
A Deep Reinforcement Learning Approach to Marginalized Importance Sampling with the Successor Representation

Scott Fujimoto¹ David Meger¹ Doina Precup¹

Abstract

Marginalized importance sampling (MIS), which measures the density ratio between the state-action occupancy of a target policy and that of a sampling distribution, is a promising approach for off-policy evaluation. However, current state-of-the-art MIS methods rely on complex optimization tricks and succeed mostly on simple toy problems. We bridge the gap between MIS and deep reinforcement learning by observing that the density ratio can be computed from the successor representation of the target policy. The successor representation can be trained through deep reinforcement learning methodology and decouples the reward optimization from the dynamics of the environment, making the resulting algorithm stable and applicable to high-dimensional domains. We evaluate the empirical performance of our approach on a variety of challenging Atari and MuJoCo environments.

1. Introduction

Off-policy evaluation (OPE) is a reinforcement learning (RL) task where the aim is to measure the performance of a target policy from data collected by a separate behavior policy (Sutton & Barto, 1998). As it can often be difficult or costly to obtain new data, OPE offers an avenue for re-using previously gathered data, making OPE an important challenge for applying RL to real-world domains (Zhao et al., 2009; Mandel et al., 2014; Swaminathan et al., 2017; Gauci et al., 2018).

Marginalized importance sampling (MIS) (Liu et al., 2018; Xie et al., 2019; Nachum et al., 2019a) is a family of OPE methods which re-weight sampled rewards by directly learning the density ratio between the state-action occupancy of the target policy and the sampling distribution. This

approach can have significantly lower variance than traditional importance sampling methods (Precup et al., 2001), which consider a product of ratios over trajectories, and is amenable to deterministic policies and behavior agnostic settings where the sampling distribution is unknown. However, the body of MIS work is largely theoretical, and as a result, empirical evaluations of MIS have mostly been carried out on simple low-dimensional tasks, such as mountain car (state dim. of 2) or cartpole (state dim. of 4). In comparison, deep RL algorithms have shown successful behaviors in high-dimensional domains such as Humanoid locomotion (state dim. of 376) and Atari (image-based).

In this paper, we present a straightforward approach for MIS that can be computed from the successor representation (SR) (Dayan, 1993) of the target policy by directly optimizing the reward function. Our algorithm, the Successor Representation Distribution Correction Estimation (SR-DICE), is the first method that allows MIS to scale to high-dimensional systems, far outperforming previous approaches. In comparison to previous algorithms which rely on minimax optimization or kernel methods (Liu et al., 2018; Nachum et al., 2019a; Uehara & Jiang, 2019; Mousavi et al., 2020; Yang et al., 2020), SR-DICE requires only a simple convex loss applied to a linear function, after computing the SR. Similar to the deep RL methods which can learn in high-dimensional domains, the SR can be computed easily using behavior-agnostic temporal-difference (TD) methods. This makes our algorithm highly amenable to deep learning architectures and applicable to complex tasks.

The SR, which measures the expected future occupancy of states for a given policy, has a clear relationship to MIS methods, which estimate the ratio between the occupancy of state-action pairs and the sampling distribution. However, this relationship is muddled in a deep RL context, where the deep SR measures the expected future sum of feature vectors. Our approach, SR-DICE, provides a straightforward and principled method for extracting density ratios from the SR without any modifications to the standard learning procedure of the SR. Access to these density ratios is valuable as they have a wide range of possible applications such as policy regularization (Nachum et al., 2019b; Touati et al., 2020), imitation learning (Kostrikov et al., 2019), off-policy

¹Mila, McGill University. Correspondence to: Scott Fujimoto <scott.fujimoto@mail.mcgill.ca>.

policy gradients (Imani et al., 2018; Liu et al., 2019b; Zhang et al., 2019), non-uniform sampling procedures (Sinha et al., 2020), or for mitigating distributional shift in offline RL (Fujimoto et al., 2019; Kumar et al., 2019).

We highlight the value of the MIS density ratios for one reason in particular—in our theoretical analysis we prove that SR-DICE and the deep SR produce *exactly the same value estimate*. This is surprising as SR-DICE takes a distinct approach for value estimation by re-weighting every reward in the dataset with an importance sampling ratio while the deep SR estimates the value in a similar fashion to TD learning. This theoretical result extends to the deep RL setting and is consistent in our experimental results. This result is a double-edged sword which (negatively) implies there is no discernible difference of using our MIS approach for policy evaluation, but (positively) implies the estimated density ratios are accurate enough to match the performance of TD methods. This is an important observation as our empirical results demonstrate that previous MIS methods scale very poorly in comparison to TD methods to high dimensions, which is consistent with prior results (Voloshin et al., 2019; Fu et al., 2021). Even if there is no difference for OPE, a MIS method which matches the performance of TD-based methods is desirable if we are concerned with estimating the density ratios of the target policy.

We benchmark the performance of SR-DICE on several high-dimensional domains in MuJoCo (Todorov et al., 2012) and Atari (Bellemare et al., 2013), against several recent MIS methods (Nachum et al., 2019a; Zhang et al., 2020a). Our results demonstrate several key findings regarding high-dimensional tasks.

Current MIS methods underperform deep RL at high-dimensional tasks. While previous results have shown that MIS methods can produce competitive results to TD methods, our empirical results show that MIS methods scale poorly to challenging tasks. In Atari we find that the baseline MIS method exhibit unstable estimates, often reaching errors with many orders of magnitude. Comparatively, the baseline deep RL methods, which rely on TD learning and have a history of achieving high performances in the control setting (Mnih et al., 2015; Schulman et al., 2017; Fujimoto et al., 2018), outperform the MIS baselines at every task and often by a wide margin.

SR-DICE outperforms current MIS methods at policy evaluation and therefore density ratio estimation. Our empirical results confirm our theoretical analysis, which state that SR-DICE and the standard deep SR approach should produce identical value estimates (with differences due only to changes to the optimization process). While this result may initially sound discouraging, given the direct SR approach is comparable to TD learning, and TD learning significantly outperforms current MIS methods, this result

also implies that SR-DICE is a much stronger technique for estimating density ratios than previous methods.

Ultimately, while SR-DICE produces a similar result to existing deep RL approaches for policy evaluation, it does provide a practical, scalable, and state-of-the-art approach for estimating state-action occupancy density ratios, while highlighting connections between the SR, reward function optimization, and state-action occupancy estimation. For ease of use and reproduction, our code is open-sourced (<https://github.com/sfujim/SR-DICE>).

2. Background

Reinforcement Learning. RL is a framework for maximizing accumulated reward of an agent interacting with its environment (Sutton & Barto, 1998). This problem is typically framed as a Markov Decision Process (MDP) $(\mathcal{S}, \mathcal{A}, \mathcal{R}, p, d_0, \gamma)$, with state space \mathcal{S} , action space \mathcal{A} , reward function \mathcal{R} , dynamics model p , initial state distribution d_0 and discount factor γ . An agent selects actions according to a policy $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$. In this paper we address the problem of off-policy evaluation (OPE) problem where the aim is to measure the normalized expected per-step reward of the policy $R(\pi) = (1 - \gamma)\mathbb{E}_\pi[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t)]$. An important notion in OPE is the value function $Q^\pi(s, a) = \mathbb{E}_\pi[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) | s_0 = s, a_0 = a]$, which measures the expected sum of discounted rewards when following π , starting from the state-action pair (s, a) .

We define $d^\pi(s, a)$ as the discounted state-action occupancy, the probability of seeing (s, a) under policy π with discount γ : $d^\pi(s, a) = (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t \int_{s_0} d_0(s_0) p_\pi(s_0 \rightarrow s, t) \pi(a|s) ds_0$, where $p_\pi(s_0 \rightarrow s, t)$ is the probability of arriving at the state s after t time steps when starting from an initial state s_0 . This distribution is important as $R(\pi)$ equals the expected reward $r(s, a)$ under d^π :

$$R(\pi) = \mathbb{E}_{(s,a) \sim d^\pi, r(s,a)}[r(s, a)]. \quad (1)$$

A common approach for estimating $R(\pi)$ is through temporal-difference (TD) learning (Sutton, 1988) where an estimate of the value function $Q(s, a)$ is updated over individual transitions $(s, a, r(s, a), s')$ by the following:

$$Q(s, a) \leftarrow \alpha(r(s, a) + \gamma Q(s', a')) + (1 - \alpha)Q(s, a), \quad (2)$$

where a' is sampled according to the target policy π and α is the learning rate. Provided an infinite set of transitions, TD learning is known to converge to the true value function in the off-policy setting (Jaakkola et al., 1994; Sutton & Barto, 1998). TD learning can also be applied to other learning problems, such as the successor representation, where the reward $r(s, a)$ in Equation (2) is replaced with the quantity of interest.

Successor Representation. The successor representation (SR) (Dayan, 1993) of a policy is a measure of occupancy of future states. It can be viewed as a general value function that learns a vector of the expected discounted visitation for each state. The SR Ψ^π of a given policy π is defined as $\Psi^\pi(s'|s) = \mathbb{E}_\pi[\sum_{t=0}^{\infty} \gamma^t \mathbb{1}(s_t = s') | s_0 = s]$. Importantly, the value function can be recovered from the SR by summing over the expected reward of each state $V^\pi(s) = \sum_{s'} \Psi^\pi(s'|s) \mathbb{E}_{a' \sim \pi}[r(s', a')]$. For infinite state and action spaces, the SR can instead be generalized to the expected occupancy over features, known as the deep SR (Kulkarni et al., 2016) or successor features (Barreto et al., 2017). For a given encoding function $\phi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^n$, the deep SR $\psi^\pi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^n$ is defined as the expected discounted sum of features from the encoding function ϕ when following the policy from a given state-action pair:

$$\psi^\pi(s, a) = \mathbb{E}_\pi \left[\sum_{t=0}^{\infty} \gamma^t \phi(s_t, a_t) \mid s_0 = s, a_0 = a \right]. \quad (3)$$

If the encoding $\phi(s, a)$ is learned such that the original reward function is a linear function of the encoding $r(s, a) = \mathbf{w}^\top \phi(s, a)$, then similar to the original formulation of SR, the value function can be recovered from a linear function of the SR: $Q^\pi(s, a) = \mathbf{w}^\top \psi^\pi(s, a)$. The deep SR network ψ^π is trained to minimize the MSE between $\psi^\pi(s, a)$ and $\phi(s, a) + \gamma \psi'(s', a')$ on transitions (s, a, s') sampled from the dataset. A frozen target network ψ' is used to provide stability (Mnih et al., 2015; Kulkarni et al., 2016), and is updated to the current network $\psi' \leftarrow \psi^\pi$ after a fixed number of time steps. The encoding function ϕ is typically trained by an encoder-decoder network (Kulkarni et al., 2016; Machado et al., 2017; 2018a). For OPE where the reward function is learned by minimizing $(\mathbf{w}^\top \phi(s, a) - r(s, a))^2$, the SR is comparable to TD learning, as they both estimate the discounted sum of future rewards and use similar updates.

Marginalized Importance Sampling. Marginalized importance sampling (MIS) is a family of importance sampling approaches for off-policy evaluation in which the performance $R(\pi)$ is evaluated by re-weighting rewards sampled from a dataset $\mathcal{D} = \{(s, a, r, s')\} \sim p(s'|s, a) d^{\mathcal{D}}(s, a)$, where $d^{\mathcal{D}}$ is an arbitrary distribution, typically but not necessarily, induced by some behavior policy. It follows that $R(\pi)$ can be computed with importance sampling weights on the rewards $\frac{d^\pi(s, a)}{d^{\mathcal{D}}(s, a)}$:

$$R(\pi) = \mathbb{E}_{(s, a) \sim d^{\mathcal{D}}, r(s, a)} \left[\frac{d^\pi(s, a)}{d^{\mathcal{D}}(s, a)} r(s, a) \right]. \quad (4)$$

The goal of marginalized importance sampling methods is to learn the weights $w(s, a) \approx \frac{d^\pi(s, a)}{d^{\mathcal{D}}(s, a)}$, using data contained in \mathcal{D} . The main benefit of MIS is that unlike traditional importance methods, the ratios are applied to individual

transitions rather than complete trajectories, which can reduce the variance of long or infinite horizon problems. In other cases, the ratios themselves can be used for a variety of applications which require estimating the occupancy of state-action pairs.

3. A Reward Function Perspective on Distribution Corrections

In this section, we present our behavior-agnostic approach to estimating MIS ratios, called the Successor Representation Distribution Correction Estimation (SR-DICE). Our main insight is that MIS can be viewed as an optimization over a learned reward function, where the loss is uniquely optimized when the virtual reward is the MIS density ratio.

Our derived loss function is a straightforward convex loss over the learned reward and the corresponding value function of the target policy. This naturally suggests the use of the successor representation which allows us to maintain an estimate of the value estimate while directly optimizing the reward function. This disentangles the learning process, where the propagation of reward through the MDP can be learned separately from the optimization of the reward. In other words, rather than learn a reward function and value function simultaneously, we tackle each separately, changing the difficult minimax optimization of previous methods into two phases. Interestingly enough, we show that our MIS estimator produces the identical value estimate as traditional deep SR methods. This means the challenging aspect of learning has been pushed onto the computation of the SR, rather than optimizing the density ratio estimate. Fortunately, we can leverage deep RL approaches (Mnih et al., 2015; Kulkarni et al., 2016) to make learning the SR stable, giving rise to a practical MIS method for high-dimensional tasks.

This section begins with the derivation of our core ideas, which shows MIS ratios can be learned through reward function optimization. We then highlight how the SR can be used for reward function optimization in the tabular domain. Finally, we generalize our results to the deep SR setting.

3.1. Basic Derivation

In MIS, our aim is to determine the MIS ratios $\frac{d^\pi(s, a)}{d^{\mathcal{D}}(s, a)}$, using only data sampled from the dataset \mathcal{D} and the target policy π . This presents a challenge as we have direct access to neither d^π nor $d^{\mathcal{D}}$.

As a starting point, we begin by following the derivation of DualDICE (Nachum et al., 2019a). We first consider the convex function $\frac{1}{2}mx^2 - nx$, which is uniquely minimized by $x^* = \frac{n}{m}$. Now by replacing x with a virtual reward $\hat{r}(s, a)$, m with the density of the dataset $d^{\mathcal{D}}(s, a)$, and

n with the density of the target policy $d^\pi(s, a)$, we have reformulated the convex function as the following:

$$\min_{\hat{r}(s,a)\forall(s,a)} J(\hat{r}) := \frac{1}{2} \mathbb{E}_{(s,a)\sim d^{\mathcal{D}}} [\hat{r}(s, a)^2] - (1 - \gamma) \mathbb{E}_{(s,a)\sim d^\pi} [\hat{r}(s, a)]. \quad (5)$$

As Equation (5) is still the convex function with renamed variables, following Nachum et al. (2019a), we can observe the following:

Observation 1 *The objective $J(\hat{r})$ is minimized when $\hat{r}(s, a) = \frac{d^\pi(s,a)}{d^{\mathcal{D}}(s,a)}$ for all state-action pairs (s, a) .*

Equation (5) is an optimization over two expectations over $d^{\mathcal{D}}$ and d^π . While the first expectation over $d^{\mathcal{D}}$ is tractable by sampling directly from the dataset \mathcal{D} , the second expectation relies on the state-action visitation of the target policy $d^\pi(s, a)$ which is not directly accessible without a model of the MDP. At this point, we highlight our choice of notation, $\hat{r}(s, a)$, in Equation (5). Describing the objective in terms of a fictitious reward \hat{r} will allow us to draw on familiar relationships between rewards and value functions. Consider the equivalence between the value function over initial state-action pairs (s_0, a_0) and the expectation of rewards over the state-action visitation of the policy $(1 - \gamma) \mathbb{E}_{s_0, a_0 \sim \pi} [Q^\pi(s_0, a_0)] = \mathbb{E}_{d^\pi} [r(s, a)]$. It follows that the expectation over d^π in Equation (5) can be replaced with a value function \hat{Q}^π over \hat{r} :

$$\min_{\hat{r}(s,a)\forall(s,a)} J(\hat{r}) := \frac{1}{2} \mathbb{E}_{(s,a)\sim d^{\mathcal{D}}} [\hat{r}(s, a)^2] - (1 - \gamma) \mathbb{E}_{s_0, a_0 \sim \pi(\cdot|s_0)} [\hat{Q}^\pi(s_0, a_0)]. \quad (6)$$

In other words, by noting that the value function is simply the (scaled) expected reward when sampled from the state-action visitation of the target policy, we can replace the impractical expectation over d^π with a tractable value function. This form of the objective, Equation (6), is convenient because we can estimate the expectation over $d^{\mathcal{D}}$ by sampling directly from the dataset and \hat{Q}^π can be computed using any policy evaluation method.

While we can estimate both terms in Equation (6) with relative ease, the optimization problem is not directly differentiable and would require re-learning the value function \hat{Q}^π with every adjustment to the learned reward \hat{r} . Fortunately, there exists a straightforward paradigm which enables direct reward function optimization known as successor representation (SR).

3.2. Tabular SR-DICE

We will begin by discussing how we can apply the SR to MIS in the tabular setting and then generalize our method to non-linear function approximation afterwards. Consider

the relationship between the SR Ψ^π of the target policy π and its value function:

$$\mathbb{E}_{s_0, a_0 \sim \pi(\cdot|s_0)} [Q^\pi(s_0, a_0)] = \mathbb{E}_{s_0} [V^\pi(s_0)] = \mathbb{E}_{s_0} \left[\sum_s \Psi^\pi(s|s_0) \mathbb{E}_{a \sim \pi} [r(s, a)] \right]. \quad (7)$$

It follows that we can create an optimization problem directly over the reward function \hat{r} by modifying Equation (6) to use the SR:

$$\min_{\hat{r}(s,a)\forall(s,a)} J_\Psi(\hat{r}) := \frac{1}{2} \mathbb{E}_{(s,a)\sim d^{\mathcal{D}}} [\hat{r}(s, a)^2] - (1 - \gamma) \mathbb{E}_{s_0} \left[\sum_s \Psi^\pi(s|s_0) \mathbb{E}_{a \sim \pi} [\hat{r}(s, a)] \right]. \quad (8)$$

Since this optimization problem is convex, it has a closed form solution. The unique optimizer of Equation (8) is:

$$(1 - \gamma) \frac{|\mathcal{D}|}{\sum_{(s',a')\in\mathcal{D}} \mathbb{1}(s' = s, a' = a)} \cdot \mathbb{E}_{s_0} [\pi(a|s) \Psi^\pi(s|s_0)]. \quad (9)$$

By noting the relationship between the SR and the state occupancy $d^\pi(s, a) = (1 - \gamma) \mathbb{E}_{s_0} [\Psi^\pi(s|s_0) \pi(s, a)]$ and the fact that $d^{\mathcal{D}}(s, a) = \frac{\sum_{(s',a')\in\mathcal{D}} \mathbb{1}(s' = s, a' = a)}{|\mathcal{D}|}$ we can show this solution simplifies to the MIS density ratio $\frac{d^\pi(s,a)}{d^{\mathcal{D}}(s,a)}$.

Theorem 1 *Equation (9) is the optimal solution to Equation (8) and is equal to $\frac{d^\pi(s,a)}{d^{\mathcal{D}}(s,a)}$.*

A direct consequence of this result is that Equation (9) can be used with MIS policy evaluation to return the true value estimate $\frac{1}{|\mathcal{D}|} \sum_{(s,a)\in\mathcal{D}} \frac{d^\pi(s,a)}{d^{\mathcal{D}}(s,a)} r(s, a) = R(\pi)$.

Unfortunately, the form of Equation (9) relies on the true SR Ψ^π , as well as an expectation over s_0 , both of which may be unobtainable in the setting where we are sampling from a finite dataset \mathcal{D} . However, we can still show that with an inexact SR $\hat{\Psi}$ and sampled estimate of the expectation, using the set of start states \mathcal{D}_0 in the dataset, approximating the optimizer Equation (9) with

$$r^*(s, a) = (1 - \gamma) \frac{|\mathcal{D}|}{\sum_{(s',a')\in\mathcal{D}} \mathbb{1}(s' = s, a' = a)} \cdot \frac{1}{|\mathcal{D}_0|} \sum_{s_0 \in \mathcal{D}_0} \pi(a|s) \hat{\Psi}(s|s_0), \quad (10)$$

gives an MIS estimator $\frac{1}{|\mathcal{D}|} \sum_{(s,a)\in\mathcal{D}} r^*(s, a) r(s, a)$ of $R(\pi)$ which is identical to the estimate of $R(\pi)$ computed directly with the SR.

Theorem 2 Let $\bar{r}(s, a)$ be the average reward in the dataset \mathcal{D} at the state-action pair (s, a) . Let $\hat{\Psi}$ be any approximate SR. The direct SR estimator $(1 - \gamma) \frac{1}{|\mathcal{D}_0|} \sum_{s_0 \in \mathcal{D}_0} \sum_{s \in \mathcal{S}} \hat{\Psi}(s|s_0) \sum_{a \in \mathcal{A}} \pi(a|s) \bar{r}(s, a)$ of $R(\pi)$ is identical to the MIS estimator $\frac{1}{|\mathcal{D}|} \sum_{(s,a) \in \mathcal{D}} r^*(s, a) r(s, a)$.

The take-away is that even when estimating the SR, the approximate density ratio defined by r^* is of sufficiently high quality to match the performance of directly estimating the value with the SR.

3.3. SR-DICE

Now we will consider how this MIS estimator can be generalized to continuous states by considering the deep SR ψ^π over features $\phi(s, a)$ and optimizing the weights of a linear function \mathbf{w} .

SR Refresher. We begin with a reminder of the details of the deep SR algorithm. The deep SR measures the expected sum of features $\psi^\pi(s, a) = \mathbb{E}_\pi [\sum_{t=0}^{\infty} \gamma^t \phi(s_t, a_t)]$. If the reward can be defined as a linear function over the features $r(s, a) = \mathbf{w}^\top \phi(s, a)$ then the value function can be recovered via a linear function over the deep SR $Q(s, a) = \mathbf{w}^\top \psi^\pi(s, a)$. The typical deep SR pipeline follows three steps:

1. Learn the encoding ϕ .
2. Learn the deep SR ψ^π over the encoding ϕ .
3. Learn \mathbf{w}_{SR} by minimizing $(\mathbf{w}_{\text{SR}}^\top \phi(s, a) - r(s, a))^2$.

We leave the first two stages vague as there is flexibility in how they are approached. This most commonly involves training the encoding ϕ via an encoder-decoder network to reconstruct transitions and training the deep SR ψ^π using TD learning-style methods (Kulkarni et al., 2016; Machado et al., 2018a). While we follow this standard practice, specific details are unimportant for our analysis and we relegate implementation-level details to the appendix.

Given the deep SR ψ^π , we can use it to learn the MIS ratio. Recall our objective of reward function optimization (Equation (6)). In the deep SR paradigm, both the reward and value function are determined by linear functions with respect to a single weight vector \mathbf{w} . Consequently, we can modify Equation (6) with these linear functions and then optimize the linear weights \mathbf{w} directly:

$$\begin{aligned} \min_{\mathbf{w}} J(\mathbf{w}) := & \frac{1}{2} \mathbb{E}_{\mathcal{D}} [(\mathbf{w}^\top \phi(s, a))^2] \\ & - (1 - \gamma) \mathbb{E}_{s_0, a_0 \sim \pi(\cdot|s_0)} [\mathbf{w}^\top \psi^\pi(s_0, a_0)], \end{aligned} \quad (11)$$

where in practice we replace the expectations with samples:

Algorithm 1 SR-DICE

Input: SR ψ , target network ψ' , encoder ϕ , decoder D .
 At each time step sample mini-batch of N transitions (s, a, r, s') and start states s_0 from \mathcal{D} .
for $t = 1$ **to** T_1 **do** # Encoding ϕ loss
 $\min_{\phi, D} \frac{1}{2} (D(\phi(s, a)) - (s, a))^2$.
for $t = 1$ **to** T_2 **do** # Deep SR ψ^π loss
 $\min_{\psi^\pi} \frac{1}{2} (\phi(s, a) + \gamma \psi'(s', a') - \psi^\pi(s, a))^2$.
for $t = 1$ **to** T_3 **do** # Density ratio \mathbf{w} loss (Equation (12))
 $a_0 \sim \pi(\cdot|s_0)$.
 $\min_{\mathbf{w}} \frac{1}{2} (\mathbf{w}^\top \phi(s, a))^2 - (1 - \gamma) \mathbf{w}^\top \psi^\pi(s_0, a_0)$.
Output: $|\mathcal{D}|^{-1} \sum_{(s,a,r) \in \mathcal{D}} \mathbf{w}^\top \phi(s, a) r(s, a) \approx R(\pi)$.

$$\begin{aligned} \min_{\mathbf{w}} J(\mathbf{w}) := & \frac{1}{2|\mathcal{D}|} \sum_{(s,a) \in \mathcal{D}} [(\mathbf{w}^\top \phi(s, a))^2] \\ & - (1 - \gamma) \frac{1}{|\mathcal{D}_0|} \sum_{s_0 \in \mathcal{D}_0, a_0} \pi(a_0|s_0) \mathbf{w}^\top \psi^\pi(s_0, a_0). \end{aligned} \quad (12)$$

Again, since the optimization problem Equation (12) is still convex, it has a closed form solution. Let Φ be a $|\mathcal{D}| \times F$ matrix where each row is the feature vector $\phi(s, a)$ with F features. Let Ψ be a $|\mathcal{D}_0| |\mathcal{A}| \times F$ matrix where each row is the SR weighted by its probability under the policy $\pi(a_0|s_0) \psi^\pi(s_0, a_0)$. Let $\mathbf{1}$ be a $|\mathcal{D}_0| |\mathcal{A}|$ dimensional vector of all 1. The unique optimizer \mathbf{w}^* of Equation (12) is a F dimensional vector defined as follows:

$$\mathbf{w}^* = (1 - \gamma) \frac{|\mathcal{D}|}{|\mathcal{D}_0|} (\Phi^\top \Phi)^{-1} \Psi^\top \mathbf{1}. \quad (13)$$

In practice, a matrix-based solution is often undesirable and we may prefer iterative, gradient-based solutions for scalability. In this case, we can directly minimize Equation (12) by taking gradient steps with respect to \mathbf{w} .

We now introduce our algorithm Successor Representation stationary DIstribution Correction Estimation (SR-DICE). SR-DICE follows the same first two steps of the standard SR procedure, but replaces the third step with optimizing Equation (12). Given \mathbf{w} , an estimate of $R(\pi)$ can be returned by $\frac{1}{|\mathcal{D}|} \sum_{(s,a,r(s,a)) \in \mathcal{D}} \mathbf{w}^\top \phi(s, a) r(s, a)$, where $\mathbf{w}^\top \phi(s, a) \approx \frac{d^\pi(s,a)}{d^{\mathcal{D}}(s,a)}$. We summarize SR-DICE in Algorithm 1.

We now remark upon two important properties of SR-DICE. The first concerns the quality of the quality of the learned MIS ratio. Although it is difficult to make any guarantees on the accuracy of an approximate ψ^π trained with deep RL techniques, if we assume ψ^π is exact, then we can show that SR-DICE learns the least squares estimator to the desired density ratio.

Theorem 3 If the deep SR is exact, such that $(1 - \gamma) \mathbb{E}_{s_0, a_0} [\psi^\pi(s_0, a_0)] = \mathbb{E}_{(s,a) \sim d^\pi} [\phi(s, a)]$, and the sup-

port of d^π is contained in the dataset \mathcal{D} , then the optimizer \mathbf{w}^* of Equation (12), as defined by Equation (13), is the least squares estimator of $\sum_{(s,a) \in \mathcal{D}} \left(\mathbf{w}^\top \phi(s,a) - \frac{d^\pi(s,a)}{d^{\mathcal{D}}(s,a)} \right)^2$.

The take-away from Theorem 3 is that our optimization problem, at least in the idealized setting, produces the same density ratios as directly learning them. This also means that the main source of error in SR-DICE is in the first two phases: learning the encoding ϕ and the deep SR ψ^π . Notably, both of these steps are independent of the main optimization problem of learning \mathbf{w} , as we have shifted the challenging aspects of density ratio estimation onto learning the deep SR. This leaves deep RL to do the heavy lifting. The remaining optimization problem, Equation (11), only involves directly updating the weights of a linear function, and unlike many other MIS methods, requires no tricky minimax optimization.

The second important property of SR-DICE is that Theorem 2 can be extended to the deep SR setting. That is, when derived from the same approximate SR, the optimal solution to both the SR-DICE estimator and the direct SR estimator produce identical estimates of $R(\pi)$.

Theorem 4 *Given the least squares estimator \mathbf{w}_{SR} of $\sum_{(s,a) \in \mathcal{D}} \left(\mathbf{w}^\top \phi(s,a) - r(s,a) \right)^2$ and the optimizer \mathbf{w}^* of Equation (12), as defined by Equation (13), then the traditional SR estimator $\frac{1}{|\mathcal{D}_0|} \sum_{s_0 \in \mathcal{D}_0} \mathbf{w}_{SR}^\top \psi^\pi(s_0, a_0)$ of $R(\pi)$ is identical to the SR-DICE estimator $\frac{1}{|\mathcal{D}|} \sum_{(s,a,r(s,a)) \in \mathcal{D}} \mathbf{w}^{*\top} \phi(s,a) r(s,a)$ of $R(\pi)$.*

This means that SR-DICE produces the same value estimate as the traditional deep SR algorithms, up to errors in the optimization process of \mathbf{w} . In other words, SR-DICE does not suffer from the same instability issues that plague other MIS methods when tackling high-dimensional domains where deep RL methods excel (relative to more traditional methods). Although, we typically think of the objective of MIS methods as policy evaluation, since SR-DICE and traditional deep SR produce the same value estimate, there is not a strong argument for using SR-DICE for policy evaluation. However, this also suggests that the estimated density ratios are of reasonably high quality since SR-DICE achieves the same performance as deep RL approaches. Therefore, we can treat SR-DICE as a tractable method for accessing the state-action occupancy of the target policy.

4. Related Work

Off-Policy Evaluation. Off-policy evaluation (OPE) is a well-studied problem with several families of approaches. One family of approaches is based on importance sampling, which re-weights trajectories by the ratio of likelihoods under the target and behavior policy (Precup et al., 2001).

Importance sampling methods are unbiased but suffer from variance which can grow exponentially with the length of trajectories (Li et al., 2015; Jiang & Li, 2016). Consequently, research has focused on variance reduction (Thomas & Brunskill, 2016; Munos et al., 2016; Farajtabar et al., 2018) or contextual bandits (Dudík et al., 2011; Wang et al., 2017). Marginalized importance sampling methods (Liu et al., 2018) aim to avoid this exponential variance by considering the ratio in stationary distributions, giving an estimator with variance which is polynomial with respect to horizon (Xie et al., 2019; Liu et al., 2019a). Follow-up work has introduced a variety of approaches and improvements, allowing them to be behavior-agnostic (Nachum et al., 2019a; Uehara & Jiang, 2019; Mousavi et al., 2020; Yang et al., 2020) and operate in the undiscounted setting (Zhang et al., 2020a;b). In a similar vein, some OPE methods rely on emphasizing, or re-weighting, updates based on their stationary distribution (Sutton et al., 2016; Mahmood et al., 2017; Hallak & Mannor, 2017; Gelada & Bellemare, 2019), or learning the stationary distribution directly (Wang et al., 2007; 2008).

For many deep RL algorithms (Mnih et al., 2015; Lillicrap et al., 2015), off-policy evaluation is based on TD learning (Sutton, 1988) and approximate dynamic programming techniques such as Fitted Q-Iteration (Ernst et al., 2005; Riedmiller, 2005; Yang et al., 2019). While empirically successful, these approaches lose any theoretical guarantees with non-linear function approximation (Tsitsiklis & Van Roy, 1997; Chen & Jiang, 2019). Regardless, they have been shown to achieve a high performance at benchmark OPE tasks (Voloshin et al., 2019; Fu et al., 2021).

Successor Representation. Introduced originally by Dayan (1993) as an approach for improving generalization in temporal-difference methods, successor representations (SR) were revived by recent work on deep successor RL (Kulkarni et al., 2016) and successor features (Barreto et al., 2017) which demonstrated that the SR could be generalized to a function approximation setting. The SR has found applications for task transfer (Barreto et al., 2018; Grimm et al., 2019), navigation (Zhang et al., 2017; Zhu et al., 2017), and exploration (Machado et al., 2018a; Janz et al., 2019). It has also been used in a neuroscience context to model generalization and human reinforcement learning (Gershman et al., 2012; Momennejad et al., 2017; Gershman, 2018). The SR and our work also relate to state representation learning (Lesort et al., 2018) and general value functions (Sutton & Tanner, 2005; Sutton et al., 2011).

5. Experiments

To evaluate our method, we perform several off-policy evaluation (OPE) experiments on a variety of domains. The aim is to evaluate the normalized average discounted reward $\mathbb{E}_{(s,a) \sim d^{\pi,r}} [r(s,a)]$ of a target policy π . We benchmark our




Figure 1: Off-policy evaluation results on the continuous action MuJoCo domain using the *easy* experimental setting (500k time steps and $\sigma_b = 0.133$), matching the setting of previous methods (Zhang et al., 2020a). The shaded area captures one standard deviation across 10 trials. We remark that this setting can be considered easy as the behavior policy achieves a lower error, often outperforming all agents. SR-DICE significantly outperforms the other MIS methods on all environments, except for Humanoid, where GradientDICE achieves a comparable performance.




Figure 2: Off-policy evaluation results on the continuous action MuJoCo domain using the *hard* experimental setting (50k time steps, $\sigma_b = 0.2$, random actions with $p = 0.2$). The shaded area captures one standard deviation across 10 trials. This setting uses significantly fewer time steps than the *easy* setting and the behavior policy is a poor estimate of the target policy. Again, we see SR-DICE outperforms the MIS methods, demonstrating the benefits of our proposed decomposition and simpler optimization. This setting also shows the benefits of deep RL methods over MIS methods for OPE in high-dimensional domains, as deep TD performs the strongest in every environment.

algorithm against two MIS methods, DualDICE (Nachum et al., 2019a) and GradientDICE (Zhang et al., 2020b), two deep RL approaches and the true return of the behavior policy. The first deep RL method is a DQN-style approach (Mnih et al., 2015) where actions are selected by π (denoted Deep TD) and the second is the deep SR where the weight w is trained to minimize the MSE between $w^\top \phi(s, a)$ and $r(s, a)$ (Kulkarni et al., 2016). Environment-specific experimental details are presented below, and complete algorithmic and hyper-parameter details are included in the appendix.

Continuous Action Experiments. We evaluate the methods on a variety of MuJoCo environments (Brockman et al., 2016; Todorov et al., 2012). We examine two experimental settings. In both settings the target policy π and behavior policy π_b are stochastic versions of a deterministic policy π_d obtained from training the TD3 algorithm (Fujimoto et al., 2018). We evaluate a target policy $\pi = \pi_d + \mathcal{N}(0, \sigma^2)$, where $\sigma = 0.1$.

- For the *easy* setting, we gather a dataset of 500k transi-

tions using a behavior policy $\pi_b = \pi_d + \mathcal{N}(0, \sigma_b^2)$, where $\sigma_b = 0.133$. This setting roughly matches the experimental setting used by GradientDICE Zhang et al. (2020a).

- For the *hard* setting, we gather a significantly smaller dataset of 50k transitions using a behavior policy which acts randomly with $p = 0.2$ and uses $\pi_d + \mathcal{N}(0, \sigma_b^2)$, where $\sigma_b = 0.2$, with $p = 0.8$.

Unless specified otherwise, we use a discount factor of $\gamma = 0.99$ and all hyper-parameters are kept constant across environments. All experiments are performed over 10 seeds. We display the results of the *easy* setting in Figure 1 and the *hard* setting in Figure 2.

Atari Experiments. To demonstrate our approach can scale to even more complex domains, we perform experiments with several Atari games (Bellemare et al., 2013), which are challenging due to their high-dimensional image-based state space. Standard pre-processing steps are applied (Castro et al., 2018) and sticky actions are used (Machado et al., 2018b) to increase difficulty and remove determinism. Each method is trained on a dataset of one million time steps. The

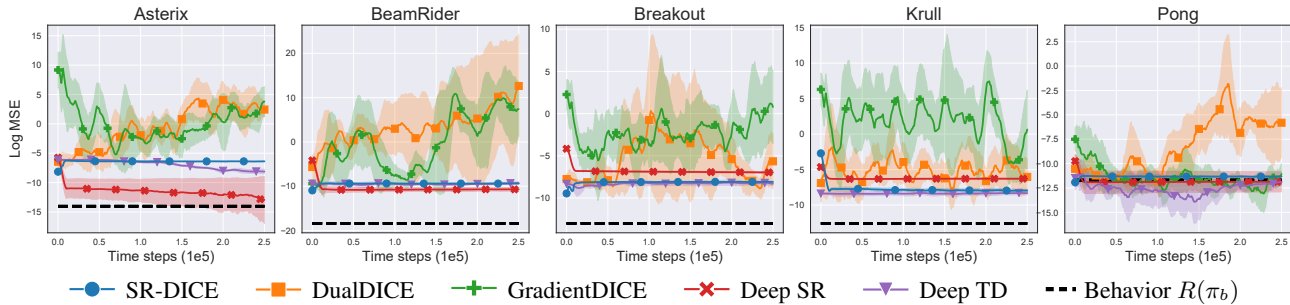


Figure 3: The log MSE for off-policy evaluation in the image-based Atari domain. This high-dimensional domain tests the ability of each method to scale to more complex environments. The shaded area captures one standard deviation across 3 trials. We can see the MIS baselines diverge on this challenging environment, while the remaining methods perform similarly. Perhaps surprisingly, on most games, the naïve baseline of using $R(\pi_b)$ from the behavior policy outperforms all methods by a fairly significant margin. Although the estimates from deep RL methods are stable, they are biased, resulting in a higher MSE.

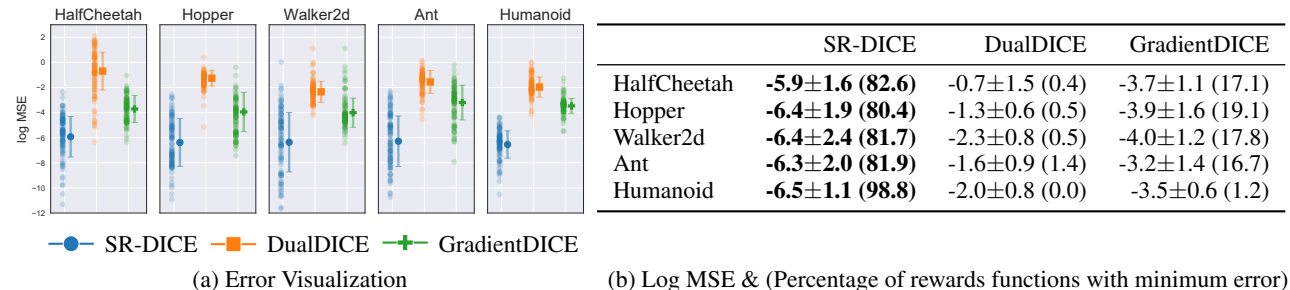


Figure 4: To evaluate the quality of the MIS ratios, we evaluate each MIS ratio with 1000 randomly sampled reward functions and compare to the ground truth on-policy value estimates. (Left) Visualization of the distribution of error. Only 100 points are displayed for visual clarity. Error bars are over the standard deviation. To normalize values across reward functions, we divide both the estimate and ground truth of $R(\pi)$ by the average reward in the dataset. (Right) Average log MSE and the standard deviation. In brackets is the percentage of reward functions where each method achieves the lowest error. We can see that SR-DICE achieves a low log MSE over a wide range of reward functions and outperforms the competing MIS methods on a high percentage of reward functions.

target policy is the deterministic greedy policy trained by Double DQN (Van Hasselt et al., 2016). The behavior policy is the ϵ -greedy policy with $\epsilon = 0.1$. We use a discount factor of $\gamma = 0.99$. Experiments are performed over 3 seeds. Results are displayed in Figure 3. Additional experiments with different behavior policies can be found in the appendix.

Evaluating the MIS ratios. To evaluate the quality of the MIS ratios themselves, we perform a randomized reward experiment. As the MIS ratio is only the value w that will return the true value of $R(\pi) = \mathbb{E}_{\mathcal{D}}[w \cdot r(s, a)]$ for all possible reward functions (Uehara & Jiang, 2019), we generate a large set of rewards functions with a randomly-initialized neural network, and evaluate the estimate of $R(\pi)$ obtained from each MIS method on each reward function. The ground-truth is estimated by a set of 100 on-policy trajectories generated by π . We generate 1000 reward functions, with scalar values in the range $[0, 10]$ and remove any

redundant reward functions from the set. The MIS ratios and dataset are taken from the *hard* setting. Experiments are performed over 5 seeds. We report the results in Figure 4.

Discussion. Across the board we find SR-DICE significantly outperforms the MIS methods. Looking at the estimated values of $R(\pi)$ in the continuous action environments, Figure 2, we can see that SR-DICE converges rapidly and maintains a stable estimate, while the MIS methods are particularly unstable, especially in the case of DualDICE. These observations are consistent in the Atari domain (Figure 3). In accordance with our theoretical analysis, Deep SR and SR-DICE perform similarly in every task, further suggesting that the limiting factor in SR-DICE is the quality of the deep successor representation, rather than learning the density ratios. In the randomized reward experiment, we find that SR-DICE vastly outperforms the other MIS methods in average log MSE, and compares favorably against

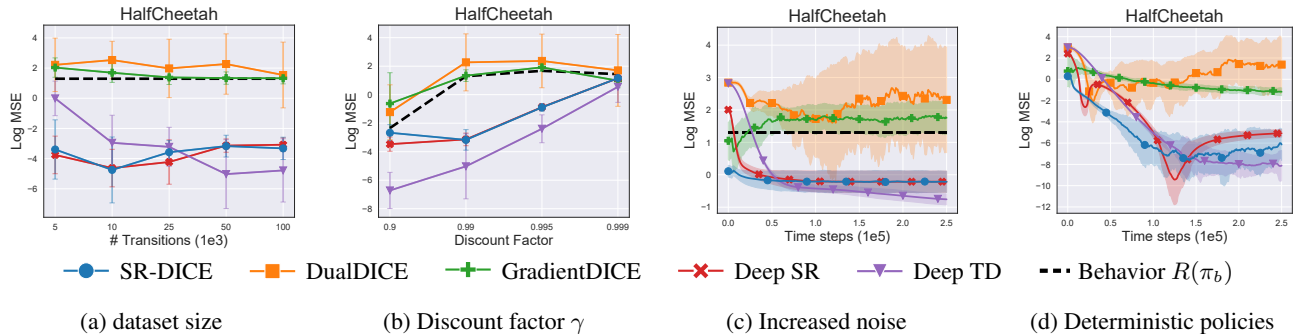


Figure 5: Ablation study results for the HalfCheetah task. We default to the *hard* setting wherever possible. Error bars and the shaded area captures one standard deviation over 10 trials. (a) We vary the size of the dataset \mathcal{D} . (b) We vary the discount factor γ . (c) We use a new behavior policy with $\mathcal{N}(0, \sigma_b^2)$ noise with $\sigma_b = 0.5$. (d) We use the same deterministic behavior and target policy.

the other MIS methods in over 80% of reward functions. In the most challenging task, Humanoid, SR-DICE is the best method in over 98% of reward functions. This suggests that SR-DICE provides much higher quality MIS ratio estimates than previous methods.

Ablation. To study the robustness of SR-DICE relative to the competing methods, we perform an ablation study and investigate the effects of dataset size, discount factor, and two different behavior policies. Unless specified otherwise, we use experimental settings matching the *hard* setting. We report the results in Figure 5. In the dataset size experiment (a), SR-DICE perform well with as few as 5k transitions (5 trajectories). In some instances, the performance is unexpectedly improved with less data, although incrementally. For small datasets, the SR methods outperform Deep TD. One hypothesis is that the encoding acts as an auxiliary reward and helps stabilize learning in the low data regime. In (b) we report the performance over changes in discount factor. The relative ordering across methods is unchanged. In (c) we use a behavior policy of $\mathcal{N}(0, \sigma_b^2)$, with $\sigma_b = 0.5$, a much larger standard deviation than either setting for continuous control. The results are similar to the original setting, with an increased bias on the deep RL methods. In (d) we use the underlying deterministic policy as both the behavior and target policy. Even though this setup should be easier since the task is no longer off-policy, the baseline MIS methods perform surprisingly poorly, once again demonstrating their weakness on high-dimensional domains.

6. Conclusion

In this paper, we introduce a method which can perform marginalized importance sampling (MIS) using the successor representation (SR) of the target policy. This is achieved by deriving an MIS formulation that can be viewed as reward function optimization. By using the SR, we effectively

disentangle the dynamics of the environment from learning the reward function. This allows us to (a) use well-known deep RL methods to effectively learn the SR in challenging domains (Mnih et al., 2015; Kulkarni et al., 2016) and (b) provide a straightforward loss function to learn the density ratios without any optimization tricks necessary for previous methods (Liu et al., 2018; Uehara & Jiang, 2019; Nachum et al., 2019a; Zhang et al., 2020b; Yang et al., 2020). Our resulting algorithm, SR-DICE, outperforms prior MIS methods in terms of both performance and stability and is the first MIS method which demonstrably scales to high-dimensional problems.

7. Acknowledgements

Scott Fujimoto is supported by a NSERC scholarship as well as the Borealis AI Global Fellowship Award. This research was enabled in part by support provided by Calcul Québec and Compute Canada. We would like to thank Wesley Chung, Pierre-Luc Bacon, Edward Smith, and Wei-Di Chang for helpful discussions and feedback.

References

- Barreto, A., Dabney, W., Munos, R., Hunt, J. J., Schaul, T., van Hasselt, H. P., and Silver, D. Successor features for transfer in reinforcement learning. In *Advances in neural information processing systems*, pp. 4055–4065, 2017.
- Barreto, A., Borsa, D., Quan, J., Schaul, T., Silver, D., Hessel, M., Mankowitz, D., Zidek, A., and Munos, R. Transfer in deep reinforcement learning using successor features and generalised policy improvement. In *International Conference on Machine Learning*, pp. 501–510, 2018.
- Bellemare, M. G., Naddaf, Y., Veness, J., and Bowling, M. The arcade learning environment: An evaluation plat-

- form for general agents. *Journal of Artificial Intelligence Research*, 47:253–279, 2013.
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. Openai gym, 2016.
- Castro, P. S., Moitra, S., Gelada, C., Kumar, S., and Bellemare, M. G. Dopamine: A research framework for deep reinforcement learning. *arXiv preprint arXiv:1812.06110*, 2018.
- Chen, J. and Jiang, N. Information-theoretic considerations in batch reinforcement learning. *arXiv preprint arXiv:1905.00360*, 2019.
- Dayan, P. Improving generalization for temporal difference learning: The successor representation. *Neural Computation*, 5(4):613–624, 1993.
- Dudík, M., Langford, J., and Li, L. Doubly robust policy evaluation and learning. In *Proceedings of the 28th International Conference on International Conference on Machine Learning*, pp. 1097–1104, 2011.
- Ernst, D., Geurts, P., and Wehenkel, L. Tree-based batch mode reinforcement learning. *Journal of Machine Learning Research*, 6(Apr):503–556, 2005.
- Farajtabar, M., Chow, Y., and Ghavamzadeh, M. More robust doubly robust off-policy evaluation. In *International Conference on Machine Learning*, pp. 1447–1456, 2018.
- Fu, J., Norouzi, M., Nachum, O., Tucker, G., Wang, Z., Novikov, A., Yang, M., Zhang, M. R., Chen, Y., Kumar, A., Paduraru, C., Levine, S., and Paine, T. Benchmarks for deep off-policy evaluation. In *International Conference on Learning Representations*, 2021.
- Fujimoto, S., van Hoof, H., and Meger, D. Addressing function approximation error in actor-critic methods. In *International Conference on Machine Learning*, volume 80, pp. 1587–1596. PMLR, 2018.
- Fujimoto, S., Meger, D., and Precup, D. Off-policy deep reinforcement learning without exploration. In *International Conference on Machine Learning*, pp. 2052–2062, 2019.
- Gauci, J., Conti, E., Liang, Y., Virochsiri, K., He, Y., Kaden, Z., Narayanan, V., Ye, X., Chen, Z., and Fujimoto, S. Horizon: Facebook’s open source applied reinforcement learning platform. *arXiv preprint arXiv:1811.00260*, 2018.
- Gelada, C. and Bellemare, M. G. Off-policy deep reinforcement learning by bootstrapping the covariate shift. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 3647–3655, 2019.
- Gershman, S. J. The successor representation: its computational logic and neural substrates. *Journal of Neuroscience*, 38(33):7193–7200, 2018.
- Gershman, S. J., Moore, C. D., Todd, M. T., Norman, K. A., and Sederberg, P. B. The successor representation and temporal context. *Neural Computation*, 24(6):1553–1568, 2012.
- Grimm, C., Higgins, I., Barreto, A., Teplyashin, D., Wulfmeier, M., Hertweck, T., Hadsell, R., and Singh, S. Disentangled cumulants help successor representations transfer to new tasks. *arXiv preprint arXiv:1911.10866*, 2019.
- Hallak, A. and Mannor, S. Consistent on-line off-policy evaluation. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 1372–1383. JMLR. org, 2017.
- Imani, E., Graves, E., and White, M. An off-policy policy gradient theorem using emphatic weightings. In *Advances in Neural Information Processing Systems*, pp. 96–106, 2018.
- Jaakkola, T., Jordan, M. I., and Singh, S. P. On the convergence of stochastic iterative dynamic programming algorithms. *Neural computation*, 6(6):1185–1201, 1994.
- Janz, D., Hron, J., Mazur, P., Hofmann, K., Hernández-Lobato, J. M., and Tschitschek, S. Successor uncertainties: exploration and uncertainty in temporal difference learning. In *Advances in Neural Information Processing Systems*, pp. 4509–4518, 2019.
- Jiang, N. and Li, L. Doubly robust off-policy value evaluation for reinforcement learning. In *International Conference on Machine Learning*, pp. 652–661, 2016.
- Kostrikov, I., Nachum, O., and Tompson, J. Imitation learning via off-policy distribution matching. *arXiv preprint arXiv:1912.05032*, 2019.
- Kulkarni, T. D., Saeedi, A., Gautam, S., and Gershman, S. J. Deep successor reinforcement learning. *arXiv preprint arXiv:1606.02396*, 2016.
- Kumar, A., Fu, J., Soh, M., Tucker, G., and Levine, S. Stabilizing off-policy q-learning via bootstrapping error reduction. In *Advances in Neural Information Processing Systems*, pp. 11784–11794, 2019.
- Lesort, T., Díaz-Rodríguez, N., Goudou, J.-F., and Filliat, D. State representation learning for control: An overview. *Neural Networks*, 108:379–392, 2018.
- Li, L., Munos, R., and Szepesvari, C. Toward minimax off-policy value estimation. In *Artificial Intelligence and Statistics*, pp. 608–616, 2015.

- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- Liu, Q., Li, L., Tang, Z., and Zhou, D. Breaking the curse of horizon: Infinite-horizon off-policy estimation. In *Advances in Neural Information Processing Systems*, pp. 5356–5366, 2018.
- Liu, Y., Bacon, P.-L., and Brunskill, E. Understanding the curse of horizon in off-policy evaluation via conditional importance sampling. *arXiv preprint arXiv:1910.06508*, 2019a.
- Liu, Y., Swaminathan, A., Agarwal, A., and Brunskill, E. Off-policy policy gradient with state distribution correction. *arXiv preprint arXiv:1904.08473*, 2019b.
- Machado, M. C., Rosenbaum, C., Guo, X., Liu, M., Tesauro, G., and Campbell, M. Eigenoption discovery through the deep successor representation. *arXiv preprint arXiv:1710.11089*, 2017.
- Machado, M. C., Bellemare, M. G., and Bowling, M. Count-based exploration with the successor representation. *arXiv preprint arXiv:1807.11622*, 2018a.
- Machado, M. C., Bellemare, M. G., Talvitie, E., Veness, J., Hausknecht, M., and Bowling, M. Revisiting the arcade learning environment: Evaluation protocols and open problems for general agents. *Journal of Artificial Intelligence Research*, 61:523–562, 2018b.
- Mahmood, A. R., Yu, H., and Sutton, R. S. Multi-step off-policy learning without importance sampling ratios. *arXiv preprint arXiv:1702.03006*, 2017.
- Mandel, T., Liu, Y.-E., Levine, S., Brunskill, E., and Popovic, Z. Offline policy evaluation across representations with applications to educational games. In *International Conference on Autonomous Agents and Multiagent Systems*, 2014.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529–533, 2015.
- Momennejad, I., Russek, E. M., Cheong, J. H., Botvinick, M. M., Daw, N. D., and Gershman, S. J. The successor representation in human reinforcement learning. *Nature Human Behaviour*, 1(9):680–692, 2017.
- Mousavi, A., Li, L., Liu, Q., and Zhou, D. Black-box off-policy estimation for infinite-horizon reinforcement learning. *arXiv preprint arXiv:2003.11126*, 2020.
- Munos, R., Stepleton, T., Harutyunyan, A., and Bellemare, M. Safe and efficient off-policy reinforcement learning. In *Advances in Neural Information Processing Systems*, pp. 1054–1062, 2016.
- Nachum, O., Chow, Y., Dai, B., and Li, L. Dualdice: Behavior-agnostic estimation of discounted stationary distribution corrections. In *Advances in Neural Information Processing Systems*, pp. 2315–2325, 2019a.
- Nachum, O., Dai, B., Kostrikov, I., Chow, Y., Li, L., and Schuurmans, D. Algaedice: Policy gradient from arbitrary experience. *arXiv preprint arXiv:1912.02074*, 2019b.
- Precup, D., Sutton, R. S., and Dasgupta, S. Off-policy temporal-difference learning with function approximation. In *International Conference on Machine Learning*, pp. 417–424, 2001.
- Riedmiller, M. Neural fitted q iteration—first experiences with a data efficient neural reinforcement learning method. In *European Conference on Machine Learning*, pp. 317–328. Springer, 2005.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Sinha, S., Song, J., Garg, A., and Ermon, S. Experience replay with likelihood-free importance weights. *arXiv preprint arXiv:2006.13169*, 2020.
- Sutton, R. S. Learning to predict by the methods of temporal differences. *Machine learning*, 3(1):9–44, 1988.
- Sutton, R. S. and Barto, A. G. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge, 1998.
- Sutton, R. S. and Tanner, B. Temporal-difference networks. In *Advances in neural information processing systems*, pp. 1377–1384, 2005.
- Sutton, R. S., Modayil, J., Delp, M., Degris, T., Pilarski, P. M., White, A., and Precup, D. Horde: a scalable real-time architecture for learning knowledge from unsupervised sensorimotor interaction. In *The 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pp. 761–768, 2011.
- Sutton, R. S., Mahmood, A. R., and White, M. An emphatic approach to the problem of off-policy temporal-difference learning. *The Journal of Machine Learning Research*, 17 (1):2603–2631, 2016.
- Swaminathan, A., Krishnamurthy, A., Agarwal, A., Dudik, M., Langford, J., Jose, D., and Zitouni, I. Off-policy evaluation for slate recommendation. In *Advances in*

- Neural Information Processing Systems*, pp. 3632–3642, 2017.
- Thomas, P. and Brunskill, E. Data-efficient off-policy policy evaluation for reinforcement learning. In *International Conference on Machine Learning*, pp. 2139–2148, 2016.
- Todorov, E., Erez, T., and Tassa, Y. Mujoco: A physics engine for model-based control. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 5026–5033. IEEE, 2012.
- Touati, A., Zhang, A., Pineau, J., and Vincent, P. Stable policy optimization via off-policy divergence regularization. *arXiv preprint arXiv:2003.04108*, 2020.
- Tsitsiklis, J. N. and Van Roy, B. Analysis of temporal-difference learning with function approximation. In *Advances in neural information processing systems*, pp. 1075–1081, 1997.
- Uehara, M. and Jiang, N. Minimax weight and q-function learning for off-policy evaluation. *arXiv preprint arXiv:1910.12809*, 2019.
- Van Hasselt, H., Guez, A., and Silver, D. Deep reinforcement learning with double q-learning. In *AAAI*, pp. 2094–2100, 2016.
- Voloshin, C., Le, H. M., Jiang, N., and Yue, Y. Empirical study of off-policy policy evaluation for reinforcement learning. *arXiv preprint arXiv:1911.06854*, 2019.
- Wang, T., Bowling, M., and Schuurmans, D. Dual representations for dynamic programming and reinforcement learning. In *2007 IEEE International Symposium on Approximate Dynamic Programming and Reinforcement Learning*, pp. 44–51. IEEE, 2007.
- Wang, T., Bowling, M., Schuurmans, D., and Lizotte, D. J. Stable dual dynamic programming. In *Advances in neural information processing systems*, pp. 1569–1576, 2008.
- Wang, Y.-X., Agarwal, A., and Dudík, M. Optimal and adaptive off-policy evaluation in contextual bandits. In *International Conference on Machine Learning*, pp. 3589–3597, 2017.
- Xie, T., Ma, Y., and Wang, Y.-X. Towards optimal off-policy evaluation for reinforcement learning with marginalized importance sampling. In *Advances in Neural Information Processing Systems*, pp. 9665–9675, 2019.
- Yang, M., Nachum, O., Dai, B., Li, L., and Schuurmans, D. Off-policy evaluation via the regularized lagrangian. *Advances in Neural Information Processing Systems*, 33, 2020.
- Yang, Z., Xie, Y., and Wang, Z. A theoretical analysis of deep q-learning. *arXiv preprint arXiv:1901.00137*, 2019.
- Zhang, J., Springenberg, J. T., Boedecker, J., and Burgard, W. Deep reinforcement learning with successor features for navigation across similar environments. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 2371–2378. IEEE, 2017.
- Zhang, R., Dai, B., Li, L., and Schuurmans, D. Gendice: Generalized offline estimation of stationary values. *arXiv preprint arXiv:2002.09072*, 2020a.
- Zhang, S., Boehmer, W., and Whiteson, S. Generalized off-policy actor-critic. In *Advances in Neural Information Processing Systems*, pp. 1999–2009, 2019.
- Zhang, S., Liu, B., and Whiteson, S. Gradientdice: Rethinking generalized offline estimation of stationary values. *arXiv preprint arXiv:2001.11113*, 2020b.
- Zhao, Y., Kosorok, M. R., and Zeng, D. Reinforcement learning design for cancer clinical trials. *Statistics in medicine*, 28(26):3294, 2009.
- Zhu, Y., Gordon, D., Kolve, E., Fox, D., Fei-Fei, L., Gupta, A., Mottaghi, R., and Farhadi, A. Visual semantic planning using deep successor representations. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 483–492, 2017.