 2023; Zhang et al. 2020; Qiu et al. 2021; Kumar, Koppel, and Ribeiro 2023; Xu, Wang, and Liang 2020b,a; Suttle et al. 2023). It can be further shown that when the value function estimation is biased, the resulting gradient estimation has a bias of $CA\epsilon_\theta$, as long as $\|\nabla \pi_\theta\| \leq C$ (Sutton et al. 1999) and $A = |\mathcal{A}|$ actions. Therefore, without loss of generality, we assume the bias of zero-th and first order gradient of Q^{π_θ} are both bounded by some ϵ_θ . We then show the convergence of 1, when the Lookback technique is applied to REINFORCE.

Env.	Algo	Avg. Return (\pm Std)			Rounds to 0.95 \times Best Variability \downarrow (%)							
		Base	LB	MU	TP	Base	LB	MU	TP	LB	MU	TP
Bipedal	DDPG	-116.1 \pm 13.5	-100.0 \pm 6.9	-47.0 \pm 27.1	-58.9 \pm 34.4	-1.0	1.3	3.7	-39	-34	-64	
	PPO	-51.9 \pm 109.2	-3.1 \pm 66.1	-13.4 \pm 8.0	8.9 \pm 39.5	6.8	9.0	-	9.0	66	68	59
	TRPO	105.0 \pm 57.0	-35.3 \pm 11.1	-17.9 \pm 7.2	-27.9 \pm 16.1	7.7	-	-	-	67	82	56
CarRacing	DDPG	-84.7 \pm 4.2	-17.6 \pm 3.6	-17.6 \pm 3.6	-17.6 \pm 3.6	-1.6	3.6	3.0	-28	-23	-26	
	PPO	-63.2 \pm 14.3	21.7 \pm 23.0	26.9 \pm 21.8	30.7 \pm 39.5	2.3	3.0	4.6	2.0	-33	-48	-45
	TRPO	-72.8 \pm 23.7	-14.7 \pm 11.0	31.1 \pm 45.2	42.4 \pm 46.3	3.5	-	3.7	3.0	57	-22	-49
LunarLander	DDPG	-117.0 \pm 45.3	-29.9 \pm 35.8	26.3 \pm 98.4	37.1 \pm 78.4	-1.0	6.7	5.7	12	23	27	
	PPO	-6.9 \pm 132.3	-19.0 \pm 113.4	-71.9 \pm 87.3	-40.6 \pm 42.5	7.3	6.7	9.0	6.8	5	4	28
	TRPO	-15.0 \pm 89.0	21.9 \pm 80.5	79.0 \pm 93.2	94.1 \pm 112.8	5.5	7.0	8.7	6.3	41	48	27
Ant	DDPG	140.6 \pm 113.9	278.9 \pm 136.8	412.4 \pm 150.0	377.7 \pm 115.4	5.0	4.0	4.2	4.8	3	4	11
	PPO	822.6 \pm 35.8	842.4 \pm 42.4	768.3 \pm 97.3	863.7 \pm 67.4	5.0	7.2	9.0	6.8	39	-31	19
	TRPO	897.6 \pm 41.2	808.8 \pm 62.0	771.2 \pm 69.1	837.6 \pm 47.1	8.7	8.0	5.0	7.7	20	-67	24
HalfCheetah	DDPG	576.3 \pm 749.9	909.0 \pm 331.4	1137.8 \pm 185.6	931.0 \pm 345.4	10.0	7.0	8.6	8.0	43	52	26
	PPO	97.1 \pm 377.9	34.2 \pm 596.5	-1004.3 \pm 409.5	95.5 \pm 561.8	5.7	7.4	10.0	6.8	-16	48	-14
	TRPO	228.4 \pm 416.1	81.2 \pm 524.3	40.9 \pm 708.6	391.9 \pm 314.9	7.3	6.0	7.0	5.3	6	-4	17
Hopper	DDPG	1189.6 \pm 600.4	1048.7 \pm 555.5	1394.5 \pm 727.7	1492.2 \pm 562.9	7.8	4.5	5.8	6.3	-3	-16	-8
	PPO	825.5 \pm 71.4	802.8 \pm 49.6	338.9 \pm 330.7	634.3 \pm 281.6	8.4	9.0	9.0	10.0	-42	-25	-88
	TRPO	1196.3 \pm 464.7	856.5 \pm 526.0	646.5 \pm 283.6	724.4 \pm 219.2	8.2	9.0	-	-	-20	19	-8
Walker2D	DDPG	113.6 \pm 90.7	299.4 \pm 74.8	248.7 \pm 145.6	267.6 \pm 88.5	8.0	5.0	4.9	5.4	-41	-80	-44
	PPO	185.6 \pm 107.2	110.8 \pm 40.2	199.3 \pm 108.2	126.4 \pm 25.1	4.4	1.8	1.8	1.2	-89	-86	-63
	TRPO	260.2 \pm 207.4	192.6 \pm 133.2	141.6 \pm 54.4	231.4 \pm 182.4	3.0	2.2	2.4	2.4	-45	-19	-56
Humanoid	DDPG	-120.1 \pm 26.7	42.6 \pm 12.5	51.6 \pm 14.3	57.1 \pm 14.3	-	3.7	2.8	3.1	50	45	48
	PPO	58.0 \pm 7.1	48.2 \pm 6.9	47.9 \pm 9.1	48.3 \pm 7.5	8.1	1.9	2.2	1.1	3	7	4
	TRPO	68.8 \pm 6.4	44.8 \pm 6.2	47.8 \pm 6.9	47.0 \pm 7.2	8.0	2.0	3.0	3.0	-21	-14	-28

Table 1: Comparison across environments and algorithms. *Avg. Return* reports final mean \pm std over seeds; *Rounds to 0.95 \times Best* counts iterations to reach 95% of the best average return; *Variability \downarrow (%)* measures relative reduction in return variance. **Bold** indicates best performance; dash (-) indicates not reached within budget. Abbreviations: Base (vanilla baseline), LB (Lookback), MU (MixUp), TP (Three-Points).

Theorem 2 We set $\eta = \mathcal{O}(1/\sqrt{T})$ and run *I(Lookback)* for T steps. If the bias ϵ_θ satisfies $\epsilon_\theta \leq \mathcal{O}(1/T)$, then we have $\min_{t \leq T} \min\{\mathbb{E}\|\nabla V^{\pi_{\theta_t}}\|_2, \mathbb{E}\|\nabla V^{\pi_{\theta_t}}\|_2^2\} \leq \mathcal{O}\left(\frac{1}{\sqrt{T}}\right)$.

The result shows that our algorithm converges to a stationary point, as long as the value estimation is accurate enough. Moreover, due to the gradient dominance property of the value function (Agarwal et al. 2021), our result also implies the global optimality of the learned policy.

6 Experiments

6.1 Simulated Gymnasium Environments

We benchmarked our reward-profiling wrapper by performing extensive experiments to stress-test policy learning across the major flavors of policy-gradient methods. We wrapped three canonical algorithms. **TRPO**, the trusted “second-order” method with hard KL-constraints, **PPO**, its “first-order” surrogate-clipping successor and current on-policy workhorse, and **DDPG**, an off-policy, replay buffer actor-critic known for sample-efficient continuous control but notoriously unstable, in our profiling layer.

Profiling framework integration. We built a modular profiling wrapper atop Stable-Baselines3 (v1.9) and SB3-Contrib (Raffin et al. 2021), leveraging Gymnasium and PyBullet (Ellenberger 2018–2019; Towers et al. 2024) for simulation.

Each variant ran for 100k environment steps across five random seeds. Each profiling round consumes the same 10k environment steps as the baseline; samples used for improvement checks are *reused* for training updates. Thus, no additional interactions are collected. All variants share the same per-round evaluation schedule and count. Every candidate set is scored with short episodes ($E = 5$), and averages are computed on the identical rollouts for all baselines for fair comparison. The policy network in SB3 is actor-critic, which was used throughout the experiments.

Benchmark environments. We evaluate reward profiling on two complementary families of continuous control tasks: **Box2D Suite** where the BipedalWalker environment challenges agents to navigate uneven terrain using 24D lidar/joint observations and 4D torque actions, with sparse rewards for forward progress. CarRacing provides pixel-based control (96 \times 96 RGB input) for lap completion, while LunarLanderContinuous tests precise thruster control (8D state, 2D actions) under sparse landing rewards. **MuJoCo/PyBullet Suite**, where we experimented with high-dimensional locomotion with Ant, HalfCheetah, Hopper, Humanoid, Walker2D, with observation state and action both having several dimensions, ranging from torque control, actions for balance, or even for forward gait.

Observation. The most significant gain is in terms of variance performance overall as shown in Figure 2. In cases it

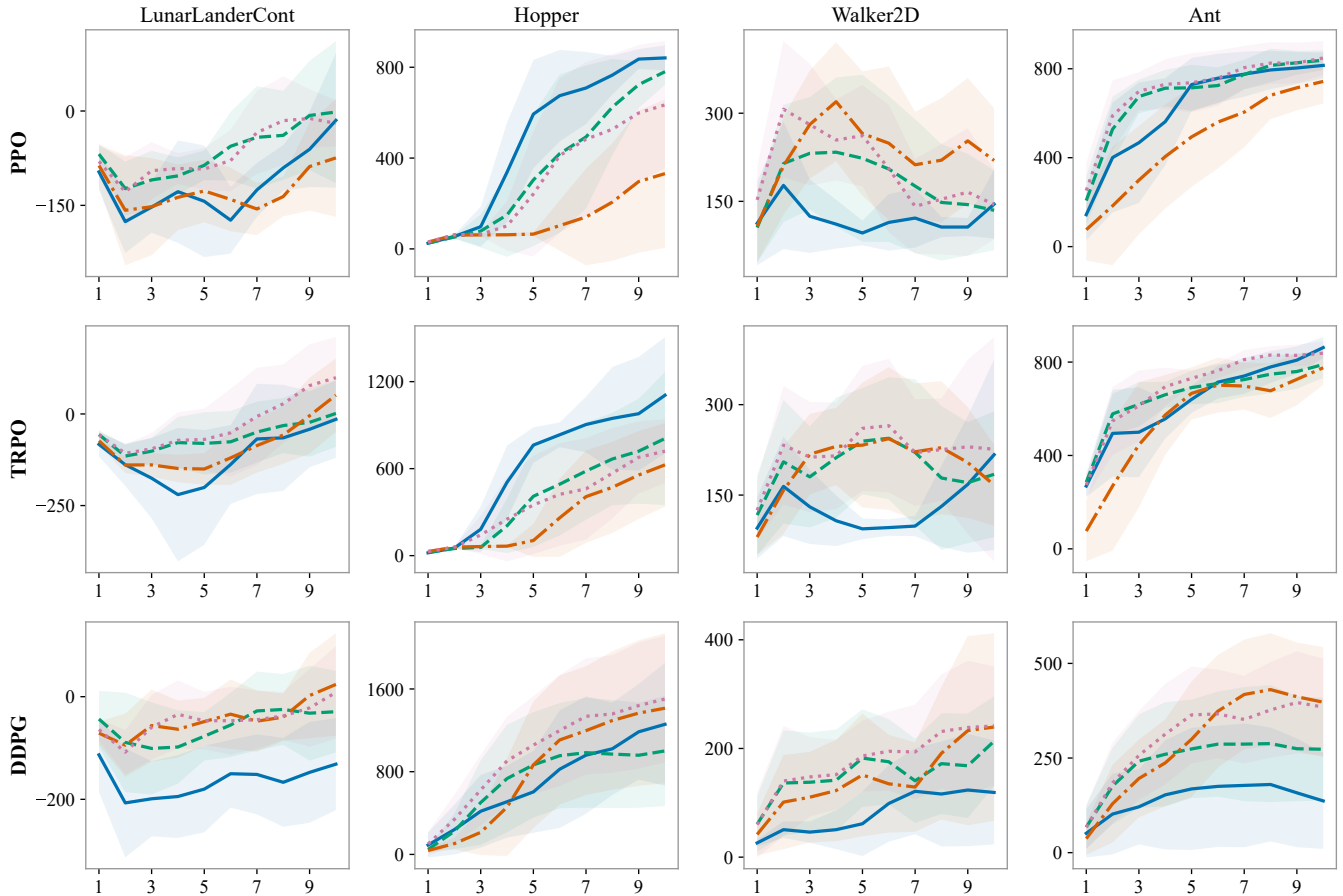


Figure 2: Average Return vs. timesteps ($\times 10^4$) for 3 policy-gradient methods, PPO (top row), TRPO (middle row), and DDPG (bottom row), across 4 continuous-control benchmarks: LunarLanderCont, Hopper, Walker2D, Ant under variants: — Vanilla, - - - Lookback, - - - Mixup, ··· Three-Points, with $E = 5$ for minimal overhead. Shaded regions indicate ± 1 standard deviation. Extra rollouts tuning yields convergence gains in complex environments.

transforms catastrophic failures into viable policies across key benchmarks. The CarRacing agent, both with TRPO and PPO, average large negative returns, crashing before completing a lap. The variants turn them into positive-return drivers. The training recovers stable driving behavior by modifying updates that worsen lap scores. DDPG remains unstable under an MLP policy, but profiling still raises the worst runs. For BipedalWalker, where vanilla TRPO already performs well (105.0 ± 57.0), profiling yields minimal gains due to inherent conservatism, though DDPG+Mixup still improves from -116.1 ± 13.5 to -47.0 ± 27.1 . It still provides competitive performance in precision tasks. Among our high-DOF benchmarks, Ant DDPG+Mixup gives modest bumps over the baseline with earlier convergence behavior, while its Three-Point variant provides superior variance characteristics (11% reduction from the baseline) with moderate final return over time. It also augments the training behavior with TRPO and PPO on several metrics, offering moderate improvement in variance. On HalfCheetah, DDPG+Mixup dominates (1137.8 ± 185.6 vs 576.3 ± 749.9) with usual variance performance in PPO. The performance on Walker2D and Humanoid remain challenging in variance stability; even

with profiling, TRPO and PPO show limited gains while baselines also show very similar results. The profiling framework improves most metrics with strong to moderate gains, except for the setups that remain challenging to learn with off-the-shelf algorithms without careful tuning.

6.2 Experiment with Multi-robot Task

For a realistic and domain-specific evaluation, we deployed it in a multi-agent robotic setting on the Unity ML-Agents Reacher task (Juliani et al. 2020; Cohen et al. 2022). Unity provides a flexible, visually realistic, and customizable simulation environment; see Figure 3. Here, the environment simulates 20 robotic arms operating in parallel, each receiving a 33-dimensional state vector (joint angles, velocities, and target position) and outputting 4D continuous torques. Agents earn a dense reward of $+0.1$ per timestep for maintaining their fingertip in a moving target zone. We plugged our profiling wrapper (the TP variant) into the platform atop a decentralized DDPG architecture, with policy updates evaluated every 50 episodes through 5-agent consensus voting.

Observation. Three-Point profiling (TP) demonstrates substantially more stable learning than vanilla DDPG, as evi-




Figure 3: Multi-agent Reacher task in UnityML: multiple arms coordinate to reach and manipulate randomly placed targets on a raised platform, testing control under interaction.

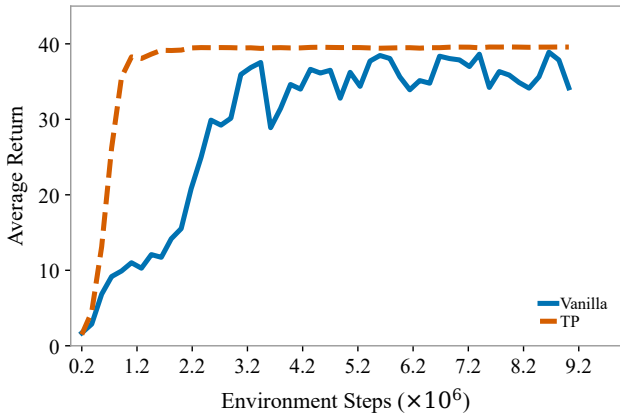


Figure 4: Performance in the UnityML Reacher environment. Profiling-enhanced DDPG converges faster and maintains greater stability than the vanilla counterpart.

denced by the training performance in Figure 4. TP-enhanced agent achieves a higher return from early training with fewer fluctuations, achieving smoother progression. Only five sampled returns ($E = 5$) make the computation overhead minimal while solving for the optimal policy at early training iterations. This implies the suitability of a reward-profiling framework for physical deployment, where stable, predictable learning behavior is a requirement.

7 Evaluation Rollout Sensitivity

Using an NVIDIA V100 GPU, baseline training of 100K steps took 10–30 min, with the TP variant adding only about 20% wall-clock time. Figure 5 shows how the evaluation budget shapes learning. At small E (10–20), the Monte Carlo estimates exhibit noisy, spurious updates and wide confidence bands. This results in large fluctuations and wide confidence bands. At the other extreme, very large E (e.g., 200), the algorithm advances only slowly once the noise floor is suppressed. Intermediate values ($E = 50$ –100) suppress the worst of the noise without over-constraining the update rule, yielding both rapid early gains and a much narrower variance envelope. So roughly, pushing E well beyond this critical scale offers diminishing returns. Once the estimation error




Figure 5: Sensitivity to evals E for DDPG+TP on Hopper. Smaller E (10–20) causes noisy, unstable updates; large E (200) improves stability but slows progress; mid-range E (50–100) balances both.

is negligible, further rollouts merely add computational cost and slow down true policy updates. This reflects the batch-size trade-off in SGD (McCandlish et al. 2018), where a task-dependent *critical batch size* marks the point beyond which additional samples no longer improve data efficiency. By analogy, one can think of a *critical evaluation budget* E_{crit} , and choose $E \approx E_{\text{crit}}$, to optimally balance variance reduction against responsiveness. Our framework complements other recent monotonic improvement goals (Xie et al. 2025) without relying on stochastic density ratios, hence applies uniformly to both stochastic and deterministic settings.

8 Conclusion and Open Challenges

We introduced *Reward Profiling*, a general-purpose wrapper applicable to policy gradient-based methods to enforce monotonic improvements. By “looking back” at the pre-update performance after gradient backpropagation and conditionally accepting, rejecting, or blending candidate updates, catastrophic collapses are reduced through stable learning. Empirically, across both on-policy (PPO, TRPO) and off-policy (DDPG) algorithms with actor-critic frameworks, and across dense-reward control tasks and high-DOF locomotion benchmarks, Reward Profiling mostly matched or improved returns along with the variance, sped up convergence relative to vanilla baselines. A key trade-off governed by the number of evaluation rollouts E for practical implementation: moderate values of E tend to strike the best balance between stability and responsiveness. Watching DDPG’s triumph with it, the integration of reward profiling into the Reacher task demonstrates its effectiveness in a realistic continuous-control setting. Profiling yielded markedly smoother learning curves under DDPG, motivating its use in off-policy, replay-buffer cases. While Reward Profiling incurs only $\mathcal{O}(\log T)$ rollouts over T iterations, the cost may be nontrivial in expensive simulators. On sparse-reward and large discrete-action tasks (e.g., Atari), this could be tested, adjusting E dynamically based on some uncertainty measure, or incorporating more selective updates for the most informative candidates.

Acknowledgments

The work of S.A and Y.W is partially supported by DARPA under Agreement No. HR0011-24-9-0427.

References

- Agarwal, A.; Kakade, S. M.; Lee, J. D.; and Mahajan, G. 2021. On the theory of policy gradient methods: Optimality, approximation, and distribution shift. *Journal of Machine Learning Research*, 22(98): 1–76.
- Brockman, G.; Cheung, V.; Pettersson, L.; Schneider, J.; Schulman, J.; Tang, J.; and Zaremba, W. 2016. OpenAI Gym. *arXiv preprint arXiv:1606.01540*.
- Chen, T.; Sun, Y.; and Yin, W. 2021. Closing the Gap: Tighter Analysis of Alternating Stochastic Gradient Methods for Bilevel Problems. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems*, volume 34, 25294–25307. Curran Associates, Inc.
- Chen, X.; Duan, J.; Liang, Y.; and Zhao, L. 2023. Global convergence of two-timescale actor-critic for solving linear quadratic regulator. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 7087–7095.
- Chen, Y.-J.; Huang, N.-C.; Lee, C.-P.; and Hsieh, P.-C. 2024. Accelerated Policy Gradient: On the Convergence Rates of the Nesterov Momentum for Reinforcement Learning. *arXiv:2310.11897*.
- Chung, W.; Thomas, V.; Machado, M. C.; and Roux, N. L. 2021. Beyond Variance Reduction: Understanding the True Impact of Baselines on Policy Optimization. In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, 1999–2009. PMLR.
- Cohen, A.; Teng, E.; Berges, V.-P.; Dong, R.-P.; Henry, H.; Mattar, M.; Zook, A.; and Ganguly, S. 2022. On the Use and Misuse of Absorbing States in Multi-agent Reinforcement Learning. *RL in Games Workshop AAAI 2022*.
- Du, S.; Lee, J.; Li, H.; Wang, L.; and Zhai, X. 2019. Gradient descent finds global minima of deep neural networks. In *International conference on machine learning*, 1675–1685. PMLR.
- Dutta, A.; Das, S.; Nielsen, J.; Chakraborty, R.; and Shah, M. 2024. Multiview Aerial Visual Recognition (MAVREC): Can Multi-view Improve Aerial Visual Perception? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22678–22690.
- Ellenberger, B. 2018–2019. PyBullet Gymperium. <https://github.com/benlot/pybullet-gym>. Accessed: 2025-05-27.
- Fujimoto, S.; van Hoof, H.; and Meger, D. 2018. Addressing Function Approximation Error in Actor-Critic Methods. In Dy, J.; and Krause, A., eds., *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, 1587–1596. PMLR.
- Ganesh, S.; and Aggarwal, V. 2024. An accelerated multi-level monte carlo approach for average reward reinforcement learning with general policy parametrization. *arXiv e-prints*, arXiv-2407.
- Haarnoja, T.; Zhou, A.; Abbeel, P.; and Levine, S. 2018. Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor. *arXiv:1801.01290*.
- Ilyas, A.; Engstrom, L.; Santurkar, S.; Tsipras, D.; Janoos, F.; Rudolph, L.; and Madry, A. 2018. A closer look at deep policy gradients. *arXiv preprint arXiv:1811.02553*.
- Islam, R.; Henderson, P.; Gomrokchi, M.; and Precup, D. 2017. Reproducibility of Benchmarked Deep Reinforcement Learning Tasks for Continuous Control. *arXiv:1708.04133*.
- Juliani, A.; Berges, V.-P.; Teng, E.; Cohen, A.; Harper, J.; Elion, C.; Goy, C.; Gao, Y.; Henry, H.; Mattar, M.; and Lange, D. 2020. Unity: A general platform for intelligent agents. *arXiv preprint arXiv:1809.02627*.
- Konda, V.; and Tsitsiklis, J. 1999. Actor-critic algorithms. *Advances in neural information processing systems*, 12.
- Kumar, H.; Koppel, A.; and Ribeiro, A. 2023. On the sample complexity of actor-critic method for reinforcement learning with function approximation. *Machine Learning*, 112(7): 2433–2467.
- Lehmann, M. 2024. The definitive guide to policy gradients in deep reinforcement learning: Theory, algorithms and implementations. *arXiv preprint arXiv:2401.13662*.
- Lillicrap, T. P.; Hunt, J. J.; Pritzel, A.; Heess, N.; Erez, T.; Tassa, Y.; Silver, D.; and Wierstra, D. 2015. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*.
- McCandlish, S.; Kaplan, J.; Amodei, D.; and Team, O. D. 2018. An empirical model of large-batch training. *arXiv preprint arXiv:1812.06162*.
- Miyato, T.; Kataoka, T.; Koyama, M.; and Yoshida, Y. 2018. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*.
- Moré, J. J.; Garbow, B. S.; and Hillstom, K. E. 1994. User Manual for MINPACK-2. In *Technical Report ANL-94/21*, Argonne National Laboratory.
- Muehlebach, M.; and Jordan, M. I. 2021. Optimization with Momentum: Dynamical, Control-Theoretic, and Symplectic Perspectives. *arXiv:2002.12493*.
- Neyshabur, B. 2017. Implicit regularization in deep learning. *arXiv preprint arXiv:1709.01953*.
- Olshesky, A.; and Ghahesifard, B. 2023. A Small Gain Analysis of Single Timescale Actor Critic. *arXiv:2203.02591*.
- Omidshafiei, S.; Pazis, J.; Amato, C.; How, J. P.; and Vian, J. 2017. Deep Decentralized Multi-task Multi-Agent Reinforcement Learning under Partial Observability. In Precup, D.; and Teh, Y. W., eds., *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, 2681–2690. PMLR.
- Palma, P.; Smith, G.; Kaşk, M.; and Leite, T. 2018. Multi-Agent Reinforcement Learning with Hysteretic Q-Learning. In *Proceedings of the 17th European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD)*, 304–319.
- Peters, J.; and Schaal, S. 2006. Policy gradient methods for robotics. In *2006 IEEE/RSJ international conference on intelligent robots and systems*, 2219–2225. IEEE.
- Qiu, S.; Yang, Z.; Ye, J.; and Wang, Z. 2021. On finite-time convergence of actor-critic algorithm. *IEEE Journal on Selected Areas in Information Theory*, 2(2): 652–664.

- Raffin, A.; Hill, A.; Gleave, A.; Kanervisto, A.; Ernestus, M.; and Dormann, N. 2021. Stable-Baselines3: Reliable Reinforcement Learning Implementations. *Journal of Machine Learning Research*, 22(268): 1–8.
- Schulman, J.; Levine, S.; Abbeel, P.; Jordan, M.; and Moritz, P. 2015. Trust region policy optimization. In *International conference on machine learning*, 1889–1897. PMLR.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Sutskever, I.; Martens, J.; Dahl, G.; and Hinton, G. 2013. On the importance of initialization and momentum in deep learning. In Dasgupta, S.; and McAllester, D., eds., *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, 1139–1147. Atlanta, Georgia, USA: PMLR.
- Suttle, W. A.; Bedi, A.; Patel, B.; Sadler, B. M.; Koppel, A.; and Manocha, D. 2023. Beyond exponentially fast mixing in average-reward reinforcement learning via multi-level monte carlo actor-critic. In *International Conference on Machine Learning*, 33240–33267. PMLR.
- Sutton, R. S. 1984. *Temporal credit assignment in reinforcement learning*. University of Massachusetts Amherst.
- Sutton, R. S.; Barto, A. G.; et al. 1998. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge.
- Sutton, R. S.; McAllester, D.; Singh, S.; and Mansour, Y. 1999. Policy gradient methods for reinforcement learning with function approximation. *Advances in neural information processing systems*, 12.
- Tian, H.; Olshevsky, A.; and Paschalidis, Y. 2023. Convergence of actor-critic with multi-layer neural networks. *Advances in neural information processing systems*, 36: 9279–9321.
- Todorov, E.; Erez, T.; and Tassa, Y. 2012. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ international conference on intelligent robots and systems*, 5026–5033. IEEE.
- Towers, M.; Kwiatkowski, A.; Terry, J.; Balis, J. U.; Cola, G. D.; Deleu, T.; Goulão, M.; Kallinteris, A.; Krimmel, M.; KG, A.; Perez-Vicente, R.; Pierré, A.; Schulhoff, S.; Tai, J. J.; Tan, H.; and Younis, O. G. 2024. Gymnasium: A Standard Interface for Reinforcement Learning Environments. *arXiv:2407.17032*.
- Wang, L.; Cai, Q.; Yang, Z.; and Wang, Z. 2019a. Neural policy gradient methods: Global optimality and rates of convergence. *arXiv preprint arXiv:1909.01150*.
- Wang, Y.; He, H.; Tan, X.; and Gan, Y. 2019b. Trust Region-Guided Proximal Policy Optimization. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; d'Alché-Buc, F.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Wang, Y.; Wang, Y.; Zhou, Y.; and Zou, S. 2024. Non-asymptotic analysis for single-loop (natural) actor-critic with compatible function approximation. *arXiv preprint arXiv:2406.01762*.
- Williams, R. J. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8: 229–256.
- Wu, Y.; Mansimov, E.; Grosse, R. B.; Liao, S.; and Ba, J. 2017. Scalable trust-region method for deep reinforcement learning using kronecker-factored approximation. *Advances in neural information processing systems*, 30.
- Xiao, L. 2022. On the Convergence Rates of Policy Gradient Methods. *arXiv:2201.07443*.
- Xie, Z.; Zhang, Q.; Yang, F.; Hutter, M.; and Xu, R. 2025. Simple Policy Optimization. In Singh, A.; Fazel, M.; Hsu, D.; Lacoste-Julien, S.; Berkenkamp, F.; Maharaj, T.; Wagstaff, K.; and Zhu, J., eds., *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, 68813–68824. PMLR.
- Xu, T.; Liu, Q.; and Peng, J. 2018. Stochastic Variance Reduction for Policy Gradient Estimation. *arXiv:1710.06034*.
- Xu, T.; Wang, Z.; and Liang, Y. 2020a. Improving sample complexity bounds for (natural) actor-critic algorithms. *Advances in Neural Information Processing Systems*, 33: 4358–4369.
- Xu, T.; Wang, Z.; and Liang, Y. 2020b. Non-asymptotic convergence analysis of two time-scale (natural) actor-critic algorithms. *arXiv preprint arXiv:2005.03557*.
- Zhang, H.; Cisse, M.; Dauphin, Y. N.; and Lopez-Paz, D. 2017. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*.
- Zhang, S.; Liu, B.; Yao, H.; and Whiteson, S. 2020. Provably convergent two-timescale off-policy actor-critic with function approximation. In *International Conference on Machine Learning*, 11204–11213. PMLR.
- Zhou, M.; and Lu, J. 2023. Single timescale actor-critic method to solve the linear quadratic regulator with convergence guarantees. *Journal of Machine Learning Research*, 24(222): 1–34.