

Low Variance Off-policy Evaluation with State-based Importance Sampling

1st David M. Bossens

IHPC

Agency for Science, Technology and Research

Singapore

CFAR

Agency for Science, Technology and Research

Singapore

bossensdm@cfar.a-star.edu.sg

2nd Philip S. Thomas

Manning College of Information and Computer Sciences

University of Massachusetts

Amherst, United States

pthomas@cs.umass.edu

Abstract—In many domains, the exploration process of reinforcement learning will be too costly as it requires trying out suboptimal policies, resulting in a need for off-policy evaluation, in which a target policy is evaluated based on data collected from a known behaviour policy. In this context, importance sampling estimators provide estimates for the expected return by weighting the trajectory based on the probability ratio of the target policy and the behaviour policy. Unfortunately, such estimators have a high variance and therefore a large mean squared error. This paper proposes state-based importance sampling estimators which reduce the variance by dropping certain states from the computation of the importance weight. To illustrate their applicability, we demonstrate state-based variants of ordinary importance sampling, weighted importance sampling, per-decision importance sampling, incremental importance sampling, doubly robust off-policy evaluation, and stationary density ratio estimation. Experiments in four domains show that state-based methods consistently yield reduced variance and improved accuracy compared to their traditional counterparts.

Index Terms—off-policy reinforcement learning, off-policy evaluation, importance sampling, variance reduction

I. INTRODUCTION

In reinforcement learning [1], an agent explores an environment to learn a policy that optimises a utility function, such as the expected cumulative reward, in that environment. In many domains, including medical interventions (e.g. surgery, diabetes or sepsis), financial decision-making (e.g. advertising or trading), and safe navigation, it would be preferable to avoid costly trial-and-error and instead use a known expert or safe policy to gather the data needed for reinforcement learning. Such domains motivate the use of off-policy evaluation [2], [3], a class of methods which estimate the utility of a target policy based on trajectories from a behaviour policy.

Importance sampling (IS) is a traditional statistical technique which estimates an integral over one distribution of interest while only having samples from another distribution. The expected cumulative reward can be seen as such an integral, and therefore IS is an important technique for off-policy evaluation [4]–[6]. Such techniques have been used

successfully in the context of diabetes treatment [6] and digital marketing [5]. Unfortunately, the variance of IS based off-policy evaluation grows exponentially with the horizon H of the decision process, a problem also known as “the curse of horizon” [7]. This makes IS based off-policy evaluation perform more poorly in domains requiring long-term planning (e.g. safe navigation).

To reduce the variance of off-policy evaluation, this paper proposes a technique that can be used on a wide variety of estimators. The paper makes the following contributions:

- state-based importance sampling, a class of state-based importance sampling estimators which remove “negligible states” from the importance weight computation;
- two methods to identify negligible states, one based on covariance testing [8] and one based on state-action values;
- implementations of state-based variants of ordinary importance sampling, weighted importance sampling [4], [9], per-decision importance sampling [4], incremental importance sampling [8], doubly robust off-policy evaluation [10], and stationary density ratio estimation [7]; and
- a set of empirical experiments demonstrating the performance of these state-based estimators compared to their traditional counterparts.

II. PROBLEM STATEMENT AND PRELIMINARIES

This present paper is based on finite-horizon undiscounted Markov decision processes (MDPs), which can be defined based on the tuple $(\mathcal{S}, \mathcal{A}, r, \mathcal{T})$, where \mathcal{S} is the state space, \mathcal{A} is the action space, $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is a reward function, and $\mathcal{T} : \mathcal{S} \times \mathcal{A} \rightarrow \Delta^{\mathcal{S}}$ is the transition dynamics model defined over the probability simplex over the state space, $\Delta^{\mathcal{S}}$. The goal of the reinforcement learning agent is to select actions that maximise the expected cumulative reward, which is also called the expected return. The decision-making is represented by the policy of the agent, which takes the form $\pi : \mathcal{S} \rightarrow \Delta^{\mathcal{A}}$; that is, the policy π outputs the action distribution for each given state.

DMB was supported by UKRI Trustworthy Autonomous Systems Hub, EP/V00784X/1 and PST was supported by NSF grant no. CCF-2018372.

Off-policy evaluation computes the expected return \mathcal{G} of an evaluation policy π_e based on samples from a behaviour policy π_b . This can be done using importance sampling,

$$\begin{aligned}\mathcal{G} &= \mathbb{E}_{\pi_e} \left[\sum_{t=1}^H r_t \right] \\ &= \mathbb{E}_{\pi_b} \left[\sum_{t=1}^H r_t \left(\prod_{t=1}^H \frac{\pi_e(a_t|s_t)}{\pi_b(a_t|s_t)} \right) \right],\end{aligned}$$

which samples from π_b and then corrects for the probability under π_e , using the product $\prod_{t=1}^H \frac{\pi_e(a_t|s_t)}{\pi_b(a_t|s_t)}$ as the importance weight of the trajectory.

A long-standing problem with importance sampling in this manner is that if the horizon is large, then the variance on the estimate of \mathcal{G} is prohibitively large. To see why, define the estimate \hat{G}_{IS}

$$\hat{G}_{IS} = \frac{1}{n} \sum_{i=1}^n \sum_{t=1}^H r_t^{(i)} \left(\prod_{t=1}^H \frac{\pi_e(a_t^{(i)}|s_t^{(i)})}{\pi_b(a_t^{(i)}|s_t^{(i)})} \right),$$

where i indexes the trajectory of the state $s_t^{(i)}$, action $a_t^{(i)}$, and reward $r_t^{(i)}$ at time t . By sample mean of a random variable $X = \sum_{t=1}^H r_t \left(\prod_{t=1}^H \frac{\pi_e(a_t|s_t)}{\pi_b(a_t|s_t)} \right)$ and Popoviciu's inequality [11],

$$\begin{aligned}\text{Var}(\hat{G}_{IS}) &= \frac{1}{n} \text{Var}(X) \\ &\leq \frac{1}{4n} \left(\sum_{t=1}^H r_{\max} \left(\prod_{t=1}^H \rho_{\max} \right) \right)^2 \\ &= \frac{1}{4n} (H r_{\max} \rho_{\max}^H)^2 \\ &= \mathcal{O}(\exp(H)),\end{aligned}\quad (1)$$

where $\rho_{\max} = \max_{(s,a) \in \mathcal{S} \times \mathcal{A}} \pi_e(a|s)/\pi_b(a|s)$ and $[0, r_{\max}]$ is the reward range. In addition to the general upper bound, it is also possible, with some additional assumptions, to show that the variance is lower bounded by $\Omega(\exp(H))$ [12]. Due to the bias-variance decomposition of the mean squared error,

$$\text{MSE}(\hat{G}_{IS}) = \text{Bias}(\hat{G}_{IS})^2 + \text{Var}(\hat{G}_{IS}),$$

where $\text{Bias}(\hat{G}_{IS}) = \mathbb{E}_{\pi_b}[\hat{G}_{IS}] - \mathcal{G}$, the above results imply an exponential dependency of the MSE on the horizon of the decision problem.

III. RELATED WORK

Due to the curse of horizon, variance reduction is one of the important problems in off-policy evaluation. We review variance reduction techniques and provide exemplary off-policy evaluation applications.

A. Variance reduction in importance sampling

While ordinary importance sampling takes a simple average across trajectories, dividing by a denominator n , the weighted importance sampling (WIS) technique [4], [9] divides by the sum of the importance weights of the different trajectories.

This mitigates the impact of excessively large importance ratios. Its resulting variance converges to 0 for $n \rightarrow \infty$ [9] but remains exponential in the horizon [7]. While state-based importance sampling does not work similarly to weighted importance sampling, it shares the commonality that it serves as an easy plug-in to a wide range of estimators.

Another common technique is to remove some of the time steps from the cumulative product, which results in a lower exponent in Eq. 1 and mitigates the curse of horizon. In per-decision importance sampling (PDIS) [4], one forms different importance weights for each reward $r_{t'}$ by relating it to the past time steps until that point, i.e. $\rho_{t'} = \prod_{t=1}^{t'} \frac{\pi_e(a_t^{(i)}|s_t^{(i)})}{\pi_b(a_t^{(i)}|s_t^{(i)})}$. Incremental importance sampling (INCRIS) [8] drops all but the k most recent action probability ratios from the importance weight of the per-decision importance sampling technique, where the number k is selected based on a covariance test. Removing time steps in this manner is not always a natural solution since the i 'th decision in one trajectory may not be related to the i 'th decision in another trajectory. In state-based importance sampling, one drops action probability ratios associated with a particular set of states rather than a particular set of time steps. In doing so, a similar covariance test is proposed as one technique to identify the set of states to drop.

It is also possible to avoid products of action probability ratios in favour of sums over state density ratios. Stationary density ratio estimation (SDRE) [7] avoids the cumulative product across time and instead is based on the average visitation distribution of state-action pairs. Denoting the average visitation distribution of a state $s_t^{(i)}$ as $d_t^{\pi_e}(s_t^{(i)})$ and its estimate as $\hat{d}_t^{\pi_e}(s_t^{(i)})$, SDRE estimates the return as

$$\hat{G}_{SDRE} = \frac{1}{n} \sum_{i=1}^n \sum_{t=1}^H r_t^{(i)} \frac{\hat{d}_t^{\pi_e}(s_t^{(i)})}{\hat{d}_t^{\pi_b}(s_t^{(i)})}, \quad (2)$$

thereby removing the curse of horizon such that the variance depends polynomially, rather than exponentially, on H . The approach is thereby suitable for large and even infinite-horizon problems but requires additional assumptions of ergodicity, a finite state space, and a clearly defined stationary distribution. In similar spirit, Marginalised importance sampling (MIS) [13] is formulated as

$$\hat{G}_{MIS} = \frac{1}{n} \sum_{i=1}^n \sum_{t=1}^H \hat{r}_t^{\pi_e}(s_t^{(i)}) \frac{\hat{d}_t^{\pi_e}(s_t^{(i)})}{\hat{d}_t^{\pi_b}(s_t^{(i)})}, \quad (3)$$

where one can note the difference with SDRE by indexing the state distribution $\hat{d}_t^{\pi_e}$ at a particular time t as well as estimating the reward $\hat{r}_t^{\pi_e}(s_t^{(i)})$ for a particular time t . The estimator of the state distribution is updated according to $\hat{d}_t^{\pi_e} = \hat{\mathcal{T}}_t^{\pi_e}(s_t|s_{t-1})\hat{d}_{t-1}^{\pi_e}$, where $\hat{\mathcal{T}}_t^{\pi_e}(s_t|s_{t-1})$ is the empirical importance-weighted visitation of $s_{t-1} \rightarrow s_t$. The estimator $\hat{r}_t^{\pi_e}$ is the importance-weighted reward for transitions $s_{t-1} \rightarrow s_t$. Within the formalism of a tabular MDP, i.e. where state space and actions space are limited such that each state-action pair can be frequently visited, Tabular MIS [14] provides an alternative approximation to $\hat{r}_t^{\pi_e}(s_t^{(i)})$ and $\hat{d}_t^{\pi_e}(s_t^{(i)})$ based on

the observed visitation frequencies. Tabular MIS provides a lower MSE bound in line with the Cramer-Rao bound of $\Omega(H^2/n)$. Compared to techniques using state density ratios, which require stationary state distributions and various function approximators, state-based estimators are applicable under varying assumptions and algorithmic frameworks, depending on which estimator it is plugged into (e.g. ordinary importance sampling vs stationary density ratio estimation).

Conditional importance sampling (CIS) uses conditional Monte Carlo over the importance weight to yield lowered variance. Denoting $\rho_{1:t'} = \prod_{t=1}^{t'} \frac{\pi_e(a_t|s_t)}{\pi_b(a_t|s_t)}$ and ϕ_t as a statistic obtained from the trajectory, CIS can be written as

$$\hat{G}_{CIS} = \sum_{i=1}^n \sum_{t=1}^H r_t^{(i)} \mathbb{E}_{\pi_b}[\rho_{1:H}|\phi_t], \quad (4)$$

which yields $\text{Var}(G\mathbb{E}_{\pi_b}[\rho_{1:H}|\phi_t]) \leq \text{Var}(G\rho_{1:H})$. Liu et al. (2020) [12] make use of this property to better understand IS, PDIS, and marginalised IS, noting that in this framework the PDIS estimate has $\phi_t = s_{1:t}, a_{1:t}$ and the marginalised IS estimate has $\phi_t = s_t, a_t$. Rowland et al. (2020) [15] present a related framework for conditional importance sampling with emphasis on trajectories given particular start state-action pair. They introduce novel algorithms arising within the framework, namely return-conditioned, reward-conditioned, and state-conditioned importance sampling. When the equations cannot be solved analytically, the technique often comes with a requirement to form an importance weight regression model, which may increase variance and reduce accuracy. Rather than conditioning on states, state-based methods drop selected states from the importance weight every time they occur in the trajectory.

State-based techniques can also be applied when an accurate model is available. In particular, doubly robust (DR) methods [10], [16]–[18] use estimators that combine importance sampling with the direct method, which uses an estimated model to infer the return of the evaluation policy, using traditional statistical techniques for missing data (e.g. [19]). The doubly robust estimator benefits from the low variance of the direct method and the low bias – or unbiasedness, if the behaviour policy is known – of importance sampling methods. It is a consistent estimator [10], implying its variance converges to 0 for $n \rightarrow \infty$.

Other variance reduction techniques not investigated in this paper are of further interest, including options [8], [20], multiple importance sampling [21], [22], and estimator selection [2].

B. Applying off-policy evaluation

Off-policy evaluation has two primary uses. First, when few trajectories are available and the designer wants to make a decision on whether or not to use the evaluation policy instead of a given expert policy; primary examples include medical treatment [6] and digital marketing [23], which would incur significant costs and safety risks when the evaluation policy's performance is not estimated correctly. Second, off-policy evaluation can be used to optimise the policy, where key

challenges include how to provide monotonic improvement and how to provide confidence estimates on improvement. For example, high-confidence policy improvement (HCOPE) [24] uses confidence estimates on the evaluation [25] to ensure with high probability that the new policy improves on the old; policy optimisation via importance sampling (POIS) [26] optimises surrogate objective that effectively captures the trade-off between the estimated improvement and the uncertainty due to importance sampling [22]; and Uniform OPE [27] simultaneously evaluates all policies in a policy class.

IV. STATE-BASED IMPORTANCE SAMPLING

State-based importance sampling (SIS) mitigates the above issue by constructing an estimator \hat{G}_{SIS} that selectively drops the action probability ratios of a select state set $\mathcal{S}^A \subset \mathcal{S}$ from the product to compute the importance weight. If the dropped term is equal to 1 and does not covary with the estimator \hat{G}_{SIS} , this yields lower-variance estimates. In this section, we demonstrate the state-based formalism for ordinary importance sampling while Sec. VI shows how to apply the state-based formalism for other estimators with lower variance and mean squared error. We extend ordinary importance sampling here for its simplicity and the fact that its variance is more significantly reduced (i.e. by reducing the exponent).

Let $\mathcal{S}^A \subset \mathcal{S}$ and $\mathcal{S}^B = \mathcal{S} \setminus \mathcal{S}^A$ its complement. Let $G = \sum_{t=1}^H r_t$. Also for any given trajectory $\tau = \{s_1, a_1, \dots, s_H, a_H\}$ taken by a policy π_b on an MDP \mathcal{M} , let

$$A = \begin{cases} \prod_{t \in [H]: s_t \in \mathcal{S}^A} \frac{\pi_e(a_t|s_t)}{\pi_b(a_t|s_t)}, & \text{if } \exists s \in \mathcal{S}^A \cap \tau \\ 1, & \text{otherwise} \end{cases}$$

$$B = \begin{cases} \prod_{t \in [H]: s_t \in \mathcal{S}^B} \frac{\pi_e(a_t|s_t)}{\pi_b(a_t|s_t)}, & \text{if } \exists s \in \mathcal{S}^B \cap \tau \\ 1, & \text{otherwise,} \end{cases}$$

where $[H] = \{1, 2, \dots, H\}$.

Given a selected state-set \mathcal{S}^A of dropped states, the state-based variant of ordinary importance sampling computes the action probability ratios only over its complement \mathcal{S}^B :

$$\hat{G}_{SIS}(\mathcal{S}^A) = \frac{1}{n} \sum_{i=1}^n G^{(i)} \left(\prod_{t \in [H]: s_t^{(i)} \in \mathcal{S}^B} \frac{\pi_e(a_t^{(i)}|s_t^{(i)})}{\pi_b(a_t^{(i)}|s_t^{(i)})} \right). \quad (5)$$

Its expected return can be rewritten as

$$\mathcal{G} = \mathbb{E}_{\pi_b} [A] \mathbb{E}_{\pi_b} [BG] + \text{Cov}(A, BG), \quad (6)$$

since $\text{Cov}(X, Y) = \mathbb{E}[XY] - \mathbb{E}[X]\mathbb{E}[Y]$. As shown in the covariance testing framework [8], any decomposition of the form (6) allows to accurately estimate \mathcal{G} based on samples of BG if $\mathbb{E}_{\pi_b} [A] \approx 1$ and $\text{Cov}(A, BG)$ is sufficiently small. For the first condition, the random variable A can be seen as the importance weight under a different evaluation policy π_e^A , which we will call the **\mathcal{S}^A -reduced evaluation policy** of π_e , which is equal to π_e for $s \in \mathcal{S}^A$ and equal to π_b for $s \in \mathcal{S}^B$; therefore it follows that $\mathbb{E}[A] = 1$ (see Appendix A). The second condition does not follow automatically, but depends

on the chosen \mathcal{S}^A ; the definition below specifies the criterion for \mathcal{S}^A to form an accurate state-based estimator.

Definition 1. (π_e, π_b, ϵ) -*decomposable MDP and ϵ -negligible state set.* Let $\epsilon > 0$, \mathcal{M} be an MDP $\langle \mathcal{S}, \mathcal{A}, r, \mathcal{T} \rangle$, and let π_e and π_b two policies in \mathcal{M} . Then \mathcal{M} is (π_e, π_b, ϵ) -decomposable iff there exists a set \mathcal{S}^A for which $\text{Cov}(A, BG) < \epsilon$. Any such set \mathcal{S}^A is called an ϵ -negligible state set for the off-policy evaluation problem $\langle \mathcal{M}, \pi_e, \pi_b \rangle$.

Applying bias-variance decomposition, it follows that

$$\text{MSE}(\hat{G}_{SIS}(\mathcal{S}^A)) = \text{Cov}(A, BG)^2 + \text{Var}(\hat{G}_{SIS}(\mathcal{S}^A)). \quad (7)$$

Denoting the subtrajectory $\tau_B := \{t \in [H] : s_t^{(i)} \in \mathcal{S}^B\}$ and $M_B := \max_{\tau \in \tau_B} |\tau|$, we have

$$\text{Var}(\hat{G}_{SIS}(\mathcal{S}^A)) \leq \frac{1}{4n} (Hr_{\max}\rho_{\max}^{M_B})^2. \quad (8)$$

The resulting variance upper bound of \hat{G}_{SIS} is exponential in the maximal number of occurrences of states in \mathcal{S}^B rather than in the full trajectory length H . Consequently, the SIS estimator will significantly improve on IS if many visited states can be dropped from the computation of the importance weight. When \mathcal{S}^A is an ϵ -negligible state set, this implies via (7) and (8) that

$$\text{MSE}(\hat{G}_{SIS}(\mathcal{S}^A)) \leq \epsilon^2 + \frac{1}{4n} (Hr_{\max}\rho_{\max}^{M_B})^2. \quad (9)$$

V. OPTIMISING THE DROPPED STATE SET

We propose two search algorithms to optimise the dropped state set, one based on covariance testing and one based on state-action values (a.k.a. Q-values). Additional derivations are presented in Appendix B showing that even with some additional bias, the performance of a state-based variant of ordinary importance sampling (SIS) is still a significant improvement over ordinary importance sampling.

A. Covariance testing

Following the covariance testing approach [8], this algorithm applies a bias-variance decomposition based on empirical estimates of the covariance and variance, and compares the MSEs of different parameter choices. In our case, the parameter choice is the state set \mathcal{S}^A for a decomposition

$$\widehat{\text{MSE}}(\hat{G}_{SIS}(\mathcal{S}^A)) = \widehat{\text{Cov}}(A, BG)^2 + \widehat{\text{Var}}(\hat{G}_{SIS}(\mathcal{S}^A)), \quad (10)$$

and the algorithm encourages larger negligible states sets using the criterion $\text{best} \leftarrow \mathcal{S}^A$ if and only if $\widehat{\text{Cov}}(A, BG) < \epsilon$ and $\widehat{\text{MSE}}(\mathcal{S}^A) < \widehat{\text{MSE}}(\text{best})$ or $\widehat{\text{MSE}}(\mathcal{S}^A) < \widehat{\text{MSE}}(\text{best}) * (1 + \epsilon_s) \wedge |\mathcal{S}^A| > |\text{best}|$ based on user-defined parameters $\epsilon_s > 0$ and $\epsilon > 0$. In this paper, we use exhaustive search over subsets while it is also possible to form heuristic search algorithms.

B. Q-value based identification

In the Q-value based identification, a model of the MDP is formulated based on the trajectories from π_b to compute time-specific Q-values for π_e via stochastic dynamic programming:

$$\hat{Q}_t(s, a) = \hat{r}(s, a) + \sum_{s' \in \mathcal{S}} \hat{\mathcal{T}}(s, a, s') \pi_e(\cdot | s') \cdot \hat{Q}_{t+1}(s', \cdot), \quad (11)$$

where \cdot indicates vectorisation. The algorithm identifies a state $s \in \mathcal{S}$ as negligible if for all $t \in [H]$, $|\hat{Q}_t(s, a) - \hat{Q}_t(s, a')| < \epsilon$ for all $a, a' \in \mathcal{A}$. This criterion leads to an SIS estimator with limited absolute bias depending on the user-defined ϵ , the maximal number of dropped weights, and the model approximation error ϵ_2 (see Appendix C). Consistent with the analysis in Appendix B, the MSE of SIS is $\mathcal{O}(\exp(M_B))$ with both state identification methods. The Q-value based identification algorithm is computationally inexpensive and can include a large subset of the state space as negligible states. Moreover, it is easy to interpret and independent of the estimator making it applicable to a variety of estimators.

VI. STATE-BASED IMPORTANCE SAMPLING VARIANTS

While so far we have focused on ordinary importance sampling, it is easy to extend SIS to use more advanced off-policy evaluation estimators as its base-estimator.

First, methods such as PDIS and INCRIS, and even the weighted counterparts of these algorithms, are easily extended by analogously setting the action probability ratio to one for negligible states.

Second, for doubly robust off-policy evaluation [10], [17], one can use a model using the same formalism as the Q-value based identification. While traditional DR uses an importance weight $w_t^i = \frac{1}{n} \prod_{j=1}^t \frac{\pi_e(a_j^i | s_j^i)}{\pi_b(a_j^i | s_j^i)}$ for a trajectory i at time t and n total trajectories, doubly robust state-based importance sampling (DRSIS) uses the importance weight $w_t^i = \frac{1}{n} \prod_{j=1}^t \frac{\pi_e^A(a_j^i | s_j^i)}{\pi_b(a_j^i | s_j^i)}$ where π_e^A is the \mathcal{S}^A -reduced evaluation policy of π_e .

Third, techniques that rely on importance sampling over state visitation frequencies such as SDRE and MIS can also benefit from a state-based formulation. For example, the stationary density ratio estimator (2) can be turned into a state-based formulation according to

$$\hat{G}_{SSDRE} = \frac{1}{n} \sum_{i=1}^n \sum_{t=1}^H r_t^{(i)} \frac{\hat{d}^{\pi_e^A}(s_t^{(i)})}{\hat{d}^{\pi_b}(s_t^{(i)})}, \quad (12)$$

where $\hat{d}^{\pi_e^A}(s_t^{(i)})$ is equal to $\hat{d}^{\pi_e}(s_t^{(i)})$ for $s_t^{(i)} \in \mathcal{S}^B$ and $\hat{d}^{\pi_b}(s_t^{(i)})$ for $s_t^{(i)} \in \mathcal{S}^A$. Such techniques are polynomial in H due to replacing the product over action probability ratios by a sum of state visitation density ratios. As such the variance reduction due to reducing polynomial factors is smaller than variance reductions due to reducing the exponent, as is the case in traditional IS methods. However, as the experiments in Sec. VII show, this can still yield significant MSE improvements and an excellent overall performance in environments with stationary distributions.

VII. EXPERIMENTS

The experimental validation of state-based estimators comprises 4 domains. The first two domains include lift states in the form of a lift transporting the agent from one side to another, while the third and fourth domain are common RL benchmarks, namely the inventory management domain [28] and the taxi domain variant used in [7]. In each of the domains,

TABLE I
MEAN SQUARED ERROR (MSE) FOR OFF-POLICY EVALUATION ESTIMATORS OF THE EXPECTED RETURN IN THE DETERMINISTIC LIFT DOMAIN BASED ON 50 INDEPENDENT RUNS OF 1,000 EPISODES FOR EACH DOMAIN SIZE. THE BEST AND SECOND BEST ESTIMATORS ARE HIGHLIGHTED IN BOLD UNDERLINED AND BOLD, RESPECTIVELY, FOR EACH DOMAIN SIZE. ALL SIS ESTIMATORS SHOWN USE THE Q-VALUE BASED ALGORITHM FOR NEGLIGIBLE STATE SET IDENTIFICATION.

Domain size	IS	SIS	PDIS	SPDIS	INCRIS	SINCRIS	SDRE	SSDRE
7	0.0067	0.0019	0.0398	0.0121	4.0024	4.0007	532.7567	536.3787
9	0.0153	0.0019	0.1766	0.0215	6.1964	6.1765	577.3089	584.4657
11	0.0353	0.0029	0.7087	0.0600	8.9408	8.9225	674.0772	680.1959
13	0.0691	0.0026	1.8038	0.0688	12.5475	12.4779	652.6004	660.2387
15	0.1389	0.0033	5.3474	0.1378	16.1902	16.1558	803.3450	813.4997
17	0.3026	0.0031	15.1566	0.1303	23.7084	20.1534	758.2237	764.6283

TABLE II
MEAN SQUARED ERROR (MSE) FOR OFF-POLICY EVALUATION ESTIMATORS OF THE EXPECTED RETURN IN THE STOCHASTIC LIFT DOMAIN BASED ON 200 INDEPENDENT RUNS OF 1,000 EPISODES FOR EACH DOMAIN SIZE. THE BEST, SECOND BEST, AND THIRD BEST ESTIMATORS ARE HIGHLIGHTED IN BOLD DOUBLE-UNDERLINED, BOLD UNDERLINED, AND BOLD, RESPECTIVELY, FOR EACH DOMAIN SIZE.

Domain size	WIS	WSIS	WPDIS	WSPDIS	WINCRIS	WSINCRIS	SDRE	SSDRE	WDR	WDRSIS
7	10.2238	8.4832	6.8271	5.3853	1.0958	0.2618	208.7413	209.8354	5.2572	5.0844
9	9.4737	7.7326	5.3924	4.0441	0.2973	0.6929	243.1366	245.1004	2.8067	2.7486
11	9.7285	6.2044	5.7219	3.0552	0.2002	2.3050	250.3149	251.5843	2.3016	1.9838
13	8.4436	9.4967	4.4449	3.2916	0.8958	4.8960	266.7911	268.6634	1.3459	1.3354
15	18.0453	4.2831	6.2259	1.7004	2.6550	9.0162	280.6434	282.6849	1.2258	0.8797
17	8.2474	4.8129	5.2761	1.8151	4.7157	13.0056	305.7465	307.9191	0.8385	0.4820

we evaluate the expected return of a near-optimal policy π_e based on sub-optimal policy π_b .

The experiments validate state-based importance sampling when implemented on IS, PDIS, INCRIS, DR, and SDRE. The corresponding state-based implementations are denoted with the prefix **S**; for instance, SPDIS denotes state-based PDIS. State-based doubly robust off-policy evaluation is denoted as **DRSIS** to distinguish it from SDRE. With the exception of the first domain, which has short trajectory lengths, and SDRE for which this is not applicable, the experiments use additional variance reduction with weighted importance sampling, which is indicated by the prefix **W**; for instance, **WIS** denotes weighted IS. Further, the effectiveness of the negligible state identification is demonstrated on the first two domains as these contain known lift states (see Appendix D). Due to its ability to identify the lift states and yield lower MSE, the Q-based identification algorithm is used throughout this section. Additional experimental details are provided in Appendix E and source code is available at <https://github.com/bosssdm/ImportanceSampling>.

A. Lift domains

Based on the ϵ -negligibility definition, one would expect state-based techniques to be advantageous on domains where there are some states for which the action taken has no or limited impact on the return. This intuition is captured by the notion of “lift states”.

Definition 2. Lift states. A lift state $s \in \mathcal{S}$ is a state for which $\mathcal{T}(s, a) = \mathcal{T}(s, a')$ and $r(s, a) = r(s, a')$ for all $a, a' \in \mathcal{A}$. The set of lift states for a domain is denoted by \mathcal{S}^L .

a) Deterministic lift domain: The deterministic lift domain involves an agent moving on a straight line, with the state space $\mathcal{S} = \{-b, \dots, -1, 0, 1, \dots, b\}$ representing the coordinates along the line and the action space $\mathcal{A} = \{\text{left}, \text{right}\}$ representing left or right movements on the line. With the exception of the starting state and states within a distance of 1 from the terminal state, the domain has lift states which force the agent to move left on $s < 0$ and right on $s > 0$. The episode stops when the agent hits the bound, and then the terminal reward of $r = s$ is given (equal to either $-b$ or b). All non-terminal rewards are $r = -1$ to penalise the length of the path. The experiments manipulate this bound for an increasing effective horizon. An analysis comparing state-based estimators to their traditional counterparts (see Tab. I) demonstrates a near universal benefit in the deterministic domain, with the only exception being SDRE. SIS is the best performing estimator across all domain sizes, followed by IS and SPDIS depending on the domain size. SDRE and SSDRE have particularly poor estimates because the domain does not have a stationary distribution. Analysis of the residuals (see Fig. 1a in Appendix F) also shows that: a) state-based methods have successfully reduced the variance compared to their traditional counterparts; b) SPDIS strongly reduces the residual of PDIS for large domain sizes; and c) INCRIS and SINCRIS have large residuals, which implies that dropping recent time steps may not be suitable for the deterministic domain (e.g. the first time step always determines the final reward). Due to the deterministic domain being easy to model, the doubly robust methods have MSEs close to zero in this domain so are not further analysed.

TABLE III
 NORMALISED MEAN SQUARED ERROR (MSE) FOR OFF-POLICY EVALUATION ESTIMATORS OF THE EXPECTED RETURN IN THE INVENTORY MANAGEMENT DOMAIN BASED ON 50 INDEPENDENT RUNS FOR EACH NUMBER OF EPISODES. NORMALISATION IS BASED ON DIVISION BY THE SQUARED RANGE. THE BEST, SECOND BEST, AND THIRD BEST ESTIMATORS ARE HIGHLIGHTED IN BOLD DOUBLE-UNDERLINED, BOLD UNDERLINED, AND BOLD, RESPECTIVELY, FOR EACH NUMBER OF EPISODES.

Episodes	WIS	WSIS	WPDIS	WSPDIS	WINCRIS	WSINCRIS	SDRE	SSDRE	WDR	WDRSIS
100	0.0490	<u>0.0088</u>	0.0278	<u>0.0074</u>	0.0641	0.0109	0.0160	<u>0.0077</u>	0.0173	0.0138
250	0.0651	<u>0.0080</u>	0.0335	<u>0.0080</u>	0.0546	0.0118	0.0161	<u>0.0070</u>	0.0132	0.0137
500	0.0652	0.0113	0.0340	<u>0.0105</u>	0.0456	0.0124	0.0163	<u>0.0067</u>	<u>0.0098</u>	0.0114
1000	0.0820	0.0122	0.0293	0.0119	0.0313	0.0128	0.0162	<u>0.0066</u>	<u>0.0117</u>	<u>0.0104</u>

TABLE IV
 MEAN SQUARED ERROR (MSE) FOR OFF-POLICY ESTIMATORS OF THE EXPECTED RETURN IN THE TAXI DOMAIN BASED ON 20 INDEPENDENT RUNS FOR EACH EFFECTIVE HORIZON. THE BEST, SECOND BEST, AND THIRD BEST ESTIMATORS ARE HIGHLIGHTED IN BOLD DOUBLE-UNDERLINED, BOLD UNDERLINED, AND BOLD, RESPECTIVELY, FOR EACH EFFECTIVE HORIZON.

Effective horizon	WIS	WSIS	WPDIS	WSPDIS	WINCRIS	WSINCRIS	SDRE	SSDRE	WDR	WDRSIS
10	0.6080	<u>0.4525</u>	0.7234	0.5084	<u>0.4896</u>	<u>0.4767</u>	1.2764	1.2790	53.2003	53.1011
50	15.8227	<u>15.9281</u>	11.6239	<u>10.7777</u>	35.6350	35.8282	<u>9.4938</u>	<u>9.4435</u>	188.8620	187.3876
250	440.6435	473.4986	281.9518	<u>273.3647</u>	1753.7938	1752.6067	<u>25.3142</u>	<u>25.2920</u>	993.3100	1055.8603
1000	8233.7773	7948.8306	4876.1759	<u>4652.7797</u>	27675.6630	27635.4627	<u>231.6101</u>	<u>231.4313</u>	11982.1668	12176.5372

b) *Stochastic lift domain:* In the stochastic lift domain, the evaluation policy and the state transitions are stochastic rather than deterministic. The evaluation policy π_e takes the best action with probability $1 - \delta$ and the other action with probability δ , and the state transitions dynamics are formulated such that the action effect is reversed with probability δ . An analysis comparing state-based estimators to their traditional counterparts (see Tab. II) demonstrates a benefit in the stochastic domain. This time doubly robust methods are also analysed as the error of the DR estimator is now non-zero. WSIS and WSPDIS consistently outperform their traditional counterparts. For large domain sizes, WDRSIS is the best performing estimator followed by WDR and WSPDIS. State-based methods are not adversely affected by the domain size, which is explained by all the lift states being dropped and therefore the maximal number of steps M_B in non-negligible states remaining constant regardless of domain size. For small domain sizes, WINCRIS and WSINCRIS have the best performance, indicating the benefit of dropping time steps; the difference to the deterministic domain maybe due to the visitation of multiple lifts in the same episode. The lower performance of WINCRIS and WSINCRIS in larger domain sizes is attributed to the limited history window $k \in \{1, \dots, 10\}$, which is needed for computational feasibility but may not be sufficient for longer trajectories. Fig. 1b in Appendix F further demonstrates that the state-based methods have successfully reduced the variance compared to their traditional counterparts. As in the deterministic lift domain, the poor performance of SDRE estimators in the stochastic lift domain is attributed to the domain not having a defined stationary state distribution.

B. Inventory management

The task of the agent in the inventory management problem is to purchase items to make optimal profits selling the items. The agent must balance supply with demand while having at most S items in the inventory. The state is the current inventory while the action is the chosen number of items to add. States corresponding to an almost full inventory are similar to lift states in the sense that the actions chosen are unimportant. As the domain is high variance due to the varying demand, we opt for a relatively high number of 50 experiment runs and instead of manipulating the effective horizon we analyse the improvement in performance across different numbers of episodes. An analysis comparing state-based estimators to their traditional counterparts (see Tab. III) demonstrates a near universal benefit in the inventory management domain, with 2- to 8-fold improvements being observed (see Tab. III) – with the exception of DR which already has high accuracy. WSPDIS performs best for 100 episodes while SSDRE performs best for 250 to 1,000 episodes. The high performance of SSDRE is attributed to a clearly defined stationary distribution. The performance of SSDRE and SDRSIS dependent on the sample size suggests that these state-based methods provide improved convergence and efficiency. Other high-performing algorithms are WSIS, WSINCRIS, WDR, and WDRSIS. Fig. 2 in Appendix F demonstrates that state-based techniques reduce the variance and yield a smaller residual.

C. Taxi

The taxi domain has a much larger state space with $|\mathcal{S}| = 2000$ and can have infinite horizon. A taxi agent has to pick up and drop off passengers who arrive at random corners in a 5×5 grid based on 6 actions, namely going a step north, west, south, east, picking up a passenger, and

dropping off a passenger. The taxi agent receives a reward of 20 upon a drop-off or pick-up at the correct location, and a reward of -1 otherwise. The domain is of interest to investigate the scalability to large state-action spaces and large effective horizon, and it is of interest to manipulate the effective horizon in an additional domain which in contrast to the lift domains has a stationary distribution. The results for the taxi domain (see Tab. IV) also show a near universal benefit for state-based estimators compared to the traditional counterparts, although the effect is not as pronounced as for example in the inventory management domain. For an effective horizon of 10, WSIS is the highest performing followed by WSINCRIS and WINCRIS. For effective horizons between 50 and 1,000, SSDRE is the highest performer followed by SDRE and WSPDIS. The high performance of SDRE based methods, and their relative improvement across effective horizon, is attributed to their MSE being polynomial in the effective horizon whereas the MSE of traditional importance sampling techniques based on products of action probability ratios is exponential in the effective horizon. As illustrated in Fig. 3 in Appendix F, the SDRE methods have a lower variance than WDR methods and their residuals do not depend as strongly on the effective horizon as WIS, WPDIS, and WINCRIS.

VIII. CONCLUSION

Off-policy reinforcement learning comes with the significant challenge of high-variance estimates for the expected return. This paper introduces state-based importance sampling, a class of off-policy evaluation techniques which reduces the variance of importance sampling by eliminating states that do not affect the return from the importance weight computation. The state-based technique is shown to improve performance when integrated with a variety of off-policy evaluation methods, making it a useful plug-in for many applications. Future work directions include theoretical studies on the presence of negligible states, implementing policy optimisation techniques based on state-based importance sampling, and exploring the links with other off-policy evaluation techniques.

ACKNOWLEDGEMENT

We thank Shreyas Chaudhari for comments on a draft of the paper.

REFERENCES

- [1] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, vol. 9. MIT Press, 1998.
- [2] T. Udagawa, H. Kiyohara, Y. Narita, Y. Saito, and K. Tateno, "Policy-Adaptive Estimator Selection for Off-Policy Evaluation," in *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI 2023)*, vol. 37, pp. 10025–10033, nov 2023.
- [3] S. Levine, A. Kumar, G. Tucker, and J. Fu, "Offline Reinforcement Learning: Tutorial, Review, and Perspectives on Open Problems," *arXiv preprint*, pp. 10–16, may 2020.
- [4] D. Precup, R. S. Sutton, and S. P. Singh, "Eligibility Traces for Off-Policy Policy Evaluation," in *Proceedings of the International Conference on Machine Learning (ICML 2000)*, pp. 759–766, 2000.
- [5] P. S. Thomas, G. Theodorou, and M. Ghavamzadeh, "High confidence off-policy evaluation," in *Proceedings of the National Conference on Artificial Intelligence (AAAI 2015)*, pp. 3000–3006, 2015.
- [6] Y. Chandak, S. M. Jordan, G. Theodorou, M. White, and P. S. Thomas, "Towards safe policy improvement for non-stationary MDPs," in *Advances in Neural Information Processing Systems (NeurIPS 2020)*, pp. 1–13, 2020.
- [7] Q. Liu, Z. Tang, L. Li, and D. Zhou, "Breaking the curse of horizon: Infinite-horizon off-policy estimation," in *Advances in Neural Information Processing Systems (NeurIPS 2018)*, pp. 5356–5366, 2018.
- [8] Z. D. Guo, P. S. Thomas, and E. Brunskill, "Using options and covariance testing for long horizon off-policy policy evaluation," in *Advances in Neural Information Processing Systems (NeurIPS 2017)*, pp. 2493–2502, 2017.
- [9] D. Precup, R. S. Sutton, and S. Dasgupta, "Off-policy temporal-difference learning with function approximation," in *Proceedings of the International Conference on Machine Learning (ICML 2001)*, pp. 417–424, 2001.
- [10] P. S. Thomas and E. Brunskill, "Data-efficient off-policy policy evaluation for reinforcement learning," in *Proceedings of the International Conference on Machine Learning (ICML 2016)*, pp. 3158–3167, 2016.
- [11] T. Popoviciu, "Sur les équations algébriques ayant toutes leurs racines réelles," *Mathematica*, pp. 129–145, 1935.
- [12] Y. Liu, P. L. Bacon, and E. Brunskill, "Understanding the curse of horizon in off-policy evaluation via conditional importance sampling," in *Proceedings of the International Conference on Machine Learning (ICML 2020)*, pp. 6140–6149, 2020.
- [13] T. Xie, Y. Ma, and Y. X. Wang, "Towards optimal off-policy evaluation for reinforcement learning with marginalized importance sampling," in *Advances in Neural Information Processing Systems (NeurIPS 2019)*, pp. 1–11, 2019.
- [14] M. Yin and Y.-X. Wang, "Asymptotically Efficient Off-Policy Evaluation for Tabular Reinforcement Learning," in *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS 2020)*, pp. 1–10, 2020.
- [15] M. Rowland, A. Harutyunyan, H. van Hasselt, D. Borsa, T. Schaul, R. Munos, and W. Dabney, "Conditional Importance Sampling for Off-Policy Learning," in *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS 2020)*, pp. 1–10, 2020.
- [16] M. Dudik, J. Langford, and L. Li, "Doubly Robust Policy Evaluation and Learning," in *Proceedings of the International Conference on Machine Learning (ICML 2011)*, pp. 1–8, 2011.
- [17] N. Jiang and L. Li, "Doubly robust off-policy value evaluation for reinforcement learning," in *Proceedings of the International Conference on Machine Learning (ICML 2016)*, pp. 1022–1035, 2016.
- [18] M. Farajtabar, Y. Chow, and M. Ghavamzadeh, "More Robust Doubly Robust Off-policy Evaluation," in *Proceedings of the International Conference on Machine Learning (ICML 2018)*, pp. 2342–2371, 2018.
- [19] H. Bang and J. M. Robins, "Doubly robust estimation in missing data and causal inference models," *Biometrics*, vol. 61, no. 4, pp. 962–973, 2005.
- [20] R. S. Sutton, D. Precup, and S. Singh, "Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning," *Artificial Intelligence*, vol. 112, no. 1, pp. 181–211, 1999.
- [21] A. B. Owen, *Monte Carlo theory, methods and examples*. 2013.
- [22] A. M. Metelli, M. Papini, N. Montali, and M. Restelli, "Importance sampling techniques for policy optimization," *Journal of Machine Learning Research*, vol. 21, pp. 1–75, 2020.
- [23] G. Theodorou, P. S. Thomas, and M. Ghavamzadeh, "Personalized Ad Recommendation Systems for Life-Time Value Optimization with Guarantees," in *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI 2015)*, pp. 1806–1812, 2015.
- [24] P. S. Thomas, G. Theodorou, and M. Ghavamzadeh, "High confidence policy improvement," in *Proceedings of the International Conference on Machine Learning (ICML 2015)*, pp. 2370–2378, 2015.
- [25] P. S. Thomas, G. Theodorou, and M. Ghavamzadeh, "High confidence off-policy evaluation," in *Proceedings of the National Conference on Artificial Intelligence (AAAI 2015)*, pp. 3000–3006, 2015.
- [26] A. M. Metelli, M. Papini, F. Francesco, and M. Restelli, "Policy Optimisation via Importance Sampling," in *Proceedings of the International Conference on Machine Learning (ICML 2019)*, pp. 1–13, 2019.
- [27] M. Yin, Y. Bai, and Y. X. Wang, "Near-Optimal Provable Uniform Convergence in Offline Policy Evaluation for Reinforcement Learning," in *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS 2021)*, vol. 130, pp. 1567–1575, 2021.
- [28] M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. 2005.

Supplementary Information for *Low Variance Off-policy Evaluation with State-based Importance Sampling*

David M. Bossens

Philip S. Thomas

APPENDIX A: PROOF OF $\mathbb{E}[A] = 1$

Denoting P_1 as the starting distribution, $\tau := s_1, a_1, \dots, s_H, a_H$, $\tau_s := s_1, \dots, s_H$, and $\tau_a := a_1, \dots, a_H$ and π_e^A as the S^A -reduced evaluation policy of π_e , we have

$$\begin{aligned} \mathbb{E}_{\pi_b}[A] &= \int_{\tau} \pi_e^A(\tau_a|\tau_s)/\pi_b(\tau_a|\tau_s)\mathbb{P}(\tau_s|\pi_b)d\tau \\ &= \int_{\tau} \pi_e^A(\tau_a|\tau_s)/\pi_b(\tau_a|\tau_s) \\ &\quad P_1(s_1) \prod_{t=1}^H \pi_b(a_t|s_t)\mathcal{T}(s_t, a_t, s_{t+1}) \\ &= \int_{\tau} \pi_e^A(\tau_a|\tau_s)P_1(s_1) \prod_{t=1}^H \mathcal{T}(s_t, a_t, s_{t+1})d\tau \\ &= \int_{\tau} \mathbb{P}(\tau_s|\pi_e^A)d\tau \\ &= 1. \end{aligned}$$

APPENDIX B: ANALYSIS OF MSE UNDER INACCURATE STATE-SET IDENTIFICATION

The variance reduction of state-based methods depends crucially on the state-set identification. In covariance testing, the covariance estimate may be highly inaccurate, and in Q-value based identification, the model on which the Q-value is based may be highly inaccurate. Below we compare the true MSE of the IS estimator to that of the covariance testing based SIS estimator under the assumption that the state-identification algorithm has identified a state-set that is not negligible at all but instead has an arbitrary covariance with the return estimate (thereby yielding a significant bias in the estimate).

Note that the true MSE of ordinary importance sampling is upper bounded by

$$\begin{aligned} \text{MSE}(\hat{G}_{IS}) &\leq \frac{1}{4n} (Hr_{\max}\rho_{\max}^H)^2 \\ &= \mathcal{O}(\exp(H)). \end{aligned}$$

While for SIS the chosen state-set has been observed to be ϵ -negligible and the covariance and variance have been estimated, for the sake of the argument assume that the population covariance and variance may take on any value. Via Cauchy-Schwartz inequality and Popoviciu's inequality with $G \in [0, Hr_{\max}]$ and $A, B \in [0, \rho_{\max}]$, we have

$$|\text{Cov}(A, BG)| \leq \text{Var}(A)\text{Var}(BG) \leq \frac{\rho_{\max}^4 H^2 r_{\max}^2}{16}.$$

With the bias equal to the covariance, it follows via bias-variance decomposition that

$$\begin{aligned} \text{MSE}(\hat{G}_{SIS}(S^A)) &= \text{Var}(\hat{G}_{SIS}(S^A)) + \text{Cov}(A, BG)^2 \\ &\leq \frac{1}{4n} (Hr_{\max}\rho_{\max}^{M_B})^2 + \frac{\rho_{\max}^8 H^4 r_{\max}^4}{256} \\ &= \mathcal{O}(\exp(M_B)). \end{aligned}$$

Therefore the SIS has a significantly improved upper bound on the true MSE provided the maximal number of steps in S^B is significantly smaller than H . If, however, the state-set identification was accurate, such that $\text{Cov}(A, BG) < \epsilon$, then any reduction in variance implies via the bias-variance decomposition that SIS has successfully optimised the true MSE compared to IS.

APPENDIX C: PROOF OF LIMITED BIAS OF Q-BASED STATE IDENTIFICATION

Let $s_t \in S^A = \{s \in S : |\hat{Q}_{t'}(s, a) - \hat{Q}_{t'}(s, a')| < \epsilon \quad \forall a, a' \in \mathcal{A} \quad \forall t' \in [H]\}$. This implies that the SIS estimator has the following expected return from starting state s_t

$$\begin{aligned} \mathbb{E}_{\pi_b}[\hat{G}_{SIS}|s_t] &= \mathbb{E}_{\pi_b}[\rho_{t+1:t+H}G_t|s_t] \\ &= \sum_a \pi_b(a|s_t) (r(s_t, a) + \mathbb{E}_{\pi_b}[\rho_{t+1:t+H}G_{t+1}|s_t]) \\ &= \sum_a \pi_b(a|s_t)Q_t(s_t, a) \\ &\approx \sum_a \pi_e(a|s_t)Q_t(s_t, a) \\ &= \mathbb{E}_{\pi_e}[G_t|s_t], \end{aligned}$$

where the approximate equality follows since $|\hat{Q}_t(s, a) - \hat{Q}_t(s, a')| < \epsilon$ implies $Q_t(s, a) \approx Q_t(s, a')$ for all $a, a' \in \mathcal{A}$. Under a perfect model, the bias of the estimator comes only from the approximate equality, in which case the absolute bias induced by state being dropped is at most ϵ . Under an erroneous model, which yields $|\hat{Q}_t(s, a) - Q_t(s, a)| < \epsilon_2$ for all t and all $a \in \mathcal{A}$, the absolute bias induced by a state being dropped is at most $\epsilon + \epsilon_2$. Repeatedly applying this argument for a maximal number of M_A dropped weights in a trajectory, the absolute bias induced is at most $M_A(\epsilon + \epsilon_2)$. Filling in $\mathcal{O}(\text{poly}(H))$ as the bias and $\frac{1}{4n} (Hr_{\max}\rho_{\max}^{M_B})^2$ as the variance in the bias-variance decomposition, the MSE is bounded by $\mathcal{O}(\exp(M_B))$. \square

APPENDIX D: COMPARISON OF STATE SET IDENTIFICATION ALGORITHMS

The comparison for traditional importance sampling includes **SIS (Lift states)**, which uses the lift states of the domain (i.e., where $S^A = S^L$); **SIS (Covariance testing)**, which uses the negligible states identified by covariance testing; and **SIS (Q-based)**, which uses the negligible states identified using Q-values. Similarly, we evaluate the effectiveness of state identification for the doubly robust off-policy evaluation method, which are analogously denoted as **SDR (Lift states)**, **SDR (Covariance testing)**, and **SIS (Q-based)**.

Analysing the different negligible state set identification methods reveals that SIS (Q-based) identifies the lift states as negligible states and therefore yields the same performance as SIS (Lift states). Tab. I demonstrates the superior performance of SIS (Q-based) and SIS (Lift states) compared to IS, which never drops any states, and SIS (Covariance testing), which typically drops two lift states (the maximal cardinality given the computational budget).

TABLE I
MEAN SQUARED ERROR (MSE) COMPARING STATE-BASED IMPORTANCE SAMPLING METHODS WITH DIFFERENT NEGLIGIBLE STATE SETS ON 50 INDEPENDENT RUNS OF THE DETERMINISTIC LIFT DOMAIN FOR EACH DOMAIN SIZE. ESTIMATES ARE OBTAINED BASED ON 1,000 EPISODES. THE BEST AND SECOND BEST ESTIMATORS ARE HIGHLIGHTED IN BOLD UNDERLINED AND BOLD, RESPECTIVELY, FOR EACH NUMBER OF EPISODES; IN CASE OF A TIE ONLY THE BEST ESTIMATORS ARE HIGHLIGHTED IN BOLD UNDERLINED.

	IS	SIS (Lift states)	SIS (Covariance testing)	SIS (Q-based)
Domain size				
7	0.0067	<u>0.0019</u>	0.0056	<u>0.0019</u>
9	0.0153	<u>0.0019</u>	0.0138	<u>0.0019</u>
11	0.0353	<u>0.0029</u>	0.0345	<u>0.0029</u>
13	0.0691	<u>0.0026</u>	0.0705	<u>0.0026</u>
15	0.1389	<u>0.0033</u>	0.1386	<u>0.0033</u>
17	0.3026	<u>0.0031</u>	0.3035	<u>0.0031</u>

Results for state set identification are analogous to the deterministic lift domain: WSIS (Q-based) is successful at identifying the lift states as negligible states and together with WSIS (Lift states) has a superior performance compared to WSIS (Covariance testing) and WIS (see Tab. II; and WSIS (Covariance testing) typically drops two lift states (the maximal cardinality given the computational budget). Tab. III demonstrates how dropping states is also beneficial for doubly robust methods and again the Q-based state identification works to identify the lift states.

TABLE II

MEAN SQUARED ERROR (MSE) COMPARING STATE-BASED IMPORTANCE SAMPLING METHODS WITH DIFFERENT NEGLIGIBLE STATE SETS ON 200 INDEPENDENT RUNS OF THE STOCHASTIC LIFT DOMAIN FOR EACH DOMAIN SIZE. ESTIMATES ARE OBTAINED BASED ON 1,000 EPISODES. THE BEST AND SECOND BEST ESTIMATORS ARE HIGHLIGHTED IN BOLD UNDERLINED AND BOLD, RESPECTIVELY, FOR EACH DOMAIN SIZE; IN CASE OF A TIE ONLY THE BEST ESTIMATORS ARE HIGHLIGHTED IN BOLD UNDERLINED.

	WIS	WSIS (Lift states)	WSIS (Covariance testing)	WSIS (Q-based)
Domain size				
7	10.2238	<u>8.4832</u>	10.2186	<u>8.4832</u>
9	9.4737	<u>7.7326</u>	9.4489	<u>7.7326</u>
11	9.7285	<u>6.2044</u>	9.6775	<u>6.2044</u>
13	<u>8.4436</u>	9.5003	<u>8.4058</u>	9.4967
15	18.0453	<u>4.2803</u>	17.9633	<u>4.2831</u>
17	8.2474	<u>4.7482</u>	7.7646	<u>4.8129</u>

TABLE III

MEAN SQUARED ERROR (MSE) COMPARING DOUBLY ROBUST METHODS WITH DIFFERENT NEGLIGIBLE STATE SETS ON 200 INDEPENDENT RUNS OF THE STOCHASTIC LIFT DOMAIN FOR EACH DOMAIN SIZE. ESTIMATES ARE OBTAINED BASED ON 1,000 EPISODES. THE BEST AND SECOND BEST ESTIMATORS ARE HIGHLIGHTED IN BOLD UNDERLINED AND BOLD, RESPECTIVELY, FOR EACH DOMAIN SIZE; IN CASE OF A TIE ONLY THE BEST ESTIMATORS ARE HIGHLIGHTED IN BOLD UNDERLINED.

	WDR	WDRSIS (Lift states)	WDRSIS (Covariance testing)	WDRSIS (Q-based)
Domain size				
7	5.2572	<u>5.0844</u>	5.2662	<u>5.0844</u>
9	2.8067	<u>2.7486</u>	2.8235	<u>2.7486</u>
11	2.3016	<u>1.9838</u>	2.3058	<u>1.9838</u>
13	1.3459	<u>1.3359</u>	1.3612	<u>1.3354</u>
15	1.2258	<u>0.8788</u>	1.1675	<u>0.8797</u>
17	0.8385	<u>0.4817</u>	0.6620	<u>0.4820</u>

APPENDIX E: EXPERIMENTAL DETAILS


In the experiments, covariance testing is applied with $\epsilon_s = \epsilon = 0.01$ and rather than the full power set, we limit the search to sets with cardinalities of at most 2. The Q-value based identification is suitable for all the domains and the parameter settings are $\epsilon = 1$ in the grid world lift domains, $\epsilon = 50$ in the inventory management domain, and $\epsilon = 2$ in the taxi domain.

In the lift domains, π_e is the optimal policy and π_b is the uniform random policy. The expected return is 1.0 in the deterministic domain and negative in the stochastic domain (with magnitude increasing with domain size). The failure probability is set to $\delta = 0.05$ for the stochastic domain.

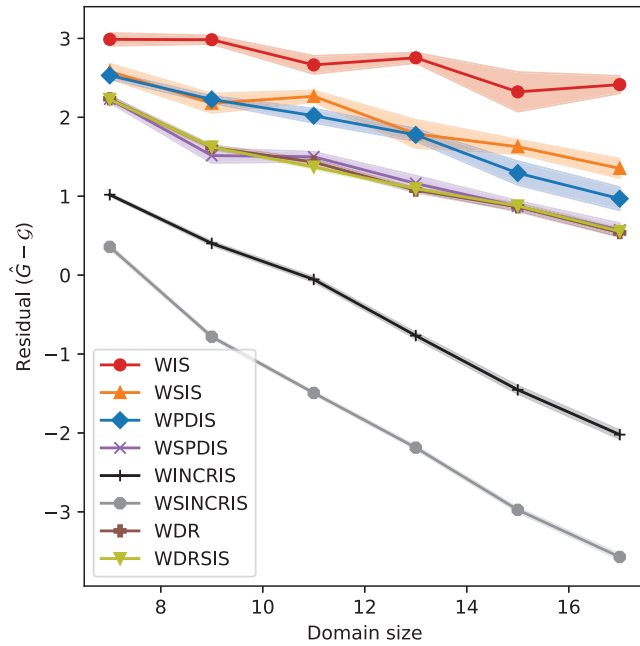
In inventory management, the purchasing cost is set to 2.49, the sale price to 3.99, and the holding cost to 0.03. Demand for items follows a Gaussian with mean $S/4$ and standard deviation $S/6$, and $S = 10$ in the experiments. Each episode consists of 100 steps, and consistent with the rest of our experiments an undiscounted setting ($\gamma = 1$) is considered. The evaluation policy, π_e , was trained for 500 episodes on the MDP and has an expected return of $\mathcal{G} = 1471.1041$. The behaviour policy, π_b , is uniform random.

In taxi, the implementation of the domain is based on code of [1] (see <https://github.com/zt95/infinite-horizon-off-policy-estimation/tree/master/taxi>). The evaluation policy represents a policy formed after 1,000 iterations of Q-learning and the behaviour policy represents a policy formed after 950 iterations of Q-learning. The expected return is positive and increases with effective horizon.

APPENDIX F: RESIDUAL PLOTS



(a) Deterministic lift domain



(b) Stochastic lift domain lift domain

Fig. 1. Residuals (y -axis), defined as $\hat{G} - \mathcal{G}$, as a function of the domain size (x -axis) in the lift domains. Estimates are based on 1,000 episodes. Residuals are represented by their mean \pm standard error over 50 independent runs for the deterministic lift domain and over 200 independent runs for the stochastic domain. Note that SDRE and SSDRE are not included for improved visibility as their residuals are extremely large.




Fig. 2. Residuals (y -axis), defined as $\hat{G} - \mathcal{G}$, as a function of the number of episodes (x -axis) in the inventory management domain. Residuals are represented by their mean \pm standard error over 50 independent runs.

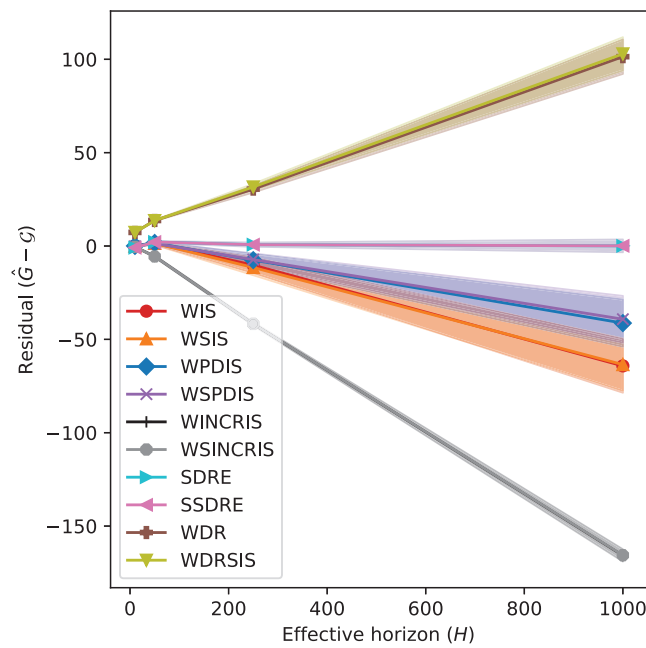


Fig. 3. Residuals (y -axis), defined as $\hat{G} - \mathcal{G}$, as a function of the effective horizon (H ; x -axis) of the taxi domain. Residuals are represented by their mean \pm standard error over 20 independent runs.

REFERENCES

- [1] Q. Liu, Z. Tang, L. Li, and D. Zhou, "Breaking the curse of horizon: Infinite-horizon off-policy estimation," in *Advances in Neural Information Processing Systems (NeurIPS 2018)*, pp. 5356–5366, 2018.