

---

# Compatible features for Monotonic Policy Improvement

---

**Marcin B. Tomczak\***  
PROWLER.io  
72 Hills Road, Cambridge, CB2 1LA, UK  
marcin@prowler.io

**Sergio Valcarcel Macua**  
PROWLER.io  
72 Hills Road, Cambridge, CB2 1LA, UK  
sergio@prowler.io

**Enrique Munoz de Cote**  
PROWLER.io  
72 Hills Road, Cambridge, CB2 1LA, UK  
enrique@prowler.io

**Peter Vrancx**  
PROWLER.io  
72 Hills Road, Cambridge, CB2 1LA, UK  
peter@prowler.io

## Abstract

Recent policy optimization approaches have achieved substantial empirical success by constructing surrogate optimization objectives. The Approximate Policy Iteration objective [1, 2] has become a standard optimization target for reinforcement learning problems. Using this objective in practice requires an estimator of the advantage function. Policy optimization methods such as those proposed in [3] estimate the advantages using a parametric critic. In this work we establish conditions under which the parametric approximation of the critic does not introduce bias to the updates of surrogate objective. These results hold for a general class of parametric policies, including deep neural networks. We obtain a result analogous to the compatible features derived for the original Policy Gradient Theorem [4]. As a result, we also identify a previously unknown bias that current state-of-the-art policy optimization algorithms [1, 5] have introduced by not employing these compatible features.

## 1 Introduction

Model-free reinforcement learning has been applied successfully to a variety of problems [6]. Nevertheless, many of the state-of-the-art approaches introduce bias to the policy updates and hence do not guarantee convergence. This might lead to oversensitivity of the algorithms to changes in initial conditions and has an adversarial influence on reproducibility [7].

Policy updates often use values learned by a critic to the variance of stochastic gradients. However, arbitrary parametrization of the critic can introduce bias into policy updates. Early work in [4] provides restrictions on the form of the function approximator under which the policy gradient remains unbiased and the convergence of learning is guaranteed. The notion of compatible features introduced in [4] also has connections to the natural policy gradient [8]. Compatible features remain an active area of research [9].

Modern policy optimization algorithms focus on optimizing a surrogate objective that approximates the value of the policy [1, 5]. This surrogate objective remains an accurate optimization target when the policy used to gather the data is close to the policy being the optimized [1, 2]. In this work we investigate the conditions on the parametrization of critic under which the updates of the surrogate objective remain unbiased.

---

\*Correspondence to mbt27@cam.ac.uk

## 2 Preliminaries

We assume a classical MDP formulation as in [4]. We consider an MDP being a tuple  $\langle S, A, P, R, \rho_0, \gamma \rangle$ ,  $S$  is a set of states,  $A$  is a set of actions,  $P : S \times A \times S \rightarrow \mathbb{R}_+$  is a transition model,  $R : S \times A \rightarrow \mathbb{R}$  is a reward function,  $\rho_0$  is an initial distribution over states  $S$  and  $\gamma \in (0, 1)$  is a discount factor. We denote a trajectory as  $\tau = (s_0, a_0, s_1, a_1, \dots)$ . Given a stochastic policy  $\pi : S \times A \rightarrow \mathbb{R}_+$ , we use  $\tau \sim \pi$  to denote that the trajectory has been generated by following policy  $\pi$ , i.e.  $s_{t+1} \sim P(\cdot | s_t, a_t)$ ,  $a_t \sim \pi(\cdot | s_t)$ , and  $s_0 \sim \rho_0(\cdot)$ .

We denote unnormalized discounted state occupancy measure as  $\rho_\pi(s) = \mathbb{E}_{\tau \sim \pi} \sum_{t \geq 0} \gamma^t P(s_t = s)$ . The value function is defined as  $V^\pi(s) = \mathbb{E}_{\tau \sim \pi} [\sum_{t \geq 0} \gamma^t R(s_t, a_t) | s_0 = s]$  and the state-value function is  $Q^\pi(s, a) = \mathbb{E}_{\tau \sim \pi} [\sum_{t \geq 0} \gamma^t R(s_t, a_t) | s_0 = s, a_0 = a]$ . The advantage function is defined as  $A^\pi(s, a) = Q^\pi(s, a) - V^\pi(s)$ . To emphasize that the policy  $\pi$  is parametrized by parameters  $\theta$  we use notation  $\pi_\theta$ . We consider policies such that  $\frac{\partial \pi_\theta(a|s)}{\partial \theta}$  exist and is continuous for every  $s \in S$  and  $a \in A$ . This is a general class of policies that includes, e.g., expressive deep neural networks. Under this discounted cost setting, reinforcement learning algorithms seek a policy  $\pi$  that maximizes its value defined as:

$$J(\pi) = \int \rho_0(s) V^\pi(s) ds. \quad (1)$$

When dealing with two different parametric policies,  $\pi_\theta$  and  $\pi_{\tilde{\theta}}$ , with parameters  $\theta$  and  $\tilde{\theta}$ , respectively, we follow the notation in [1] and define a surrogate policy optimization objective as

$$L_{\pi_\theta}(\pi_{\tilde{\theta}}) := J(\pi_\theta) + \mathbb{E}_{s \sim \rho^{\pi_\theta(\cdot)}, a \sim \pi_{\tilde{\theta}}(\cdot | s)} A^{\pi_\theta}(s, a). \quad (2)$$

When the expected divergence between  $\pi_\theta$  and  $\pi_{\tilde{\theta}}$  is small  $L_{\pi_\theta}(\pi_{\tilde{\theta}})$  can be treated as a good approximation of  $J(\pi_{\tilde{\theta}})$  [1, 2].

## 3 Related work

We begin by restating a well-known result presented in [4] that allows to calculate the gradient of  $J(\pi_\theta)$  w.r.t policy parameters  $\theta$ :

$$\frac{\partial J(\pi_\theta)}{\partial \theta} = \int \rho_{\pi_\theta}(s) \int \pi_\theta(a|s) \frac{\partial \log \pi_\theta(a|s)}{\partial \theta} Q^{\pi_\theta}(s, a) da ds. \quad (3)$$

The remarkable property of the expression for the policy gradient given by Equation (3) is that calculating the gradient  $\frac{\partial J(\pi_\theta)}{\partial \theta}$  does not require calculating the gradient of the occupancy measure  $\frac{\partial \rho_{\pi_\theta}}{\partial \theta}$ . Thus, if we know  $Q^{\pi_\theta}$ , the gradient  $\frac{\partial J(\pi_\theta)}{\partial \theta}$  can be approximated with Monte Carlo sampling using the trajectories obtained by following policy  $\pi_\theta$ .

However, since  $Q^{\pi_\theta}(s, a)$  is not known in advance, it has to be estimated. One possible choice of the estimator is to use empirical returns:  $\hat{Q}^{\pi_\theta}(s, a) = \sum_{t \geq t'} \gamma^t r(s_t, a_t)$ . The problem with this choice of estimator  $\hat{Q}^{\pi_\theta}(s, a)$  is the large variance of the stochastic version of the gradients  $\frac{\partial J(\pi_\theta)}{\partial \theta}$  which results in poor practical performance [10].

To reduce the variance in estimation of  $\frac{\partial J(\pi_\theta)}{\partial \theta}$  many algorithms learn a parametric approximation  $f_w(s, a)$  of  $Q^{\pi_\theta}(s, a)$  by solving the regression problem:

$$w^* = \operatorname{argmin}_w \frac{1}{N} \sum_{i=1}^N (Q^{\pi_\theta}(s_i, a_i) - f_w(s_i, a_i))^2 \quad (4)$$

where  $\{(s_i, a_i)\}_{i=1}^N$  are state action pairs sampled with policy  $\pi_\theta$ . Using such  $f_w^*(s, a)$  in place of  $Q^{\pi_\theta}(s, a)$  can introduce bias to policy updates. The following Theorem derived in [4, 11] provides conditions under which function approximator  $f_w(s, a)$  of  $Q^{\pi_\theta}(s, a)$  does not introduce bias to  $\frac{\partial J(\pi_\theta)}{\partial \theta}$ .

**Theorem 1.** Let  $f_w : S \times A \rightarrow \mathbb{R}$  be differentiable function approximator that satisfies the following conditions:

$$\int \rho_{\pi_\theta}(s) \int \pi_\theta(a|s) (Q^{\pi_\theta}(s, a) - f_w(s, a)) \frac{\partial f_w(s, a)}{\partial w} da ds = 0 \quad (5)$$

and

$$\frac{\partial f_w}{\partial w}(s, a) = \frac{1}{\pi_\theta(a|s)} \frac{\partial \pi_\theta(a|s)}{\partial \theta}. \quad (6)$$

Then,

$$\frac{\partial J(\pi_\theta)}{\partial \theta} = \int \rho_{\pi_\theta}(s) \int \frac{\partial \pi_\theta(a|s)}{\partial \theta} f_w(s, a) da ds. \quad (7)$$

Note that condition (5) can be satisfied by solving the regression problem  $w^* = \operatorname{argmin}_w \int \rho_{\pi_\theta}(s) \int \pi_\theta(a|s) (Q^{\pi_\theta}(s, a) - f_w(s, a))^2 da ds$ . Satisfying the condition given by (6) can be done by setting  $f_w(s, a) = w^\top \frac{\partial \log \pi_\theta(a|s)}{\partial \theta} + c_0$ . The constant  $c_0$  can be set to  $V^\pi(s)$  to ensure that  $\mathbb{E}_{a \sim \pi_\theta(\cdot|s)} f_w(s, a) = w^\top \mathbb{E}_{a \sim \pi_\theta(\cdot|s)} \frac{\partial \log \pi_\theta(a|s)}{\partial \theta} + V^\pi(s) = V^\pi(s)$ , since the expectation  $\mathbb{E}_{a \sim \pi_\theta(\cdot|s)} \frac{\partial \log \pi_\theta(a|s)}{\partial \theta} = \int \frac{\partial \pi_\theta(a|s)}{\partial \theta} da = \frac{\partial \int \pi_\theta(a|s) da}{\partial \theta} = 0$  for every  $s \in S$ . See the discussion in [4] for details.

Performing gradient ascent on  $J(\pi_\theta)$  requires resampling the data with every policy update as the expectation in Equation (3) is taken over the distribution of trajectories induced by  $\pi_\theta$ . The approach presented in [2, 1] tackles the problem of improving policy  $\pi_\theta$  in a different way.

Given the data gathered using the current policy  $\pi_\theta$  we want to estimate the lower bound on the performance of an arbitrary policy  $\pi_{\tilde{\theta}}$  and perform maximisation w.r.t  $\tilde{\theta}$ . Since the lower bound is tight at  $\pi_{\tilde{\theta}} = \pi_\theta$  maximizing this lower bound w.r.t  $\tilde{\theta}$  guarantees an improvement in the value of policy [1]. Hence, sequential optimization of the discussed lower bound is called Monotonic Policy Improvement. The lower bound derived in [1] is summarized in the following theorem.

**Theorem 2.** *Let  $\epsilon = \max_{s,a} |A^{\pi_\theta}(s, a)|$  and  $\alpha = \max_s \frac{1}{2} \int |\pi_\theta(a|s) - \pi_{\tilde{\theta}}(a|s)| da$ . Then the  $J(\pi_{\tilde{\theta}})$  can be lower bounded as follows:*

$$J(\pi_{\tilde{\theta}}) \geq L_{\pi_\theta}(\pi_{\tilde{\theta}}) - \frac{4\epsilon\gamma}{(1-\gamma)^2} \alpha^2. \quad (8)$$

The work done in [12] extends this results so that the max operator in  $\alpha$  can be replaced with the expected value taken w.r.t  $\rho^\pi$ . Calculating the gradient  $\frac{\partial L_{\pi_\theta}(\pi_{\tilde{\theta}})}{\partial \tilde{\theta}}$  is straightforward as the occupancy measure  $\rho_{\pi_\theta}$  in the definition of  $L_{\pi_\theta}(\pi_{\tilde{\theta}})$  does not depend on  $\tilde{\theta}$ .

In the practical setting  $L_{\pi_\theta}(\pi_{\tilde{\theta}})$  is optimized by constraining the divergence between  $\pi_\theta$  and  $\pi_{\tilde{\theta}}$  [1, 5]. The algorithm derived in [1] is closely similar to the natural gradient policy optimisation [8]. Note that  $L_{\pi_\theta}(\pi_{\tilde{\theta}})$  is a biased approximation of  $J(\pi_{\tilde{\theta}})$  but the bias can be controlled by restricting the distance between  $\pi_{\tilde{\theta}}$  and  $\pi_\theta$ .

Optimizing  $L_{\pi_\theta}(\pi_{\tilde{\theta}})$  requires knowing the values of  $A^{\pi_\theta}(s, a)$ . There are various ways of estimating the critic  $\hat{A}^{\pi_\theta}(s, a)$ , for instance see [3]. However, similarly to the case of Theorem 1, using an arbitrary parametrization of the critic introduces bias to an estimate of  $\frac{\partial L_{\pi_\theta}(\pi_{\tilde{\theta}})}{\partial \pi_{\tilde{\theta}}}$ .

## 4 Compatible features for surrogate policy optimization

In this section, we seek a parametric form  $f_w(s, a)$  of an approximator of  $Q^{\pi_\theta}(s, a)$  for which  $\frac{\partial L_{\pi_\theta}(\pi_{\tilde{\theta}})}{\partial \tilde{\theta}}$  remains unbiased. To this end, we follow the approach presented in [4]. We derive the following theorem.

**Theorem 3** (Compatible features for Monotonic Policy Improvement). *Assuming the following condition is satisfied:*

$$\int \rho_{\pi_\theta}(s) \int \pi_\theta(a|s) (Q^{\pi_\theta}(s, a) - f_w(s, a)) \frac{\partial f_w(s, a)}{\partial w} da ds = 0 \quad (9)$$

and

$$f_w(s, a) = w^\top \frac{\pi_{\tilde{\theta}}(a|s)}{\pi_\theta(a|s)} \frac{\partial \log \pi_{\tilde{\theta}}(a|s)}{\partial \tilde{\theta}} + c_0. \quad (10)$$

Then,

$$\frac{\partial L_{\pi_\theta}(\pi_{\tilde{\theta}})}{\partial \tilde{\theta}} = \int \rho_{\pi_\theta}(s) \int \frac{\partial \pi_{\tilde{\theta}}(a|s)}{\partial \tilde{\theta}} f_w(s, a) da ds. \quad (11)$$

*Proof.* Firstly, we note that the value function  $V^{\pi_\theta}$  in the definition of  $L_{\pi_\theta}(\pi_{\tilde{\theta}})$  fulfils the role of the control variate, i.e. it does not influence the expectation of the gradient. To note this, we analyse the gradient  $\frac{\partial L_{\pi_\theta}(\pi_{\tilde{\theta}})}{\partial \tilde{\theta}}$ :

$$\begin{aligned}
\frac{\partial L_{\pi_\theta}(\pi_{\tilde{\theta}})}{\partial \tilde{\theta}} &= \frac{\partial}{\partial \tilde{\theta}} \int \rho_{\pi_\theta}(s) \int \pi_{\tilde{\theta}}(a|s) A^{\pi_\theta}(s, a) dads \\
&= \int \rho_{\pi_\theta}(s) \int \frac{\partial \pi_{\tilde{\theta}}(a|s)}{\partial \tilde{\theta}} A^{\pi_\theta}(s, a) dads \\
&= \int \rho_{\pi_\theta}(s) \int \frac{\partial \pi_{\tilde{\theta}}(a|s)}{\partial \tilde{\theta}} (Q^{\pi_\theta}(s, a) - V^{\pi_\theta}(s)) dads \\
&= \int \rho_{\pi_\theta}(s) \left[ \int \frac{\partial \pi_{\tilde{\theta}}(a|s)}{\partial \tilde{\theta}} Q^{\pi_\theta}(s, a) da - \int \frac{\partial \pi_{\tilde{\theta}}(a|s)}{\partial \tilde{\theta}} V^{\pi_\theta}(s) da \right] ds \\
&= \int \rho_{\pi_\theta}(s) \int \frac{\partial \pi_{\tilde{\theta}}(a|s)}{\partial \tilde{\theta}} Q^{\pi_\theta}(s, a) dads, \tag{12}
\end{aligned}$$

where the second step is allowed since  $\frac{\partial \pi_{\tilde{\theta}}(a|s)}{\partial \tilde{\theta}}$  exist and is continuous  $\forall s \in S, \forall a \in A$ ; and the last step is due to the fact that  $V^{\pi_\theta}$  does not depend on  $\tilde{\theta}$ , i.e.:  $\int \frac{\partial \pi_{\tilde{\theta}}(a|s)}{\partial \tilde{\theta}} V^{\pi_\theta}(s) da = \frac{\partial V^{\pi_\theta}(s)}{\partial \tilde{\theta}} = 0$ .

Next we define  $\mathcal{L}_{\pi_\theta, w}(\pi_{\tilde{\theta}}) := \int \rho_{\pi_\theta}(s) \int \pi_{\tilde{\theta}}(a|s) f_w(s, a) dads$ . We seek a condition under which the following equality holds:

$$\frac{\partial L_{\pi_\theta}}{\partial \tilde{\theta}} - \frac{\partial \mathcal{L}_{\pi_\theta, w}}{\partial \tilde{\theta}} = 0. \tag{13}$$

By subtracting  $\frac{\partial \mathcal{L}_{\pi_\theta, w}}{\partial \tilde{\theta}}$  and the assumption given by (9) from (12), we obtain:

$$\begin{aligned}
\frac{\partial L_{\pi_\theta}(\pi_{\tilde{\theta}})}{\partial \tilde{\theta}} - \frac{\partial \mathcal{L}_{\pi_\theta, w}(\pi_{\tilde{\theta}})}{\partial \tilde{\theta}} &= \int \rho_{\pi_\theta}(s) \frac{\partial \pi_{\tilde{\theta}}(a|s)}{\partial \tilde{\theta}} (Q^{\pi_\theta}(s, a) - f_w(s, a)) dads \\
&= \int \rho_{\pi_\theta}(s) \int \pi_{\tilde{\theta}}(a|s) \frac{\partial \log \pi_{\tilde{\theta}}(a|s)}{\partial \tilde{\theta}} (Q^{\pi_\theta}(s, a) - f_w(s, a)) dads \\
&\quad - \int \rho_{\pi_\theta}(s) \int \pi_{\tilde{\theta}}(a|s) (Q^{\pi_\theta}(s, a) - f_w(s, a)) \frac{\partial f_w(s, a)}{\partial w} dads \\
&= \int \rho_{\pi_\theta}(s) \int \pi_{\tilde{\theta}}(a|s) \frac{\pi_{\tilde{\theta}}(a|s)}{\pi_\theta(a|s)} \frac{\partial \log \pi_{\tilde{\theta}}(a|s)}{\partial \tilde{\theta}} (Q^{\pi_\theta}(s, a) - f_w(s, a)) dads \\
&\quad - \int \rho_{\pi_\theta}(s) \int \pi_{\tilde{\theta}}(a|s) (Q^{\pi_\theta}(s, a) - f_w(s, a)) \frac{\partial f_w(s, a)}{\partial w} dads \\
&= \int \rho_{\pi_\theta}(s) \int \pi_{\tilde{\theta}}(a|s) (Q^{\pi_\theta}(s, a) - f_w(s, a)) \left[ \frac{\pi_{\tilde{\theta}}(a|s)}{\pi_\theta(a|s)} \frac{\partial \log \pi_{\tilde{\theta}}(a|s)}{\partial \tilde{\theta}} - \frac{\partial f_w(s, a)}{\partial w} \right] dads. \tag{14}
\end{aligned}$$

where we have used the log trick,  $\pi_{\tilde{\theta}}(a|s) \frac{\partial \log \pi_{\tilde{\theta}}(a|s)}{\partial \tilde{\theta}} = \frac{\partial \pi_{\tilde{\theta}}(a|s)}{\partial \tilde{\theta}}$ , in the second line. Hence Equation (13) can be satisfied by requiring:

$$\frac{\pi_{\tilde{\theta}}}{\pi_\theta} \frac{\partial \log \pi_{\tilde{\theta}}}{\partial \tilde{\theta}} - \frac{\partial f_w}{\partial w} = 0. \tag{15}$$

Integrating this last equation w.r.t  $w$  yields:  $f_w(s, a) = w^\top \frac{\pi_{\tilde{\theta}}(a|s)}{\pi_\theta(a|s)} \frac{\partial \log \pi_{\tilde{\theta}}(a|s)}{\partial \tilde{\theta}} + c_0$ , which completes the proof.  $\square$

Again setting  $c_0 = V^{\pi_\theta}(s)$  ensures that:

$$\begin{aligned}
\mathbb{E}_{a \sim \pi_\theta(\cdot|s)} f_w(s, a) &= w^\top \mathbb{E}_{a \sim \pi_\theta(\cdot|s)} \frac{\pi_{\tilde{\theta}}(a|s)}{\pi_\theta(a|s)} \frac{\partial \log \pi_{\tilde{\theta}}(a|s)}{\partial \tilde{\theta}} + V^{\pi_\theta}(s) \\
&= w^\top \mathbb{E}_{a \sim \pi_{\tilde{\theta}}(\cdot|s)} \frac{\partial \log \pi_{\tilde{\theta}}(a|s)}{\partial \tilde{\theta}} + V^{\pi_\theta}(s) = w^\top \frac{\partial \int \pi_{\tilde{\theta}}(a|s) da}{\partial \tilde{\theta}} + V^{\pi_\theta}(s) \\
&= V^{\pi_\theta}(s), \quad \forall s \in S, \forall a \in A. \tag{16}
\end{aligned}$$

In reference [4] the authors conjecture that the compatible features  $\frac{\partial \log \pi_\theta(a|s)}{\partial \theta}$  might be the only choice of features that lead to an unbiased policy gradient  $\frac{\partial J(\pi_\theta)}{\partial \theta}$ . In the case of  $L_{\pi_\theta}(\pi_{\tilde{\theta}})$  we can provide another choice of features leading to an unbiased gradient  $\frac{\partial L_{\pi_\theta}(\pi_{\tilde{\theta}})}{\partial \tilde{\theta}}$ . The features derived in Theorem 3 depend on importance sampling weights  $\frac{\pi_{\tilde{\theta}}(a|s)}{\pi_\theta(a|s)}$ . To remove importance sampling weights  $\frac{\pi_{\tilde{\theta}}(a|s)}{\pi_\theta(a|s)}$  from derived compatible features we modify the condition in Equation (9) by replacing action sampling distribution  $\pi_\theta(\cdot|s)$  with  $\pi_{\tilde{\theta}}(\cdot|s)$ .

**Theorem 4.** *Assuming that the following condition is satisfied:*

$$\int \rho_{\pi_\theta}(s) \int \pi_{\tilde{\theta}}(a|s) (Q^{\pi_\theta}(s, a) - f_w(s, a)) \frac{\partial f_w(s, a)}{\partial w} dad s = 0 \quad (17)$$

and

$$f_w(s, a) = w^\top \frac{\partial \log \pi_{\tilde{\theta}}(a|s)}{\partial \tilde{\theta}} + c_0. \quad (18)$$

Then,

$$\frac{\partial L_{\pi_\theta}(\pi_{\tilde{\theta}})}{\partial \tilde{\theta}} = \int \rho_{\pi_\theta}(s) \int \frac{\partial \pi_{\tilde{\theta}}(a|s)}{\partial \tilde{\theta}} f_w(s, a) dad s. \quad (19)$$

*Proof.* We derive the result by following a similar line of thought as in proof of Theorem 3. Again we use  $\mathcal{L}_{\pi_\theta, w}(\pi_{\tilde{\theta}}) = \int \rho_{\pi_\theta}(s) \int \pi_{\tilde{\theta}}(a|s) f_w(s, a) dad s$ . We subtract  $\frac{\partial \mathcal{L}_{\pi_\theta, w}(\pi_{\tilde{\theta}})}{\partial \tilde{\theta}}$  and the assumption given by (17) from  $\frac{\partial L_{\pi_\theta}(\pi_{\tilde{\theta}})}{\partial \tilde{\theta}}$ , which yields:

$$\begin{aligned} & \frac{\partial L_{\pi_\theta}(\pi_{\tilde{\theta}})}{\partial \tilde{\theta}} - \frac{\partial \mathcal{L}_{\pi_\theta, w}(\pi_{\tilde{\theta}})}{\partial \tilde{\theta}} = \\ &= \int \rho_{\pi_\theta}(s) \int \pi_{\tilde{\theta}}(a|s) \frac{\partial \log \pi_{\tilde{\theta}}(a|s)}{\partial \tilde{\theta}} (Q^{\pi_\theta}(s, a) - f_w(s, a)) \\ & \quad - \int \rho_{\pi_\theta}(s) \int \pi_{\tilde{\theta}}(a|s) (Q^{\pi_\theta}(s, a) - f_w(s, a)) \frac{\partial f_w(s, a)}{\partial w} dad s \\ &= \int \rho_{\pi_\theta}(s) \int \pi_{\tilde{\theta}}(a|s) (Q^{\pi_\theta}(s, a) - f_w(s, a)) \left[ \frac{\partial \log \pi_{\tilde{\theta}}(a|s)}{\partial \tilde{\theta}} - \frac{\partial f_w(s, a)}{\partial w} \right] dad s. \quad (20) \end{aligned}$$

The integral in Equation (20) can be set to zero by setting  $f_w$  to solve the differential equation:  $\frac{\partial \log \pi_{\tilde{\theta}}}{\partial \tilde{\theta}} - \frac{\partial f_w}{\partial w} = 0$ , which results in desired compatible features  $f_w(s, a) = w^\top \frac{\partial \log \pi_{\tilde{\theta}}(a|s)}{\partial \tilde{\theta}} + c_0$ .  $\square$

Note the features derived in Theorem 4 are analogous to ones derived in Theorem 1 with the difference that the occupancy measure is taken w.r.t the data gathering policy  $\pi_\theta$  in the case of Theorem 4.

Assumption (17) in Theorem 4 has a natural interpretation. Intuitively, the critic values  $f_w(s, a)$  should be more accurate for actions with high probability under policy  $\pi_{\tilde{\theta}}$ . More formally, the condition in (17) can be satisfied by finding a parameter  $w^*$  that minimises the weighted quadratic error:  $w^* = \arg \min_w \int \rho_{\pi_\theta}(s) \int \pi_\theta(a|s) \frac{\pi_{\tilde{\theta}}(a|s)}{\pi_\theta(a|s)} (Q^{\pi_\theta}(s, a) - f_w(s, a))^2 dad s$ , where weights  $\frac{\pi_{\tilde{\theta}}(a|s)}{\pi_\theta(a|s)}$  ensure improved accuracy for likely actions  $a$  given  $s$  under target policy  $\pi_{\tilde{\theta}}(\cdot|s)$ . Similarly, when target policy  $\pi_{\tilde{\theta}}(\cdot|s)$  assigns low probability to an action  $a$  the error in the critic  $f_w(s, a)$  estimation is not relevant as it will not influence the gradient  $\frac{\partial L_{\pi_\theta}(\pi_{\tilde{\theta}})}{\partial \tilde{\theta}}$ .

Given a sequence of state action pairs  $\{(s_t, a_t)\}_{t=1}^T$  gathered by following policy  $\pi_\theta$ , the integrals in condition (17) can be also approximated with samples, which leads to a weighted regression problem emerging from Theorem 4:

$$w^* = \operatorname{argmin}_w \frac{1}{T} \sum_{s, a} \frac{\pi_{\tilde{\theta}}(a|s)}{\pi_\theta(a|s)} (Q^{\pi_\theta}(s, a) - f_w(s, a))^2. \quad (21)$$

## 5 Experiments

In this section we use the NChain [13] environment to compare the gradient  $\frac{\partial L_{\pi_{\tilde{\theta}}}(\pi_{\tilde{\theta}})}{\partial \tilde{\theta}}$  calculated in two different ways: i) using a standard linear critic of the form  $f_w(s, a) = w_2s + w_1a + w_0$ , learned with the standard least squares regression approach given by (4); and ii) using a linear critic with compatible features given by (18), learned by solving the weighted the least squared regression problem (21). We use the following parametric policy function  $\pi_{\theta}(a|s) = \sigma(\theta_1s + \theta_2a)$ . Hence, the proposed compatible features are given by  $\frac{\partial \log \pi_{\tilde{\theta}}(a|s)}{\partial \tilde{\theta}} = [(1 - \sigma(\theta_1s + \theta_2a))s, (1 - \sigma(\theta_1s + \theta_2a))a]^\top$ , since  $\frac{d\sigma(x)}{dx} = \sigma(x)(1 - \sigma(x))$  and  $\frac{d \log \sigma(f(x))}{dx} = (1 - \sigma(f(x))) \frac{df(x)}{dx}$ . We provide true state action value function  $Q^{\pi_{\theta}}(s, a)$  as targets to learn critics. We estimate the gradient  $\frac{\partial L_{\pi_{\tilde{\theta}}}(\pi_{\tilde{\theta}})}{\partial \tilde{\theta}}$  using rollouts from policy  $\pi_{\theta}$ . We set policy parameters to  $\theta_1 = 0.2$ ;  $\theta_2 = 0.5$  and  $\tilde{\theta}_1 = 0.3$ ;  $\tilde{\theta}_2 = 0.6$ . We use the expression for  $\frac{\partial L_{\pi_{\tilde{\theta}}}(\pi_{\tilde{\theta}})}{\partial \tilde{\theta}}$  given by Equation (12).

We compare the obtained gradients to ground truth calculated using the true state action value function  $Q^{\pi_{\theta}}(s, a)$ . In both cases, we provide an increasing number of rollouts from policy  $\pi_{\theta}$  to both approaches and analyse the errors in estimation. We report bias, variance and RMSE of estimated gradients in Figure 1.




Figure 1: Comparison of gradient  $\frac{\partial L_{\pi_{\tilde{\theta}}}(\pi_{\tilde{\theta}})}{\partial \tilde{\theta}}$  derived with compatible parametrization and by using standard critic. Employing a critic with derived compatible features allows to provide unbiased gradient of  $\frac{\partial L_{\pi_{\tilde{\theta}}}(\pi_{\tilde{\theta}})}{\partial \tilde{\theta}}$ . The results are averaged over 250 trials.

As expected, the bias in the estimation of  $\frac{\partial L_{\pi_{\tilde{\theta}}}(\pi_{\tilde{\theta}})}{\partial \tilde{\theta}}$  introduced by the standard critic cannot be removed by increasing the number of provided trajectories, as it is caused by employing a noncompatible function approximator. The derived compatible features are unbiased by construction, hence provide substantially improved quality of estimation of  $\frac{\partial L_{\pi_{\tilde{\theta}}}(\pi_{\tilde{\theta}})}{\partial \tilde{\theta}}$ . Interestingly, using compatible features also provides slightly lower variance than using the standard critic.

## 6 Conclusions

We have analysed the use of parametric critics to estimate the policy optimization surrogate objective in a systematical way. As a consequence, we can provide conditions under which the approximation does not introduce bias to the policy updates. This result holds for a general class of policies, including policies parametrized by deep neural networks. We have shown that for the investigated surrogate objective there exists two different choices of compatible features. We empirically demonstrated that the compatible features allow to estimate the gradient of surrogate objective more accurately.

## References

- [1] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 1889–1897, Lille, France, 07–09 Jul 2015. PMLR.
- [2] Sham Kakade and John Langford. Approximately optimal approximate reinforcement learning. In *Proceedings of the Nineteenth International Conference on Machine Learning*, ICML '02, pages 267–274, San Francisco, CA, USA, 2002. Morgan Kaufmann Publishers Inc.
- [3] John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. High-dimensional continuous control using generalized advantage estimation. *arXiv preprint arXiv:1506.02438*, 2015.
- [4] Richard S. Sutton, David McAllester, Satinder Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. In *Proceedings of the 12th International Conference on Neural Information Processing Systems*, NIPS'99, pages 1057–1063, Cambridge, MA, USA, 1999. MIT Press.
- [5] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *CoRR*, abs/1707.06347, 2017.
- [6] Yan Duan, Xi Chen, Rein Houthoofd, John Schulman, and Pieter Abbeel. Benchmarking deep reinforcement learning for continuous control. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, ICML'16, pages 1329–1338. JMLR.org, 2016.
- [7] Peter Henderson, Riashat Islam, Philip Bachman, Joelle Pineau, Doina Precup, and David Meger. Deep reinforcement learning that matters. *CoRR*, abs/1709.06560, 2017.
- [8] Sham Kakade. A natural policy gradient. In *Proceedings of the 14th International Conference on Neural Information Processing Systems: Natural and Synthetic*, NIPS'01, pages 1531–1538, Cambridge, MA, USA, 2001. MIT Press.
- [9] Joni Pajarinen, Hong Linh Thai, Riad Akrou, Jan Peters, and Gerhard Neumann. Compatible natural gradient policy search. *CoRR*, abs/1902.02823, 2019.
- [10] Kamil Ciosek and Shimon Whiteson. Expected policy gradients. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [11] Vijay R. Konda and John N. Tsitsiklis. On actor-critic algorithms. *SIAM J. Control Optim.*, 42(4):1143–1166, April 2003.
- [12] Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. Constrained policy optimization. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 22–31, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR.
- [13] Malcolm J. A. Strens. A bayesian framework for reinforcement learning. In *Proceedings of the Seventeenth International Conference on Machine Learning*, ICML '00, pages 943–950, San Francisco, CA, USA, 2000. Morgan Kaufmann Publishers Inc.