
Off-Policy Policy Gradient with State Distribution Correction

Yao Liu

Stanford University
yaoliu@stanford.edu

Alekh Agarwal

Microsoft Research
alekha@microsoft.com

Adith Swaminathan

Microsoft Research
adswamin@microsoft.com

Emma Brunskill

Stanford University
ebrun@cs.stanford.edu

Abstract

We study the problem of off-policy policy optimization in Markov decision processes, and develop a novel off-policy policy gradient method. Prior off-policy policy gradient approaches have generally ignored the mismatch between the distribution of states visited under the behavior policy used to collect data, and what would be the distribution of states under the learned policy. Here we build on recent progress for estimating the ratio of the state distributions under behavior and evaluation policies for policy evaluation, and present an off-policy policy gradient optimization technique that can account for this mismatch in distributions. We present an illustrative example of why this is important and a theoretical convergence guarantee for our approach. Empirically, we compare our method in simulations to several strong baselines which do not correct for this mismatch, significantly improving in the quality of the policy discovered.

1 INTRODUCTION

The ability to use data about prior decisions and their outcomes to make counterfactual inferences about how alternative decision policies might perform, is a cornerstone of intelligent behavior. It also has immense practical potential – it can enable the use of electronic medical record data to infer better treatment decisions for patients, the use of prior product recommendations to inform more effective strategies for presenting recommendations, and previously collected data from students using educational software to better teach those and future students. Such counterfactual reasoning, particularly when one is deriving decision policies that will be

used to make not one but a sequence of decisions, is important since online sampling during a learning procedure is both costly and dangerous, and not practical in many of the applications above. While amply motivated, doing such counterfactual reasoning is also challenging because the data is censored – we can only observe the result of providing a particular chemotherapy treatment policy to a particular patient, not the counterfactual of if we were then to start with a radiation sequence.

We focus on the problem of performing such counterfactual inferences in the context of sequential decision making in a Markov decision process (MDP). We assume that data has been previously collected using some fixed and known behavior policy, and our goal is to learn a new decision policy with good performance for future use. This problem is often known as batch off-policy policy optimization. We assume that the behavior policy used to gather the data is stochastic: if it is deterministic, without any additional assumptions, we will not be able to estimate the performance of any other policy.

In this paper we consider how to perform batch off-policy policy optimization (OPPO) using a policy gradient method. While there has been increasing interest in batch off-policy reinforcement learning (RL) over the last few years [Thomas et al., 2015, Jiang and Li, 2016, Thomas and Brunskill, 2016], much of this has focused on off-policy policy evaluation, where the goal is to estimate the performance of a particular given target decision policy. Ultimately we will very frequently be interested in the optimization question, which requires us to determine a good new policy for future potential deployment, given a fixed batch of prior data.

To do batch off-policy policy optimization, value function methods (like Deep Q-Network [Mnih et al., 2015] or Fitted Q-Iteration [Ernst et al., 2005]) can be used alone, but there are many cases where we might prefer to focus on policy gradient or actor-critic methods. Policy gradient methods have seen substantial success

in the last few years [Schulman et al., 2015] in the on-policy setting, and they can be particularly appealing for cases where it is easier to encode inductive bias in the policy space, or when the actions are continuous (see e.g. Abbeel and Schulman [2016] for more discussion). However, existing approaches incorporating offline information into online policy gradients [Gu et al., 2017b,a, Metelli et al., 2018] have shown limited benefit, in part due to the variance in gradients incurred due to incorporating off-policy data. One approach is to correct exactly for the difference between the sampling data distribution and the target policy data distribution, by using importance sampling to re-weight every sample according to the likelihood ratio of behavior policy and evaluation policy up to that step. Unfortunately the variance of this importance sampling ratio will grow exponentially with the problem horizon.

To avoid introducing variance in gradients, off-policy actor critic (Off-PAC) [Degris et al., 2012] ignores the state distribution difference between the behavior and target policies, and instead only uses a one step of importance sampling to reweight the action distributions. Many practical off-policy policy optimization algorithms including DDPG [Silver et al., 2014], ACER [Wang et al., 2016], and Off-PAC with emphatic weightings [Imani et al., 2018] are based on the gradient expression in the Off-PAC algorithm [Degris et al., 2012]. However as we will demonstrate, not correcting for this mismatch in state distributions can result in poor performance in general, both in theory and empirically.

Instead, here we introduce an off-policy policy gradient algorithm that can be used with batch data and that accounts for the difference in the state distributions between the current target and behavior policies during each gradient step. Our approach builds on recent approaches for policy evaluation that avoid the exponential blow up in importance sampling weights by instead directly computing the ratio of the distributions of state visitation under the target and behavior policies [Hallak and Mannor, 2017, Liu et al., 2018a, Gelada and Belle-mare, 2019]. We incorporate these ideas within an off-policy actor critic method to do batch policy optimization. We first provide an illustrative example to demonstrate the benefit of this approach over Off-PAC [Degris et al., 2012], and show that correcting for the mismatch in state distributions of the behavior and target policies can be critical for getting good estimates of the policy gradient, and we also provide convergence guarantees for our algorithm under certain assumptions. We then compare our approach and Off-PAC experimentally on two simulated domains, cart pole and a HIV patient simulator [Ernst et al., 2005]. Our results show that our approach is able to learn a substantially higher performing

policy than both Off-PAC and the behavior policy that is used to gather the batch data. While they have not been previously evaluated in the batch setting, we also compare our method against two more recent policy optimization approaches which also make use of available off-policy data. We find that without the several enhancements on top of vanilla actor-critic method which these methods incorporate, we are able to outperform them in our simulations by virtue of the state distribution correction. We further demonstrate that we can use the recently proposed off-policy evaluation technique of Liu et al. [2018a] to reliably identify good policies found during the policy gradient optimization run. Our results suggest that directly accounting for the state distribution mismatch can be done without prohibitively increasing the variance during policy gradient evaluations, and that doing so can yield significantly better policies. These results are promising for enabling us to learn better policies given batch data or improving the sample efficiency of online policy gradient methods by being able to better incorporate past data.

Related Work Many prior works focus on the off-policy policy evaluation (OPPE) problem, as it is a foundation for downstream policy learning problems. These approaches often build on importance sampling techniques to correct for distribution mismatch in the trajectory space, pioneered by the early work on eligibility traces [Precup et al., 2000], and further enhanced with a variety of variance reduction techniques [Thomas et al., 2015, Jiang and Li, 2016, Thomas and Brunskill, 2016]. Some consider model-based approaches to OPPE [Farajtabar et al., 2018, Liu et al., 2018b], which usually perform better than importance sampling approaches empirically in policy evaluation settings. But those methods do not extend easily to our OPPO setting, as well as introduce additional challenges due to bias in the models and typically require fitting a separate model for each target policy. The recent work of Liu et al. [2018a] partially alleviates the variance problem for model-free OPPE by reweighting the state visitation distributions, which can result in as just as high a variance in the worst case, but is often much smaller. Our work incorporates this recent estimator in policy optimization methods to enable learning from off-policy collected data.

In the off-policy policy optimization setting, many works study value-function based approaches (like fitted Q iteration [Ernst et al., 2005] and DQN [Mnih et al., 2015]), as they are known to be more robust to distribution mismatch. Some recent works aim to further incorporate reweighting techniques within off-policy value function learning [Hallak and Mannor, 2017, Gelada and Belle-mare, 2019]. These methods hint at the intriguing po-

tential of value-function based techniques for off-policy learning, and we are interested in similarly understanding the viability of using direct policy optimization techniques in the off-policy setting.

Off-policy actor critic method [Degris et al., 2012, Imani et al., 2018] propose a partial answer to this question by learning the critic using off-policy data and reweighting actor gradients by correcting the conditional action probabilities, but ignore the mismatch between the state distributions. A different research thread on trust region policy optimization method [Schulman et al., 2015], while requiring the on-policy setting, incorporates robustness to the mismatch between the data collection and gradient evaluation policies. However this is not a fully off-policy scenario and learning from an offline dataset is still strongly motivated by many applications. Somewhat related, many recent off-policy RL algorithms [Silver et al., 2014, Wang et al., 2016, Gu et al., 2017a,b, Lillicrap et al., 2015, Haarnoja et al., 2018, Dai et al., 2018] improve the empirical sample efficiency by using more off-policy samples from the previous iterations. In contrast to our setting with a fixed dataset collected with a behavior policy, these methods are online off-policy in that they still collect data with each policy found in the optimization procedure, but also use previous data for added sample efficiency. We nevertheless compare to two such methods in our empirical evaluation.

2 PRELIMINARIES

We consider finite horizon MDPs $M = \langle \mathcal{S}, \mathcal{A}, P, r, \gamma \rangle$, with a continuous state space \mathcal{S} , a discrete action space \mathcal{A} , a transition probability distribution $P : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \mapsto [0, 1]$ and an expected reward function $r : \mathcal{S} \times \mathcal{A} \mapsto [0, 1]$. We observe tuples of state, action, reward and next state: (s_t, a_t, r_t, s_{t+1}) , where s_0 is drawn from a initial state distribution $p_0(s)$, action a is drawn from a stochastic behavior policy $\mu(a|s)$ and the reward and next state are generated by the MDP. Given a discount factor $\gamma \in (0, 1]$, the goal is to maximize the expected return of policy:

$$R_M^\pi = \mathbb{E}_\pi \left[\lim_{T \rightarrow \infty} \frac{1}{\sum_{t=0}^T \gamma^t} \sum_{t=0}^T \gamma^t r_t \right] \quad (1)$$

When $\gamma = 1$ this becomes the average reward case and $\gamma < 1$ is called the discounted reward case. In this paper, we focus on $\gamma < 1$ for most of the results. Given any fixed policy π the MDP becomes a Markov chain and we can define the state distribution at time step t : $d_t^\pi(s)$, and the stationary state distribution across time: $d^\pi(s) = \lim_{T \rightarrow \infty} \frac{1}{\sum_{t=0}^T \gamma^t} \sum_{t=0}^T \gamma^t d_t^\pi(s)$ We consider infinite horizon for the convenience of defining stationary

distribution. A similar scheme should work for finite horizon, but requires time-dependent state distribution and we omit this here. To make sure the optimal policy is learnable from collected data, we assume the following about the support set of behavior policy:

Assumption 1. *For at least one optimal policy π^* , $d^\mu(s) > 0$ for all s such that $d^{\pi^*}(s) > 0$, and $\mu(a|s) > 0$ for all a such that $\pi^*(a|s) > 0$ when $d^{\pi^*}(s) > 0$.*

3 AN OFF-POLICY POLICY GRADIENT ESTIMATOR

Note that Assumption 1 is quite weak when designing a policy evaluation or optimization scheme, since it only guarantees that μ adequately visits all the states and actions visited by some π^* . However, a policy optimization algorithm might require off-policy policy gradient estimates at arbitrary intermediate policy it produces along the way, which might visit states not reached by μ . A strong assumption to handle such scenarios is that Assumption 1 holds not just for some π^* , but any possible policy π . Instead of making such a strong assumption, we start by defining an augmented MDP where Assumption 1 suffices for obtaining pessimistic estimates of policy values and gradients.

Constructing an Augmented MDP Given a data collection policy μ , let its support set be $\mathcal{S}_\mu = \{s : d^\mu(s) > 0\}$ and $\mathcal{SA}_\mu = \{(s, a) : d^\mu(s)\mu(a|s) > 0\}$. Consider a modified MDP $M_\mu = \langle \mathcal{S}_\mu \cup \{s_{abs}\}, \mathcal{A}, P_\mu, r_\mu, \gamma \rangle$. Any state-action pairs not in \mathcal{SA}_μ will essentially transition to s_{abs} which is a new absorbing state where all actions will lead to a zero reward self-loop. Concretely, $P_\mu(s_{abs}|s_{abs}, a) = 1$ and $r(s_{abs}, a) = 0$ for any a . For all other states, the transition probabilities and rewards are defined as: For $(s, a) \in \mathcal{SA}_\mu$, $P_\mu(s'|s, a) = P(s'|s, a)$ for all $s' \in \mathcal{S}_\mu$, and $P_\mu(s_{abs}|s, a) = \int_{s' \notin \mathcal{S}_\mu} P(s'|s, a) ds'$. For all $s \in \mathcal{S}_\mu$ but $(s, a) \notin \mathcal{SA}_\mu$, $P_\mu(s_{abs}|s, a) = 1$. $r_\mu(s, a) = r(s, a)$ for $(s, a) \in \mathcal{SA}_\mu$, and $r_\mu(s, a) = 0$ otherwise. First we prove that the optimal policy π^* of the original MDP remains optimal in augmented MDP as a consequence of Assumption 1.

Theorem 1. *The expected return of all policies π in the original MDP is larger than the expected return in the new MDP: $R_M^\pi \geq R_{M_\mu}^\pi$. For any optimal π^* that satisfies Assumption 1 we have that $R_M^{\pi^*} = R_{M_\mu}^{\pi^*}$.*

That is, policy optimization in M_μ has at least one optimal solution identical to the original MDP M with the same policy value since M_μ lower bounds the policy value in M , so sub-optimal policies remain sub-optimal. Proof in Appendix A.

Off-Policy Policy Gradient in Augmented MDP We will now use the expected return in the modified MDP, $R_{M_\mu}^\pi$, as a surrogate for deriving policy gradients. According to the policy gradient theorem in Sutton et al. [2000], for a parametric policy π with parameters θ :

$$\frac{\partial R_{M_\mu}^\pi}{\partial \theta} = \mathbb{E}_{d^\pi} \left[\int_a \pi(a|s) \frac{\partial \log \pi(a|s)}{\partial \theta} Q_{M_\mu}^\pi(s, a) da \right],$$

From here on, $d^\pi(s)$ is with respect to the new MDP. The definition of $Q_{M_\mu}^\pi$ follows Sutton et al. [2000].¹

Now we will show that we can get an unbiased estimator of this gradient using importance sampling from the stationary state distribution $d^\mu(s)$ and the action distribution $\mu(a|s)$. According to the definition of M_μ , we have that for all s, a such that $d^\mu(s)\mu(a|s) = 0$, (s, a) is not in \mathcal{SA}_μ . Hence $Q_{M_\mu}^\pi(s, a) = 0$ for any policy π since (s, a) will receive zero reward and lead to a zero reward self-loop. So we have:

$$\begin{aligned} \frac{\partial R_{M_\mu}^\pi}{\partial \theta} &= \sum_s d^\pi(s) \sum_a \pi(a|s) \frac{\partial \log \pi(a|s)}{\partial \theta} Q_{M_\mu}^\pi(s, a) \\ &= \sum_{s: d^\mu(s) > 0} \frac{d^\pi(s)}{d^\mu(s)} d^\mu(s) \\ &\quad \sum_{a: \mu(a|s) > 0} \frac{\pi(a|s)}{\mu(a|s)} \mu(a|s) \frac{\partial \log \pi(a|s)}{\partial \theta} Q_{M_\mu}^\pi(s, a) \\ &= \mathbb{E}_{d^\mu, \mu} \left[\frac{d^\pi(s)}{d^\mu(s)} \frac{\pi(a|s)}{\mu(a|s)} \frac{\partial \log \pi(a|s)}{\partial \theta} Q_{M_\mu}^\pi(s, a) \right] \quad (2) \end{aligned}$$

Note that according to the definition of M_μ , the Markov chain induced by M and μ is exactly the same as M_μ and μ . Thus the distribution of (s_t, a_t, s_{t+1}) generated by executing μ in M is the same as executing μ in M_μ . So, we can estimate this policy gradient using the data we collected from μ in M . We conclude the section by pointing out that working in the augmented MDP allows us to construct a reasonable off-policy policy gradient estimator under the mild Assumption 1, while all prior works in this vein either explicitly or implicitly require the coverage of all possible policies.

Note that in the average reward case, such an augmented MDP would not be helpful for policy optimization since all policies that potentially reach s_{abs} will have a value of zero, and the stationary state distribution will be a single mass in the absorbing state. That would not induce a practical policy optimization algorithm. In the average reward case, either we need a stronger assumption that μ

¹For discounted case, our definition of expected return differs from the definition of Sutton et al. [2000] by a normalization factor $(1 - \gamma)$. This is because the definitions of stationary distributions are scaled differently in the two cases.

covers the entire state-action space or we must approximate the problem by setting a discount factor $\gamma < 1$ for the policy optimization algorithm, which is a common approach for deriving practical algorithms in an average reward (episodic) environment.

4 ALGORITHM: OPPOSD

Given the off-policy policy gradient derived in (2), how can we efficiently estimate it from samples collected from μ ? Notice that most quantities in the gradient estimator (2) are quite intuitive and also present in prior works such as Off-PAC. The main difference is the state distribution reweighting $d^\pi(s)/d^\mu(s)$, which we would like to estimate using samples collected with μ . For estimating this ratio of state distributions, we build on the recent work of Liu et al. [2018a] which we describe next.

For a policy π , let us define the shorthand $\rho_\pi(s, a) = \pi(s, a)/\mu(s, a)$. Further given a function $w : \mathcal{S} \rightarrow \mathbb{R}$, define $\Delta(w; s, a, s') := w(s)\rho_\pi(s, a) - w(s')$. Then we have the following result.

Theorem 2 ([Liu et al., 2018a]). *Given any $\gamma \in (0, 1)$, assume that $d^\mu(s) > 0$ for all s and define*

$$\begin{aligned} L(w, f) &= \gamma \mathbb{E}_{(s, a, s') \sim d^\mu} [\Delta(w; s, a, s') f(s')] \\ &\quad + (1 - \gamma) \mathbb{E}_{s \sim p_0} [(1 - w(s)) f(s)]. \end{aligned}$$

Then $w(s) = d^\pi(s)/d^\mu(s)$ if and only if $L(w, f) = 0$ for any measurable test function f .²

This result suggests a constructive procedure for estimating the state distribution ratio using samples from μ , by finding a function w over the states which minimizes $\max_f L(w, f)$. Since the maximization over all measurable functions as per Theorem 2 is intractable, Liu et al. [2018a] suggest restricting the maximization to a unit ball in an Reproducing Kernel Hilbert Space (RKHS), which has an analytical solution to the maximization problem, and we use the same procedure to approximate density ratios in our algorithm.

Applying Theorem 2 requires overcoming one final obstacle. The theorem presupposes $d^\mu(s) > 0$ for all s . In case where $\mathcal{SA}_\mu = \mathcal{S} \times \mathcal{A}$ we can directly apply the theorem. Otherwise in the MDP M_μ , this assumption indeed holds for all states, but μ never visits the absorbing state s_{abs} , or any transitions leading into this state. However, since we know this special state, as well as the dynamics leading in and out of it, we can simulate some samples for this state, effectively corresponding to a slight perturbation of μ to cover s_{abs} . Concretely, we first choose

²When $\gamma = 1$, w is only determined up to normalization, and hence an additional constraint $\mathbb{E}_{s \sim d^\mu} [w(s)] = 1$ is required to obtain the conclusion $w(s) = d^\pi(s)/d^\mu(s)$.

a small smoothing factor $\epsilon \in (0, 1)$. For any sample (s, a, s') in our data set, if there exist k actions \tilde{a} such that $\mu(\tilde{a}|s) = 0$, then we will keep the old samples with probability $1 - \epsilon$ and sample any one of the k actions with probability ϵ/k uniformly and change the next state s' to s_{abs} . If we sampled \tilde{a} , consequently, we would also change all samples after this transition to a self-loop in s_{abs} . Thus we create samples drawn according to a new behavior policy which covers all the state action pairs: $\tilde{\mu} = (1 - \epsilon)\mu + \epsilon U(s)$ where $U(s)$ is a uniform distribution over the k actions not chosen by μ in state s . Now we can use Theorem 2 and the algorithm from Liu et al. [2018a] to estimate $d^\pi(s)/d^{\tilde{\mu}}(s)$. Note that the propensity scores and policy gradients computed on this new dataset correspond to the behaviour policy $\tilde{\mu}$ and not μ . Formally, in place of using (2), we now estimate:

$$\mathbb{E}_{d^{\tilde{\mu}}, \tilde{\mu}} \left[\frac{d^\pi(s)}{d^{\tilde{\mu}}(s)} \frac{\pi(a|s)}{\tilde{\mu}(a|s)} \frac{\partial \log \pi(a|s)}{\partial \theta} Q_{M_\mu}^\pi(s, a) \right] \quad (3)$$

Note that we can estimate the expectation in (3) from the smoothed dataset by construction, since the ratio $\pi(s, a)/\tilde{\mu}(s, a)$ in all states are known.

Now that we have an algorithm for estimating policy gradients from (3), we can plug this into any policy gradient optimization method. Following prior work, we incorporate our off-policy gradients into an actor-critic algorithm. For learning the critic $Q_{M_\mu}^\pi(s, a)$, we can use any off-policy Temporal Difference [Bhatnagar et al., 2009, Maei, 2011] or Q-learning algorithm [Watkins and Dayan, 1992]. In our algorithm, we fit an approximate value function \hat{V} by:³

$$\hat{V}(s) \leftarrow \hat{V}(s) + \alpha_c \frac{\pi(a|s)}{\tilde{\mu}(a|s)} \left(R^\lambda(s, a) - \hat{V}(s) \right), \quad (4)$$

where α_c is the step-size for critic updates and $R^\lambda(s, a)$ is the off-policy λ -return:

$$R^\lambda(s, a) = r(s, a) + (1 - \lambda)\gamma\hat{V}(s') + \lambda\gamma\frac{\pi(a|s)}{\tilde{\mu}(a|s)}R^\lambda(s', a'),$$

and (s, a, s', a') is generated by executing $\tilde{\mu}$. After we learn \hat{V} , R^λ serves the role of $Q_{M_\mu}^\pi$ in our algorithm.

Given the estimates of the state distribution ratio from Liu et al. [2018a] and the critic updates from (4), we can now update the policy by plugging these quantities in (3). It remains to specify the initial conditions to start the algorithm. Since we have data collected from a behavior policy, it is natural to also warm-start the actor policy in our algorithm to be the same as the behavior policy and correspondingly the critic and w 's to be

the value function and distribution ratios for the behavior policy. This can be particularly useful in situations where the behavior policy, while suboptimal, still gets to states with high rewards with a reasonable probability. Hence we use behavior cloning to warm-start the policy parameters for the actor, use on-policy value function learning for the critic and also fit the state ratios w for the actor obtained by behavior cloning. Note that while the ratio will be identically equal to 1 if our behavior cloning is perfect, we actually estimate the ratio to better handle imperfections in the learned actor initialization.

Algorithm 1 OPPOSD: Off-Policy Policy Optimization with State Distribution Correction

Require: $\mathcal{S}, \mathcal{A}, \mu, \mathcal{D} : \{ \{s_t^i, a_t^i, r_t^i, \mu(a_t^i|s_t^i)\}_{t=0}^T \}_{i=0}^n$

Require: Hyperparameters $\lambda, \gamma, N_w, \alpha_w, N_c, \alpha_c, \alpha$

- 1: Warm start $\pi_\theta, \hat{V}_{\theta_c}, w_{\theta_w}$
 - 2: Pad \mathcal{D} to get \mathcal{D}' , $\tilde{\mu}$ if necessary
 - 3: **for** each step of policy update **do**
 - 4: **for** state ratio updates $i = 1, 2, \dots, N_w$ **do**
 - 5: Sample a mini-batch $B_w \sim \mathcal{D}'$ according to \hat{d}_γ ⁴
 - 6: **if** $\gamma = 1$ **then**
 - 7: Perform one update according to Algorithm 1 in Liu et al. [2018a] with stepsize α_w
 - 8: **else**
 - 9: Perform one update according to Algorithm 2 in Liu et al. [2018a] with stepsize α_w
 - 10: **end if**
 - 11: **end for**
 - 12: **for** critic updates $i = 1, 2, \dots, N_c$ **do**
 - 13: Sample a mini-batch $B_c \sim \mathcal{D}'$
 - 14: $\theta_c \leftarrow \theta_c - \alpha_c \frac{\partial \ell_c}{\partial \theta_c}$, where: $\ell_c = \frac{1}{|B_c|} \sum_{(s, a, s') \sim B_c} \frac{\pi(a|s)}{\tilde{\mu}(a|s)} \left(R^\lambda(s, a) - \hat{V}(s) \right)^2$
 - 15: **end for**
 - 16: Sample a mini-batch $B_a \sim \mathcal{D}'$ according to \hat{d}_γ
 - 17: $z_w \leftarrow \frac{1}{|B_a|} \sum_{s \sim B_a} w(s)$
 - 18: $Q^\pi(s, a) \leftarrow \mathbb{1}\{(s, a) \in \mathcal{S}\mathcal{A}_\mu\} R^\lambda(s, a)$
 - 19: $\theta \leftarrow \theta - \frac{\alpha}{|B_a|} \sum_{s \sim B_a} \frac{w(s)}{z_w} \rho(s, a) \frac{\partial \log \pi(a|s)}{\partial \theta} Q^\pi(s, a)$
 - 20: **end for**
-

A full pseudo-code of our algorithm, which we call OPPOSD for Off-Policy Policy Optimization with State Distribution Correction, is presented in Algorithm 1. We mention a couple of implementation details which we found helpful in improving the convergence properties of the algorithm. Typical actor-critic algorithms update the critic once per actor update in the on-policy setting. However, in the off-policy setting, we find that

³For simplicity, Eqn 4 views $\hat{V}(s)$ in the tabular setting. See Line 14 in Alg 1 for the function approximation case.

⁴ $\hat{d}_\gamma = \frac{1}{\sum_{t=0}^T \gamma^t} \sum_{t=0}^T \gamma^t \hat{d}_t(s)$, where $\hat{d}_t(s)$ is the empirical state distribution at time step t in dataset \mathcal{D}'

performing multiple critic updates before an actor update is helpful, since the off-policy TD learning procedure can have a high variance. Secondly, the computation of the state distribution ratio w is done in an online manner similar to the critic updates, and analogous to the critic, we always retain the state of the optimizer for w across the actor updates (rather than learning the w from scratch after each actor update). Similar to the critic, we also perform multiple w updates after each actor update. These choices are intuitively reasonable as the standard two-time scale asymptotic analysis of actor-critic methods [Borkar, 2009] does require the critic to converge faster than the actor.

5 CONVERGENCE RESULT

In this section, we present two main results to demonstrate the theoretical advantage of our algorithm. First we present a simple scenario where the prior approach of Off-PAC yields an arbitrarily biased gradient estimate, despite having access to a perfect critic. In contrast OPPOSD estimates the gradients correctly whenever the distribution ratios in (2) and the critic are estimated perfectly, by definition. We will further provide a convergence result for OPPOSD to a stationary point of the expected reward function.

A hard example for Off-PAC Many prior off-policy policy gradient methods use the policy gradient estimates proposed in Degris et al. [2012].

$$g_{\text{OffPAC}}(\theta) = \sum_s d^\mu(s) \sum_a \pi(a|s) \frac{\partial \log \pi(a|s)}{\partial \theta} Q^\pi(s, a)$$

Notice that, in contrast to the exact policy gradient, the expectation over states is taken with respect to the behavior policy μ distribution instead of π . In tabular settings this can lead to correct policy updates, as proved by Degris et al. [2012]. We now present an example where the policy gradient computed this way is problematic when using function approximators. Consider the problem instance shown in Figure 1, where the behavior policy π_b is uniformly random in any state. Now we consider policies parameterized by a parameter $\alpha \in [0, 1]$ where $\pi_\alpha(s_1, \ell) = \pi_\alpha(s_2, \ell) = \alpha$ and $\pi_\alpha(s, \ell) = 1$ for other states. Thus π_α aliases the states s_3 and s_4 as a manifestation of imperfect representation which is typical with large state spaces. Clearly the optimal policy is π_1 . With perfect critic, in Appendix C we show that for Off-PAC the gradient vanishes for any policy π_α , meaning that the algorithm can be arbitrarily sub-optimal at any point during policy optimization. Our gradient estimator (2) instead evaluates to $\partial R_{M_\mu}^\pi / \partial \alpha = (1 + \alpha)/4$, which is correctly maximized at $\alpha = 1$.




Figure 1: Hard example for Off-PAC [Degris et al., 2012]

Convergence results for OPPOSD We next ask whether OPPOSD converges, given reasonable estimates for the density ratio and the critic. To this end, we need to introduce some additional notations and assumptions. Suppose we run OPPOSD over some parametric policy class π_θ with $\theta \in \Theta$. In the sequel, we use subscripts and superscripts by θ to mean the corresponding quantities with π_θ to ease the notation. We begin by describing an abstract set of assumptions and a general result about the convergence of OPPOSD, when we run it over the policies π_θ given data collected with an arbitrary policy ν , before instantiating our assumptions for the specific structure of $\tilde{\mu}$ used in our algorithm.

Definition 1. A function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is L -smooth when

$$\|\nabla f(x) - \nabla f(y)\|_2 \leq L\|x - y\|_2, \quad \text{for all } x, y \in \mathbb{R}^d,$$

Assumption 2. $\forall (s, a)$ pairs, $\forall \theta \in \Theta$ and a data collection policy ν , we assume that the MDP guarantees:

1. $\left\| \frac{\partial \pi_\theta(a|s)}{\partial \theta} \right\| \leq G_{\max}$.
2. $Q^\theta(s, a) \leq V_{\max}$.
3. $\nu(a|s) \geq \nu_{\min}$.
4. $w(s) := d^\theta(s)/d^\nu(s)$ has a bounded second moment $\sigma_w^2 := E_\nu[w(s)^2] \leq \infty$.
5. The expected return of π_θ : R^θ is a differentiable, G -Lipschitz and L -smooth function w.r.t. θ .

Theorem 3. Assume an MDP, a data collection policy ν and function classes $\{\pi_\theta\}$ and $\{\hat{w}\}$ satisfy Assumption 2. Suppose OPPOSD with policy parameters θ_k at iteration k is provided critic estimates \hat{Q}_k and distribution ratio estimates \hat{w}_k satisfying $\mathbb{E}_{(s,a) \sim d^\nu} (w_{\theta_k}(s, a) - \hat{w}_k(s, a))^2 \leq \varepsilon_{w,k}^2$ and $\mathbb{E}_{(s,a) \sim d^\nu} (Q^{\theta_k}(s, a) - \hat{Q}_k(s, a))^2 \leq \varepsilon_{Q,k}^2$ for iterations $k = 1, 2, \dots, K$. Then

$$\frac{1}{K} \sum_{k=1}^K \mathbb{E} \left[\|\nabla_\theta R^{\theta_k}\|^2 \right] \leq \frac{2V_{\max}}{K} + \frac{\sum_{k=1}^K O((\varepsilon_{w,k}^2 V_{\max}^2 + \varepsilon_{Q,k}^2 (\sigma_w^2 + \varepsilon_{w,k}^2)) G_{\max}^2)}{K \nu_{\min}^2}$$

That is, when Assumption 2 holds, the scheme converges to an approximate stationary point given estimators \hat{w} and \hat{Q} with a small average MSE across the iterations under ν . An immediate consequence of the theorem above is that as long as we guarantee that $\lim_{K \rightarrow \infty} \frac{\sum \epsilon_{w,k}^2 + \epsilon_{Q,k}^2}{K} = 0$, which a reasonable online critic and w learning algorithm can guarantee, we have: $\lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K \mathbb{E} \left[\left\| \nabla_{\theta} R_{M_{\nu}}^{\theta_k} \right\|^2 \right] = 0$ which implies the procedure will converge to a stationary point where the true policy gradient is zero. Proof of this theorem in Appendix B.

We now discuss the validity of Assumption 2 in the specific context of the data collection policy $\tilde{\mu}$ used in OPPOSD as well as the augmented MDP M_{μ} . The first assumption, that the gradient of policy distribution is bounded, can be achieved by an appropriate policy parametrization such as a linear or a differentiable neural network-based scoring function composed with a softmax link. The second assumption on bounded value functions is standard in the literature. In particular, both these assumptions are crucial for the convergence of policy gradient methods even in an on-policy setting. The third assumption on lower bounded action probabilities holds by construction for the policy $\tilde{\mu}$ due to the ϵ -smoothing. The fourth assumption is about finite second momentum of distribution ratios. Prior works on off-policy evaluation in RL [Jiang and Li, 2016, Wang et al., 2017] both provide lower bounds for off-policy policy evaluation MSE in MDP and contextual bandit settings respectively that scale with the second moment of propensity scores, so this dependence is unavoidable for off-policy evaluation and learning. Finally the regularity assumption on the smoothness of the reward function is again standard for policy gradient methods even in an on-policy setting.

Thus we find that under standard assumptions for policy gradient methods, along with some reasonable additional conditions, we expect OPPOSD to have good convergence properties in theory.

6 EXPERIMENTAL EVALUATION

In this section we study the empirical properties of OPPOSD, with an eye towards two questions:

1. Does the state distribution correction help improve the performance of off-policy policy optimization?
2. Can we identify the best policy from the optimization path using off-policy policy evaluation?

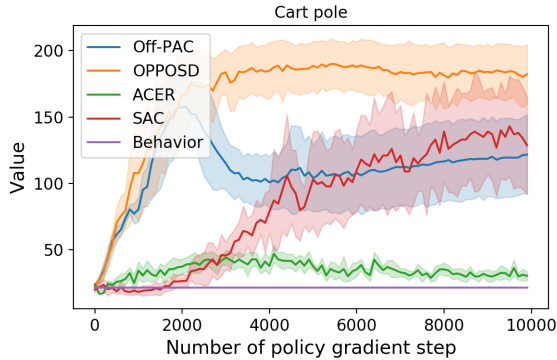
Baseline and implementation details To answer the first question, we compare OPPOSD with its closest prior

work, but without the state distribution correction, that is the Off-PAC algorithm [Degris et al., 2012].

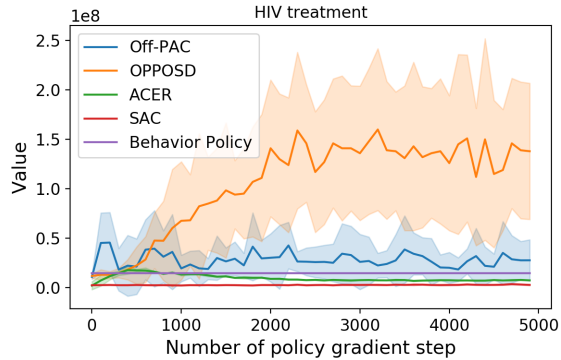
We implement both OPPOSD and Off-PAC using feed-forward neural networks for the actor and critic, with ReLU hidden layers. For state distribution ratio w , we also use a neural network with ReLU hidden layers, with the last activation function $\log(1 + \exp(x))$ to guarantee that $w(s) > 0$ for any input. To make a fair comparison, we keep the implementation of Off-PAC as close as possible to OPPOSD other than the use of w . Concretely, we also equip Off-PAC with the enhancements that improve empirical performance such as the warm start of the actor and critic, as well as several critic updates per actor update. We use the same off-policy critic learning algorithm for Off-PAC and OPPOSD. To learn w , we use Algorithm 1 (average reward) in Liu et al. [2018a] with RBF kernel for CartPole-v0 experiment, and Algorithm 2 (discounted reward) in Liu et al. [2018a] with RBF kernel for HIV experiment. We normalize the input to the networks to have 0 mean and 1 standard deviation, and in each mini-batch we normalized kernel loss of fitting w by the mean of the kernel matrix elements, to minimize the effect of kernel hyper-parameters on the learning rate. Full implementation details when omitted are provided in the Appendix D. We also compare OPPOSD with two more recent off-policy policy gradient algorithms, namely soft actor-critic (SAC) and actor-critic with experience replay (ACER).

Simulation domains We compare the algorithms in two simulation domains. The first domain is the *cart pole* control problem, where an agent needs to balance a mass attached to a pole in an upright position, by applying one of two sideways movements to a cart on a frictionless track. The horizon is fixed to 200. If the trajectory ends in less than 200 steps, we pad the episode by continuing to sample actions and repeating the last state. We use a uniformly random policy to collect $n = 500$ trajectories as off-policy data, which is a very challenging data set for off-policy policy optimization methods to learn from as this policy does not attain the desired upright configuration for any prolonged period of time.

The second domain is an *HIV treatment simulation* described in Ernst et al. [2006]. Here the states are six-dimensional real-valued vectors, which model the response of numbers of cells/virus to a treatment. Each action corresponds to whether or not to apply two types of drug, leading to a total of 4 actions. The horizon of this domain is 200 and discount factor is set by the simulator to $\gamma = 0.98$. A uniformly random policy does not visit any rewarding states often enough to collect useful data for off-policy learning. To simulate an imperfect but reasonable data collection policy, we first train an on-



(a) CartPole-v0. Optimal policy gets a reward of 200, while the uniformly random data collection policy obtains 22.



(b) HIV treatment simulator. The data collection policy obtains a reward of 1.5E7.

Figure 2: Episodic scores over length 200 episodes in CartPole-v0 [Barto et al., 1983, Brockman et al., 2016] and HIV treatment simulator Ernst et al. [2006]. Shaded region represents 1 standard deviation over 10 runs of each method.

policy actor critic method to learn a reasonable (but still far from optimal) policy $\hat{\pi}$. We then use the data collection policy $\mu = 0.7*\hat{\pi} + 0.3*U$, where U is the uniformly random policy, to collect $n = 1000$ trajectories.

Though in both domains our data collection policy is eventually able to cover the whole state-action space, the situation under finite amounts of data is different. In cart pole since an optimal policy can control the cart to stay in a small region, it is relatively easy for the uniform random policy to cover the states visited by the optimal policy. In the HIV treatment domain, it is unlikely that the logged data will cover the desirable state space. For the horizon of problem, theoretical derivation of our algorithm requires infinite horizon, but for computational feasibility, the horizon is truncated as 200 steps in the both domains.

Impact of state reweighting on policy optimization

In Figures 2a and 2b, we plot the on-policy evaluation values of the policies produced by OPPOSD and Off-PAC during the actor updates across 10 runs. Each run uses a different data set collected using the behavior policy as well as a different random seed for the policy optimization procedure. In each run we use the same dataset for each method to allow paired comparisons. We evaluate the policy after every 100 actor updates using on-policy Monte-Carlo evaluation over 20 trajectories. The results are averaged over 10 runs and error bars show the standard deviation. Along X-axis, the plot shows how the policy value changes as we take policy gradient steps.

At a high-level, we see that in both the domains our algorithm significantly improves upon the behavior policy, and eventually outperforms Off-PAC consistently. Zooming in a bit, we see that for the initial iterates on the

left of the plots, the gap between OPPOSD and Off-PAC is small as the state distribution between the learned policies is likely close enough to the behavior policy for the distribution mismatch to not matter significantly. This effect is particularly pronounced in Figure 2a. However, the gap quickly widens as we move to the right in both the figures. In particular, Off-PAC barely improves over behavior policy in Figure 2b, while OPPOSD finds a significantly better policy. Overall, we find that these results are an encouraging validation of our intuition about the importance of correcting the state distribution mismatch.

As for comparison with SAC and ACER, we would like to emphasize that we run SAC and ACER with a fixed batch of data which is not exactly the same setting as they were proposed for. In Appendix E we discuss more details on the difference of settings in this paper and SAC and ACER originally. As an illustration of the effect of this difference on algorithm performance, with the same function approximator class, both SAC and ACER perform worse than OPPOSD in our experimental domains. Since SAC and ACER are not proposed for the batch off-policy settings and this simulation does not reflect these algorithms abilities in their proposed settings.

Identifying Best Policy by Off-Policy Evaluation

While OPPOSD consistently outperforms Off-PAC in average performance across 10 runs in both the domains, there is still significant variance in both the methods across runs. Given this variance, a natural question is whether we can identify the best performing policies, during and across multiple runs of OPPOSD for a single dataset. To answer this question, we checkpoint all the policies produced by OPPOSD after every 1000 actor updates, across 5 runs of our algorithm with the same input dataset generated in the HIV domain. We then eval-

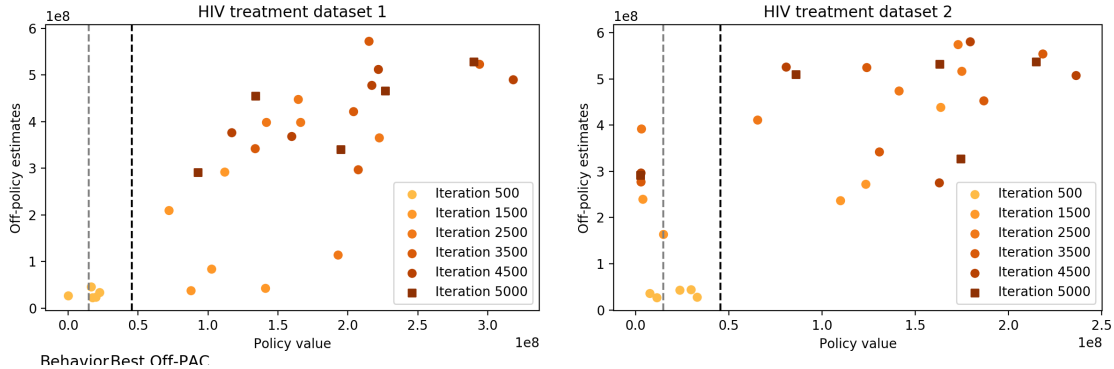


Figure 3: Off-policy policy evaluation results of saved policies from OPPOSD. The estimated and true values exhibit a high correlation (coefficient = 0.80 and 0.71 in the left and right plots respectively) for most policies. Two panels correspond to repeating the whole procedure using two datasets from the same behavior policy.

uate these policies using the off-policy policy evaluation (OPPE) method in Liu et al. [2018a]. The evaluation is performed with an additional dataset sampled from the behavior policy. This corresponds to the typical practice of sample splitting between optimization and evaluation.

We show the quality of the OPPE estimates against the true policy values for two different datasets for OPPE sampled from the behavior policy in the two panels of Figure 3. In each plot, the X-axis shows the on-policy Monte-Carlo evaluation results and Y-axis shows the OPPE estimates. We find that the OPPE estimates are generally well correlated with the on-policy values, and picking the policy with the best OPPE estimate results in a true value substantially better than both the best Off-PAC result. A closer inspection also reveals the importance of this validation step. The red squares correspond to the final iterate of OPPOSD in each of the 5 iterations, which has a very high value in some cases, but somewhat worse in other runs. Using OPPE to robustly select a good policy adds a layer of additional assurance to off-policy policy optimization procedure.

7 DISCUSSION

We presented a new off-policy actor critic algorithm, OPPOSD, based on a recently proposed stationary state distribution ratio estimator. There exist many interesting next steps, including different critic learning methods which may also leverage the state distribution ratio, and exploring alternative methods for policy evaluation or alternative stationary state distribution ratio estimators [Hallak and Mannor, 2017, Gelada and Bellemare, 2019]. Another interesting direction is to improve the sample efficiency of online policy gradient algorithms by using our corrected gradient estimates. More discussion on concurrent potential future work is in Appendix E.

In parallel with our work, Zhang et al. [2019] have presented a different off-policy policy gradient approach. While similarly motivated by the bias in the Off-PAC gradient estimator, the two works have important differences. On the methodological side, Zhang et al. [2019] start from an off-policy objective function and derive a gradient for it. In contrast, we compute an off-policy estimator for the gradient of the on-policy objective function. The latter leads to a much simpler method, both conceptually and computationally, as we do not need to compute the gradients of the visitation distribution. On the other hand, Zhang et al. [2019] focus on incorporating more general interest functions in the off-policy objective, and use the emphatic weighting machinery. In terms of settings, our approach works in the offline setting (though easily extended to online), while they require an online setting in order to compute the gradients of the propensity score function. Finally, we present convergence results quantifying the error in our critic and propensity score computations while Zhang et al. [2019] assume a perfect oracle for both and rely on a truly unbiased gradient estimator for the convergence results.

To conclude, our algorithm fixes the bias in off-policy policy gradient estimates introduced by the behavior policy’s stationary state distribution. We prove under certain assumptions our algorithm is guaranteed to converge. We also show that ignoring the bias due to the mismatch in state distributions can make an off policy gradient algorithm fail even in a simple illustrative example, and that by accounting for this mismatch our approach yields significantly better performance in two simulation domains.

Acknowledgements

We acknowledge a NSF CAREER award, an ONR Young Investigator Award, and support from Siemens.

References

- Pieter Abbeel and John Schulman. Deep reinforcement learning through policy optimization. *NeurIPS Tutorial*, 2016. URL <https://nips.cc/Conferences/2016/Schedule?showEvent=6198>.
- András Antos, Csaba Szepesvári, and Rémi Munos. Learning near-optimal policies with bellman-residual minimization based fitted policy iteration and a single sample path. *Machine Learning*, 71(1):89–129, 2008.
- Andrew G Barto, Richard S Sutton, and Charles W Anderson. Neuronlike adaptive elements that can solve difficult learning control problems. *IEEE transactions on systems, man, and cybernetics*, pages 834–846, 1983.
- Shalabh Bhatnagar, Doina Precup, David Silver, Richard S Sutton, Hamid R Maei, and Csaba Szepesvári. Convergent temporal-difference learning with arbitrary smooth function approximation. In *Advances in Neural Information Processing Systems*, pages 1204–1212, 2009.
- Vivek S Borkar. *Stochastic approximation: a dynamical systems viewpoint*, volume 48. Springer, 2009.
- Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym, 2016.
- Bo Dai, Albert Shaw, Lihong Li, Lin Xiao, Niao He, Zhen Liu, Jianshu Chen, and Le Song. Sbeed: Convergent reinforcement learning with nonlinear function approximation. In *International Conference on Machine Learning*, pages 1133–1142, 2018.
- Thomas Degris, Martha White, and Richard S Sutton. Off-policy actor-critic. *arXiv preprint arXiv:1205.4839*, 2012.
- Damien Ernst, Pierre Geurts, and Louis Wehenkel. Tree-based batch mode reinforcement learning. *Journal of Machine Learning Research*, 6(Apr):503–556, 2005.
- Damien Ernst, Guy-Bart Stan, Jorge Goncalves, and Louis Wehenkel. Clinical data based optimal sti strategies for hiv: a reinforcement learning approach. In *Proceedings of the 45th IEEE Conference on Decision and Control*, pages 667–672. IEEE, 2006.
- Mehrdad Farajtabar, Yinlam Chow, and Mohammad Ghavamzadeh. More robust doubly robust off-policy evaluation. In *Proceedings of the 35th International Conference on Machine Learning*, pages 1446–1455, 2018.
- Carles Gelada and Marc G Bellemare. Off-policy deep reinforcement learning by bootstrapping the covariate shift. In *AAAI*, 2019.
- Saeed Ghadimi and Guanghui Lan. Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23(4):2341–2368, 2013.
- Shixiang Shane Gu, Timothy Lillicrap, Zoubin Ghahramani, Richard E. Turner, and Sergey Levine. Q-prop: Sample-efficient policy gradient with an off-policy critic. In *International Conference on Learning Representations*, 2017a.
- Shixiang Shane Gu, Timothy Lillicrap, Richard E Turner, Zoubin Ghahramani, Bernhard Schölkopf, and Sergey Levine. Interpolated policy gradient: Merging on-policy and off-policy gradient estimation for deep reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 3846–3855, 2017b.
- Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, et al. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*, 2018.
- Assaf Hallak and Shie Mannor. Consistent on-line off-policy evaluation. In *Proceedings of the 34th International Conference on Machine Learning*, pages 1372–1383, 2017.
- Ehsan Imani, Eric Graves, and Martha White. An off-policy policy gradient theorem using emphatic weightings. In *Advances in Neural Information Processing Systems*, pages 96–106, 2018.
- Nan Jiang and Lihong Li. Doubly robust off-policy value evaluation for reinforcement learning. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning-Volume 48*, pages 652–661, 2016.
- 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.
- Qiang Liu, Lihong Li, Ziyang Tang, and Dengyong Zhou. Breaking the curse of horizon: Infinite-horizon off-policy estimation. In *Advances in Neural Information Processing Systems*, pages 5361–5371, 2018a.
- Yao Liu, Omer Gottesman, Aniruddh Raghu, Matthieu Komorowski, Aldo A Faisal, Finale Doshi-Velez, and Emma Brunskill. Representation balancing mdps for off-policy policy evaluation. In *Advances in Neural Information Processing Systems*, pages 2649–2658, 2018b.
- Hamid Reza Maei. *Gradient Temporal-difference Learning Algorithms*. PhD thesis, University of Alberta, 2011.

- Alberto Maria Metelli, Matteo Papini, Francesco Faccio, and Marcello Restelli. Policy optimization via importance sampling. In *Advances in Neural Information Processing Systems*, pages 5442–5454, 2018.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529, 2015.
- Doina Precup, Richard S Sutton, and Satinder P Singh. Eligibility traces for off-policy policy evaluation. In *Proceedings of the 17th International Conference on Machine Learning*, pages 759–766, 2000.
- John Schulman, Sergey Levine, Philipp Moritz, Michael Jordan, and Pieter Abbeel. Trust region policy optimization. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning*, pages 1889–1897, 2015.
- David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. Deterministic policy gradient algorithms. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning*, 2014.
- Richard S Sutton, David A McAllester, Satinder P Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems*, pages 1057–1063, 2000.
- Philip Thomas and Emma Brunskill. Data-efficient off-policy policy evaluation for reinforcement learning. In *International Conference on Machine Learning*, pages 2139–2148, 2016.
- Philip Thomas, Georgios Theodorou, and Mohammad Ghavamzadeh. High confidence policy improvement. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning*, pages 2380–2388, 2015.
- Yu-Xiang Wang, Alekh Agarwal, and Miroslav Dudik. Optimal and adaptive off-policy evaluation in contextual bandits. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3589–3597. JMLR. org, 2017.
- Ziyu Wang, Victor Bapst, Nicolas Heess, Volodymyr Mnih, Remi Munos, Koray Kavukcuoglu, and Nando de Freitas. Sample efficient actor-critic with experience replay. *arXiv preprint arXiv:1611.01224*, 2016.
- Christopher JCH Watkins and Peter Dayan. Q-learning. *Machine learning*, 8(3-4):279–292, 1992.
- Shangdong Zhang, Wendelin Boehmer, and Shimon Whiteson. Generalized off-policy actor-critic. *arXiv preprint arXiv:1903.11329*, 2019.