

---

# Abstract Reward Processes: Leveraging State Abstraction for Consistent Off-Policy Evaluation

---

**Shreyas Chaudhari**  
University of Massachusetts  
schaudhari@cs.umass.edu

**Ameet Deshpande**  
Princeton University  
asd@cs.princeton.edu

**Bruno Castro da Silva**  
University of Massachusetts  
bsilva@cs.umass.edu

**Philip S. Thomas**  
University of Massachusetts  
pthomas@cs.umass.edu

## Abstract

Evaluating policies using off-policy data is crucial for applying reinforcement learning to real-world problems such as healthcare and autonomous driving. Previous methods for *off-policy evaluation* (OPE) generally suffer from high variance or irreducible bias, leading to unacceptably high prediction errors. In this work, we introduce STAR, a framework for OPE that encompasses a broad range of estimators—which include existing OPE methods as special cases—that achieve lower mean squared prediction errors. STAR leverages state abstraction to distill complex, potentially continuous problems into compact, discrete models which we call *abstract reward processes* (ARPs). Predictions from ARPs estimated from off-policy data are provably consistent (asymptotically correct). Rather than proposing a specific estimator, we present a new framework for OPE and empirically demonstrate that estimators within STAR outperform existing methods. The best STAR estimator outperforms baselines in all twelve cases studied, and even the median STAR estimator surpasses the baselines in seven out of the twelve cases.

## 1 Introduction

Within *reinforcement learning* (RL), *off-policy evaluation* (OPE) is the foundational challenge of evaluating the performance,  $J(\pi)$ , of policies  $\pi$  that are different from the ones used to generate data. OPE methods are a general-purpose tool that can be used as part of a local policy search algorithm [45] to provide insight about policies similar to the current policy, or as a tool to evaluate policies without requiring their actual deployment for high-risk applications like those in healthcare [38], education [35, 15], and recommendation systems [5, 7]. Despite many recent advances in OPE, existing methods struggle to give accurate predictions for many real-world applications [52], showing the need for new perspectives on OPE.

Existing methods can be broadly divided into two categories: *importance sampling* (IS) based and model-based [58]. IS-based methods are typically consistent (i.e., their predictions converge probabilistically to the correct value in the limit as the amount of data approaches infinity), but have variance that increases exponentially with the horizon [29, 31]. Model-based methods have lower variance but often introduce bias due to model class mismatch and are not generally guaranteed to be consistent [36, 10]. A third set of methods, which we call *mixture methods*, combine the predictions obtained from both of these categories [22, 53]. However, in some cases, combining the predictions also combines the drawbacks—high variance and bias. This leads us to ask: *Can we develop a framework for OPE that yields predictions that are both consistent and low variance?*

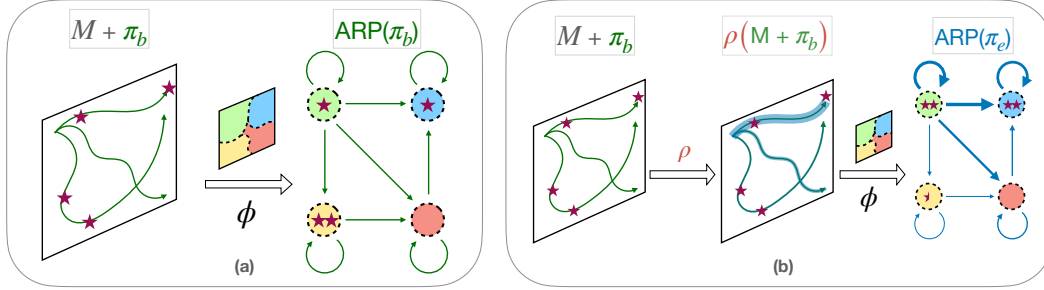


Figure 1: **(a)**: MDP  $M$  and policy  $\pi_b$  are transformed into a discrete *abstract reward process* (ARP) using a state abstraction function  $\phi$ . The ARP aggregates rewards (denoted by stars) and transition probabilities from all states that map to each abstract state. **(b)**: A model of the ARP for the evaluation policy  $\pi_e$  is constructed by: reweighting data generated by  $\pi_b$  with importance weights  $\rho$  (middle), applying the state abstraction function  $\phi$ , and performing weighted maximum likelihood estimation of the ARP (right). The expected return of a model of this ARP estimated from off-policy data is a *consistent estimator* of the expected return of  $\pi_e$ .

In this paper, we introduce a new framework that attains this goal by combining the *machinery* underlying IS-based and model-based approaches (not just their *predictions*). Our proposed framework is a fundamentally different approach to OPE that incorporates importance sampling *into* model learning for OPE. Our approach is motivated by the intuition that humans build small mental models of their environment to plan and predict, selectively abstracting away information that is not relevant to the problem at hand [54]. Similarly, complex sequential decision processes can be distilled into compact models that hold sufficient information for (off-)policy evaluation. Specifically, we propose creating small tabular models (even for continuous environments), which we call *abstract reward processes* (ARPs), customized for the problem at hand and for the policy being evaluated. We call this framework for constructing a range of ARPs *state-abstract reward processes* (STAR).

**Idea Summary:** A *Markov decision process* (MDP) combined with a policy  $\pi$  induces a Markov chain with rewards, called a *Markov reward process* (MRP). A model of this process can be estimated to evaluate policy  $\pi$ . However, two main challenges arise: (1) like any model-based approach, estimating an MRP can introduce asymptotic bias if the chosen model class cannot represent the underlying MRP; and (2) the model of the MRP must be accurately estimated from *off-policy* data, i.e., data generated by a behavior policy  $\pi_b$  that differs from the evaluation policy  $\pi_e$ . The proposed framework addresses both challenges by modeling a special instantiation of an MRP, as detailed next.

First, to address potential model class mismatch, we propose mapping the large and possibly continuous set of states of the MDP to a finite set of *abstract states* using a discrete state abstraction function. We then represent the resulting MRP defined over abstract states, referred to as the *abstract reward process* (ARP), using tabular models. Since the ARP is finite, tabular models can represent it accurately, as depicted in Figure 1(a). However, using a discrete state abstraction may lead to loss of state information that can potentially introduce modeling errors. Surprisingly, we prove that despite state information being abstracted away, the maximum likelihood model of an ARP estimated from *on-policy* data provides consistent estimates of the expected return of the behavior policy  $\pi_b$ .

Next, to estimate a model of the ARP corresponding to  $\pi_e$  from data generated by  $\pi_b$ , we reweight occurrences of abstract states in the off-policy dataset using importance sampling, see Figure 1(b). In expectation, this has the effect of updating the abstract state visitation counts to reflect those resulting from the policy being evaluated. We prove that the weighted maximum likelihood estimate of a model of the ARP estimated from this off-policy dataset provides consistent estimates of the performance of the evaluation policy. The integration of importance sampling *into* model estimation permits a favorable interpretation of weight clipping for mitigating variance (see Section 4.1).

The STAR framework offers two adjustable knobs—the state abstraction function, and the amount of weight clipping—that instantiate a range of OPE estimators. Varying the configurations of these knobs results in different bias-variance trade-offs for OPE, with existing OPE methods forming special cases of the range of estimators that lie within this framework.

**Contributions:** We empirically evaluate estimators instantiated in this framework on synthetic domains and a healthcare simulator built from real-world ICU data, where the best STAR estimator significantly outperforms baselines in all cases, and even the median STAR estimator surpasses baselines in seven out of twelve cases. It must be emphasized that this work does not propose a specific estimator for OPE; rather, it introduces a fundamentally different framework that offers fresh insights on approaches for off-policy evaluation. These insights open up exciting avenues for new research questions and future directions. In this paper, we introduce:

- (1) the first model-based approach for OPE that guarantees asymptotic correctness of the estimates without model class assumptions, even for continuous state MDPs (Theorem 4.1).
- (2) the concept of *abstract reward processes* for consistent OPE. ARPs abstract away the complexity of the underlying problem, and distill sufficient information for accurate policy evaluation (Theorem 3.1). Being finite, they can be consistently estimated.
- (3) a generalizing framework that provides a fresh perspective on OPE by merging the machinery of model-based and IS-based approaches. The framework offers two tunable knobs, various configurations of which instantiate a range of OPE estimators with varying bias-variance characteristics. Existing model-free and model-based methods are special cases in this framework.

## 2 Background and Notation

An MDP is a tuple  $M := (\mathcal{S}, \mathcal{A}, p, r, \gamma, \eta)$  where  $\mathcal{S}$  is the set of states,  $S_t$  is the state at time  $t \in \{0, 1, \dots\}$ ,  $\mathcal{A}$  is the set actions,  $A_t$  is the action at time  $t$ ,  $p : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$  is the *transition function* that characterizes state transition dynamics according to  $p(s, a, s') := \Pr(S_{t+1}=s' | S_t=s, A_t=a)$ ,  $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$  is the *reward function* that characterizes rewards according to  $r(s, a) := \mathbb{E}[R_t | S_t=s, A_t=a]$ ,  $\gamma \in [0, 1]$  is the reward discount parameter, and  $\eta : \mathcal{S} \rightarrow [0, 1]$  characterizes the initial state distribution according to  $\eta(s) := \Pr(S_0=s)$ .<sup>1</sup> A policy  $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$  characterizes how actions can be selected given the current state according to  $\pi(s, a) := \Pr(A_t=a | S_t=s)$ . We consider finite horizon MDPs [49] where episodes terminate by some (unspecified) time  $T \in \mathbb{N}$ —which is common in practical applications of OPE. For simplicity, we set  $\gamma = 1$ , allowing us to omit  $\gamma$  terms.

For OPE, a dataset  $\mathcal{D}_n^{(\pi_b)}$  is collected by deploying a behavior policy  $\pi_b$  on the MDP  $M$ . The dataset of  $n$  logged trajectories is denoted by  $\mathcal{D}_n^{(\pi_b)} := \{H^i\}_{i=1}^n$  where each  $H^i := (S_0^i, A_0^i, R_0^i, S_1^i, \dots)$  represents an independent trajectory generated by executing  $\pi_b$ . The performance of an evaluation policy  $\pi_e$  is its expected return, denoted by<sup>2</sup>

$$J(\pi_e) := \mathbb{E} \left[ \sum_{t=1}^T R_t; \pi_e \right]. \quad (1)$$

The problem of off-policy evaluation entails estimating  $J(\pi_e)$  with access only to data  $\mathcal{D}_n^{(\pi_b)}$ , generated by a behavior policy  $\pi_b$ , without additional interaction with the MDP. To ensure that samples in  $\mathcal{D}_n^{(\pi_b)}$  are sufficiently informative, we make the common assumption that any outcome under  $\pi_e$  has non-negligible probability of occurring under  $\pi_b$ .

**Assumption 2.1.** There exists an (unknown)  $\varepsilon > 0$  such that for all  $s \in \mathcal{S}$  and  $a \in \mathcal{A}$ ,  $(\pi_b(s, a) < \varepsilon) \implies (\pi_e(s, a) = 0)$ .

**Background:** For a detailed review of OPE methods, we refer the reader to surveys by Voloshin et al. [58] and Uehara et al. [56]. Concepts fundamental to this approach are briefly introduced here.

1. **Importance Sampling:** Importance sampling [25] enables unbiased estimation of the expected value,  $\mathbb{E}[f(X)]$ , of a function  $f$  applied to a random variable  $X \sim p$ , given samples of a different random variable  $Y \sim q$ . The importance sampling estimator is  $(p(Y)/q(Y))f(Y)$ , where

<sup>1</sup>For simplicity, our notation assumes that states, actions, and rewards are discrete random variables allowing for discussion of probabilities, rather than densities or measure theoretic probability. However, the methods proposed in this work extend to MDPs with continuous states, actions, and rewards. Furthermore, although we focus on OPE for MDPs, the method that we propose also applies to *partially observable* MDPs (POMDPs).

<sup>2</sup>We write  $\pi$  within statements of probability or expectations to indicate that random variables like  $S_t, A_t$ , and  $R_t$  result from the use of policy  $\pi$ .

$p(Y)/q(Y)$  is a term called an **importance weight**. This technique can provide unbiased estimates (i.e.,  $\mathbb{E}[(p(Y)/q(Y))f(Y)] = \mathbb{E}[f(X)]$ ) and has proven effective for variance reduction in Monte Carlo sampling [44] and for model-free OPE in RL [42].

2. **State Abstraction:** State abstraction aims to reduce the size of the state space by grouping together similar states in a way that does not change the essence of the underlying problem [28, 43, 1]. A state abstraction function  $\phi : \mathcal{S} \rightarrow \mathcal{Z}$  lies in the set of functions  $\phi \in \Phi$  that map each state  $s \in \mathcal{S}$  to an abstract state  $z \in \mathcal{Z}$ . We consider abstraction functions that partition  $\mathcal{S}$  into disjoint sets, where  $\mathcal{Z}$  is a finite set.

**Notation:** Indicator functions are abbreviated for clarity. For example,  $\mathbf{1}_t^i\{z, z'\} := \mathbf{1}\{\phi(S_{t+1}^{(i)}) = z', \phi(S_t^{(i)}) = z\}$  denotes the occurrence of abstract states  $z$  and  $z'$  at time steps  $t$  and  $t + 1$  in the  $i^{\text{th}}$  logged trajectory, with  $\mathbf{1}_t^i\{z\}$  defined correspondingly. The expected return of  $\pi$  when obtained from  $O$ —where  $O$  may be the MDP, or an ARP—is denoted by  $J(\pi; O)$ . The sample estimate of a variable  $y$  estimated from  $n$  samples is denoted by  $\hat{y}_n$ . Summation limits are often dropped for brevity, with  $\sum_t$  denoting  $\sum_{t=0}^T$  and  $\sum_i$  denoting  $\sum_{i=1}^n$ .

## 2.1 Markov Reward Processes

A *Markov reward process* (MRP) extends the idea of a Markov chain by associating states with rewards. Formally, an MRP is a tuple  $(\mathcal{X}, p, r, \gamma, \eta)$  where  $\mathcal{X}$  is the set of states of the MRP,  $X_t$  is the state at time  $t$ ,  $p : \mathcal{X} \times \mathcal{X} \rightarrow [0, 1]$  is the transition function where  $p(x, x') := \Pr(X_{t+1}=x' | X_t=x)$ ,  $r : \mathcal{X} \rightarrow \mathbb{R}$  is the reward function where  $r(x) := \mathbb{E}[R_t | X_t=x]$ ,  $\gamma \in [0, 1]$  is the discount factor, and  $\eta : \mathcal{X} \rightarrow [0, 1]$  is the starting state distribution. We consider finite horizon MRPs where episodes terminate by some (unspecified) timestep and set  $\gamma = 1$ .

A specific MRP is induced by the use of a fixed **policy**  $\pi$  on an **MDP**  $M$ , where  $\mathcal{X} = \mathcal{S}$ . The resulting transition and reward functions, denoted by  $p^\pi$  and  $r^\pi$  respectively, are:

$$p^\pi(x, x') = \frac{\sum_t \Pr(S_{t+1} = x', S_t = x; \pi)}{\sum_t \Pr(S_t = x; \pi)}, \quad r^\pi(x) = \frac{\sum_t \mathbb{E}[R_t | S_t = x; \pi] \Pr(S_t = x; \pi)}{\sum_t \Pr(S_t = x; \pi)}. \quad (2)$$

The Markov property [37] allows for further simplification of the above expressions (detailed in Appendix A.1), but this form is most conducive to our subsequent discussion. In this work, we focus on a specific instantiation of an MRP, described in the next section, where the set of states  $\mathcal{X}$  of the MRP are outputs of a state abstraction function  $\phi \in \Phi$ .

## 3 Abstract Reward Processes

An abstract reward process is a Markov reward process—derived from **MDP**  $M$  and **policy**  $\pi$  and defined over *abstract* states—that we use to evaluate  $\pi$ . The ARP provides two primary benefits for policy evaluation: (1) it preserves sufficient information to exactly evaluate the policy  $\pi$ , and (2) the ARP can be *consistently* estimated from data. In this section, we formalize the concept of an ARP, and highlight the theoretical and practical benefits of using ARPs for policy evaluation.

Given a state abstraction function  $\phi : \mathcal{S} \rightarrow \mathcal{Z}$ , the ARP  $\mathfrak{R}_\phi^\pi$  is defined such that  $\mathcal{X} = \mathcal{Z}$ . Formally,  $\mathfrak{R}_\phi^\pi$  is an MRP  $(\mathcal{Z}, P_\phi^\pi, R_\phi^\pi, \eta_\phi)$ , with  $\gamma = 1$  (see Appendix A.1 for a discussion on termination in ARPs and MRPs). The components of the ARP are defined over *abstract* states as:

$$P_\phi^\pi(z, z') := \frac{\sum_t \Pr(\phi(S_{t+1})=z', \phi(S_t)=z; \pi)}{\sum_t \Pr(\phi(S_t)=z; \pi)}, \quad R_\phi^\pi(z) := \frac{\sum_t \mathbb{E}[R_t | \phi(S_t)=z; \pi] \Pr(\phi(S_t)=z; \pi)}{\sum_t \Pr(\phi(S_t)=z; \pi)}, \quad (3)$$

and  $\eta_\phi(z) := \Pr(\phi(S_0) = z)$ . These expressions *cannot* be simplified further, unlike the case of an MRP [2]. Since  $\mathcal{Z}$  is a finite set, i.e., the abstract states are discrete, the components of the ARP can be represented by matrices (we use uppercase letters to emphasize this). The expected return of  $\mathfrak{R}_\phi^\pi$  can be computed efficiently using a linear solver to evaluate the expression  $J(\pi; \mathfrak{R}_\phi^\pi) := (\mathbf{I} - P_\phi^\pi)^{-1} R_\phi^\pi \eta_\phi$ , or via Monte Carlo rollouts of the reward process.

**ARPs are Performance Preserving:** The expected return of an ARP has a surprising property: even though some state information is abstracted away to create simple discrete abstract states, the finite

ARP, derived from a possibly continuous and complex MDP, preserves sufficient information about the performance of the policy that defines the ARP for all  $\phi \in \Phi$ .

**Theorem 3.1.**  $\forall \phi \in \Phi$ , the performance of a policy  $\pi$  is equal to the expected return of the abstract reward process  $\mathfrak{R}_\phi^\pi$  defined from MDP  $M$ , i.e.,  $J(\pi; \mathfrak{R}_\phi^\pi) = J(\pi; M)$ .

*Proof.* See Appendix B.1. □

The result holds for the ground-truth ARP  $\mathfrak{R}_\phi^\pi$ . In practice, a model of the ARP must be estimated from data. Next, we describe how the choice of defining an ARP over discrete abstract states eliminates model class mismatch, enabling asymptotically correct estimation of the ARP from data.

**Eliminating Model Class Mismatch:** Methods that learn models from data make an assumption about the class of models used to represent the data. A significant challenge is that of *model class mismatch*, where this assumed model class is often unable to represent the true data distribution. As an example, a neural network parameterizing a univariate Gaussian distribution cannot accurately represent data generated from a bimodal distribution. In the context of this work, the transition function of an ARP may specify arbitrary probability distributions over discrete abstract states, necessitating a careful selection of the model class. *Tabular models* are capable of representing any distribution over discrete variables. Therefore, using tabular models when estimating an ARP from data ensures that there is no model class mismatch. This is why we employ state abstraction functions  $\phi \in \Phi$  that partition the state space into a finite number of disjoint sets, or discrete abstract states. The abstraction functions may be viewed as: (a) a discrete clustering of the state space, or (b) a discretization of continuous states.

While this addresses model class mismatch, the use of a discrete state abstraction may itself be a source of modeling error. Mapping groups of (possibly continuous) states to discrete abstract states loses information about the state of the MDP. A process defined over the abstract states cannot in general capture the full complexity of the underlying MDP and policy. Nonetheless, Theorem 3.1 guarantees that the ARP is performance-preserving, ensuring that the use of discrete state abstractions is not a source of error for policy evaluation. Additionally, since we have eliminated model class mismatch, a perfect model of the ARP can be asymptotically estimated.

To estimate the ARP from  $\mathcal{D}_n^{(\pi_b)}$ , apply the state abstraction function to states in  $\mathcal{D}_n^{(\pi)}$  to map them to the abstract state space. Denote the *maximum likelihood estimate* of the model of the ARP obtained from the dataset (with abstract states) by  $\hat{\mathfrak{R}}_{n,\phi}^\pi := (\mathcal{Z}, \hat{\mathbf{P}}_{n,\phi}^\pi, \hat{\mathbf{R}}_{n,\phi}^\pi, \hat{\eta}_{n,\phi}^\pi)$ . The components take the form:

$$\hat{\mathbf{P}}_{n,\phi}^\pi(z, z') = \frac{\sum_{i,t} \mathbf{1}_t^i\{z, z'\}}{\sum_{i,t} \mathbf{1}_t^i\{z\}}; \quad \hat{\mathbf{R}}_{n,\phi}^\pi(z) = \frac{\sum_{i,t} \mathbf{1}_t^i\{z\} R_t^i}{\sum_{i,t} \mathbf{1}_t^i\{z\}}; \quad \hat{\eta}_{n,\phi}^\pi(z) = \frac{\sum_{i=1}^n \mathbf{1}^i\{z_0 = z\}}{n} \quad (4)$$

**Asymptotic Correctness:** With access to *on-policy data*  $\mathcal{D}_n^{(\pi)}$ , the following result states that ARPs enable consistent model-based estimation of the policy's performance.

**Lemma 3.2.**  $\forall \phi \in \Phi$ , the expected return of the maximum likelihood estimate  $\hat{\mathfrak{R}}_{n,\phi}^\pi$  converges almost surely to the expected return of the policy  $\pi$ , i.e.,  $J(\pi; \hat{\mathfrak{R}}_{n,\phi}^\pi) \xrightarrow{a.s.} J(\pi; M)$ .

*Proof.* See Appendix B.2. □

As the amount of data ( $n$ ) increases, the estimate of  $J(\pi; M)$  becomes increasingly accurate, i.e., the return estimate is consistent. This result holds for all  $\phi \in \Phi$ . It implies that *even an arbitrarily small model derived from a large, complex sequential decision-making problem will not introduce asymptotic bias*. However, this theoretical guarantee requires on-policy data and so does not directly assist us in off-policy evaluation. To that end, we introduce a procedure for estimation of the ARP from *off-policy data* that merges the machinery of IS-based and model-based methods.

### 3.1 Estimation from Off-Policy Data: Weighted Maximum Likelihood Estimation

We present a method for consistent estimation of the ARP corresponding to the evaluation policy  $\pi_e$  from off-policy data  $\mathcal{D}_n^{(\pi_b)}$ . It relies on the following intuition:

The expected value of the indicator function of an event represents the probability of that event. Use importance sampling to approximate the probability of that event under a different distribution.

To estimate an ARP from off-policy data, assign importance weights  $\rho_{0:t}$  to the abstract states ( $Z_t := \phi(S_t)$ ) in the dataset  $\mathcal{D}_n^{(\pi_b)}$ . Let  $H_t := (S_0, A_0, R_0, \dots, S_{t-1}, A_{t-1}, R_{t-1}, S_t, A_t)$  denote a sub-trajectory up to time  $t$ . The importance weight  $\rho_{0:t}$  is then the ratio of the probability of  $H_t$  under  $\pi_e$  and  $\pi_b$ , i.e.,  $\rho_{0:t} := \frac{\Pr(H_t; \pi_e)}{\Pr(H_t; \pi_b)} = \prod_{j=0}^t \frac{\pi_e(S_j, A_j)}{\pi_b(S_j, A_j)}$ .<sup>3</sup> The maximum likelihood estimate (MLE) of  $\mathfrak{R}_\phi^{\pi_e}$  obtained from the weighted off-policy data is denoted by  $\widehat{\mathfrak{R}}_{n,\phi}^{\pi_b \rightarrow \pi_e} := \left( \widehat{P}_{n,\phi}^{\pi_b \rightarrow \pi_e}, \widehat{R}_{n,\phi}^{\pi_b \rightarrow \pi_e}, \widehat{\eta}_{n,\phi}^{\pi_b \rightarrow \pi_e} \right)$ , where  $\widehat{\eta}_{n,\phi}^{\pi_b \rightarrow \pi_e}(z) = \frac{\sum_{i=1}^n \mathbf{1}^i\{z_0=z\}}{n}$  remains unchanged, and

$$\widehat{P}_{n,\phi}^{\pi_b \rightarrow \pi_e}(z, z') = \frac{\sum_{i,t} \mathbf{1}_t^i\{z, z'\} \rho_{0:t}}{\sum_{i,t} \mathbf{1}_t^i\{z\} \rho_{0:t}}, \quad \widehat{R}_{n,\phi}^{\pi_b \rightarrow \pi_e}(z) = \frac{\sum_{i,t} \mathbf{1}_t^i\{z\} \rho_{0:t} R_t^i}{\sum_{i,t} \mathbf{1}_t^i\{z\} \rho_{0:t}}. \quad (5)$$

This estimation is a form of weighted maximum likelihood estimation [13]. Including the importance ratios in the numerator and denominator of the estimated transition and reward functions of the ARP enables estimation from off-policy data generated by  $\pi_b$ . The estimated model of the ARP is consistent and, as shown next, allows for consistent off-policy evaluation.

**Lemma 3.3.** *Under Assumption 2.1, the weighted maximum likelihood estimate  $\widehat{\mathfrak{R}}_{n,\phi}^{\pi_b \rightarrow \pi_e}$  converges almost surely to the ground-truth ARP  $\mathfrak{R}_\phi^{\pi_e}$ , i.e.,  $\widehat{\mathfrak{R}}_{n,\phi}^{\pi_b \rightarrow \pi_e} \xrightarrow{a.s.} \mathfrak{R}_\phi^{\pi_e}$ .*

*Proof.* See Appendix B.3. □

## 4 Off-Policy Evaluation with ARPs

The expected return of  $\widehat{\mathfrak{R}}_{n,\phi}^{\pi_b \rightarrow \pi_e}$  is a consistent estimate of the performance of policy  $\pi_e$  since  $\widehat{\mathfrak{R}}_{n,\phi}^{\pi_b \rightarrow \pi_e}$  is an asymptotically correct estimate of  $\mathfrak{R}_\phi^{\pi_e}$ , as per Lemma 3.3.

**Theorem 4.1.** *The expected return of the ARP  $\widehat{\mathfrak{R}}_{n,\phi}^{\pi_b \rightarrow \pi_e}$  (built from off-policy data) converges almost surely to the expected return of  $\pi_e$ , i.e.,  $J(\pi_e; \widehat{\mathfrak{R}}_{n,\phi}^{\pi_b \rightarrow \pi_e}) \xrightarrow{a.s.} J(\pi_e; M)$ .*

*Proof.* See Appendix B.4. □

To our knowledge, this is the *first* instance of a model-based OPE method that comes with the theoretical guarantee of consistent performance estimates, even for continuous problems and without model class assumptions. So far, we have achieved one of the starting goals—that of *consistency*. However, the use of importance weights  $\rho_{0:t}$  for weighted MLE is expected to introduce high variance. Next, we discuss methods to mitigate the variance of IS.

### 4.1 Variance Reduction: Leveraging Markovness of the State Abstraction

A common technique for mitigating the variance of IS-based methods for OPE is clipping the importance weights to the  $c$  most recent ratios [18, 3, 20], i.e.,  $\rho_{(t-c+1)^+:t} := \prod_{i=(t-c+1)^+}^t \frac{\pi_e(S_i, A_i)}{\pi_b(S_i, A_i)}$ , where  $(t-c+1)^+ := \max(t-c+1, 0)$ . This is often a bad approximation for classical IS-based methods, as it implies that only the  $c$  most recent actions affect the reward distribution at any timestep, which rarely holds true in practice. In STAR, importance weights are incorporated into model estimation, resulting in a more reasonable implication of weight clipping.

By importance weighting the abstract-state occurrences as described in Equation (5), clipping importance weights, in this case, implies *that the  $c$  most recent abstract states are sufficient to determine the current abstract-state transition and reward distributions*. This allows actions from the distant past to influence the current reward, as *the effects of actions propagate through the abstract state transitions*, unlike in IS-based methods. This condition, that a recent history of abstract states is sufficient to predict the current abstract state transition distribution, often approximately holds

<sup>3</sup>The expression for the importance weight can be extended to continuous and hybrid probability measures using Radon-Nikodym derivatives. As with other terms, we hereafter denote the importance weight for the  $i^{\text{th}}$  episode as  $\rho_{0:t}^i$ .

in practice as discussed in POMDP literature [30, 39]. While the approximation may introduce asymptotic bias in exchange for reduced variance, certain abstraction functions that satisfy specific conditions can incur no asymptotic bias.

**Weight Clipping without Asymptotic Bias:** Intuitively, the use of  $c$ -clipped importance weights,  $\rho_{(t-c+1)^+:t}$ , updates the estimated distribution of the previous  $c$  abstract states—as if under the evaluation policy—while leaving the ones before unchanged.  $c$ -clipping does not introduce asymptotic bias when the previous  $c$  abstract states form a sufficient statistic for predicting the current abstract state transition distribution. This notion of conditional independence from history given the recent past is referred to as the *Markovness* of the abstraction function  $\phi$ . We posit that there exist abstraction functions that are  $c$ -th order Markov [51, 9].

**Definition 4.2** ( $c$ -th order Markov). The abstraction function  $\phi$  is  $c$ -th order Markov if  $\Pr(\phi(S_{t+1})|\phi(S_t), \dots, \phi(S_{(t-c+1)^+}); \pi) = \Pr(\phi(S_{t+1})|\phi(S_t), \dots, \phi(S_0); \pi)$  for  $\pi \in \{\pi_b, \pi_e\}$ .

Let  $\widehat{\mathfrak{R}}_{\phi,c}^{\pi_b \rightarrow \pi_e}$  denote the ARP estimated using  $c$ -clipped importance weights,  $\rho_{(t-c+1)^+:t}$ , in place of  $\rho_{0:t}$  in Equation (5).

**Theorem 4.3.** Given a  $c$ -th order Markov  $\phi$ , the expected return of the abstract reward process  $\widehat{\mathfrak{R}}_{\phi,c}^{\pi_b \rightarrow \pi_e}$  converges almost surely to the expected return of  $\pi_e$ , i.e.,  $J(\pi_e; \widehat{\mathfrak{R}}_{\phi,c}^{\pi_b \rightarrow \pi_e}) \xrightarrow{a.s.} J(\pi_e; M)$ .

*Proof.* See Appendix B.5.  $\square$

Even when  $\phi$  does not satisfy the above condition, weight clipping proves to be a practical approximation and results in low mean squared prediction error, as demonstrated empirically in Section 5. The steps for performing off-policy evaluation by estimating  $\widehat{\mathfrak{R}}_{\phi,c}^{\pi_b \rightarrow \pi_e}$  are highlighted in Algorithm 1.

## 4.2 Fantastic $\phi$ 's and Where to Find Them

Discrete abstraction functions that are  $c$ -th order Markov with small values of  $c$  represent the most suitable abstractions for enabling asymptotically correct, low-variance off-policy evaluation using STAR. An automated approach to discovering such abstraction functions, however, remains elusive. In a manner reminiscent of the options framework [50], wherein one might consider the usefulness of options before having methods for constructing options automatically, this work emphasizes the remarkable effectiveness of state abstractions used in abstract reward processes for OPE. It motivates a research area akin to option discovery: *abstraction discovery for OPE*.

We expect the following factors to play an important role in the search for good abstraction functions: (a) state-visitation distributions of  $\pi_b$  and  $\pi_e$ , determining the granularity of abstraction in different parts of the state set, and (b) the distribution shift in abstract state visitation induced by the two policies, determining the extent of weight clipping that can be applied. Both of these are affected by properties of the underlying MDP, in particular the transition function, and in our initial analyses, we observe varying effects of similar abstractions across different MDPs (Appendix C.2).

We observe that a simple approach of randomly initializing centroids and applying  $k$ -means clustering [34, 33], where each cluster denotes a discrete abstract state, results in abstractions that provide competitive OPE performance, often significantly outperforming existing methods. We call this naive clustering-based abstraction method CluSTAR, and use it for our experiments. In some cases, abstraction by aggregation of states can increase the difficulty of estimation of the transition function. For example, aggregation of two states with deterministic transitions—which can be estimated perfectly from a single observation of those transitions—creates stochastic transitions between abstract states. However, in general, state aggregation tends to simplify estimation by increasing the effective sample size [21].

**Recovering Existing OPE Methods from STAR:** Different configurations of  $(\phi, c)$  induce different ARPs,  $\widehat{\mathfrak{R}}_{\phi,c}^{\pi_b \rightarrow \pi_e}$ . For certain configurations of  $(\phi, c)$ :

- $|\mathcal{Z}| = 1$  and no weight clipping: Mapping all states to a single abstract state yields the *weighted per-decision importance sampling* (WPDIS) estimator [42].
- $\mathcal{Z} = \mathcal{S}$  and  $c = 1$ : Amounts to no state abstraction, and yields the maximum likelihood estimate of the MRP over states. The MRP is a combination of the approximate-model estimator [40] that directly estimates the model dynamics with the evaluation policy.



The recovery of these familiar estimators at the endpoints of STAR highlights the unifying nature of the framework. More importantly, the intermediate configurations of  $(\phi, c)$  uncover a whole new set of OPE estimators. The ARPs in this space often inherit a mixture of the favorable properties of both of the endpoints. Consequently, the framework yields estimators that can significantly outperform existing methods, as shown empirically in the next section.

## 5 Empirical Analysis

In this section, we **(A)** analyze the performance of the set of estimators (ARPs) induced by STAR across different configurations of  $(\phi, c)$ , and **(B)** compare the performance of the best and median ARPs from this set against existing OPE methods to demonstrate that estimators encompassed by STAR often outperform prior OPE methods. We use the following domains for OPE: **(1)** CartPole [49]: A classic control domain in OpenAI Gym [6]. **(2)** ICU-Sepsis [8]: An MDP that simulates treatment of sepsis in the ICU. ICU-Sepsis is built from real-world medical records obtained from the MIMIC-III dataset [24], using a modified version of the process described by Komorowski et al. [26]. **(3)** Asterix from the MinAtar testbed [61]: A miniaturized version of the Atari game Asterix. Details about each domain, and about the behavior and evaluation policies can be found in the Appendix C. The code is available at: <https://github.com/shreyasc-13/STAR>.

**(A) Estimator Selection:** Estimator selection presents a significant challenge for OPE [48] due to the unavailability of a *validation set*. To select STAR estimators to compare against existing methods, we first report the performance of the set of the ARPs induced by STAR, across a range of configurations of  $(\phi, c)$ . We highlight the performance of the **best** and **median** estimators from this set. For reference, we compare the mean squared prediction errors from the estimated ARPs against WPDIS and approximate-model estimator (MBased), the two endpoints of STAR, as shown in Figure 2 on the CartPole domain. The state abstraction is performed with CluSTAR with the number of centroids  $|\mathcal{Z}| \in \{2, 4, 8, 16, 32, 64, 128\}$  and the weight clipping factor  $c \in \{1, 2, 3, 4, 5\}$  defining 35 ARPs, where the performance of each is indicated by  $\bullet$ . This range of  $(\phi, c)$  is picked based on the intuition that (a) larger values of  $c$  are expected to introduce high variance, and (b) this range of  $|\mathcal{Z}|$  covers a variety of granularities of state abstraction (for a continuous problem). The results are averages across 200 trials. The prediction errors for the estimated ARPs are competitive with baselines. This suggests that *an average estimator in STAR is competitive with existing OPE methods*, and even better performance may be attained by using specialized methods for abstraction discovery and estimator selection [55]. Figure 4, in the Appendix, provides a detailed breakdown of the performance of each ARP in the set considered, presented as a heatmap. This highlights patterns observed for varying values of  $|\mathcal{Z}|$  and  $c$  across different domains. Further discussion and details about the best-performing configurations of  $(\phi, c)$  are deferred to Appendix C.2.

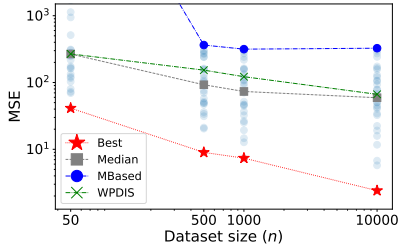


Figure 2: Mean squared prediction errors of the estimated ARPs for the set of hyperparameters swept over for CartPole.

**(B) ARPs Can Outperform Existing Methods for OPE:** We compare against representative methods from the main categories of OPE methods: (a) **IS-based methods:** Vanilla IS, Per-Decision IS, Weighted IS, Weighted Per-Decision IS [42]; (b) **Model-based methods:** that approximate a model of the MDP—MBased [49]—and directly estimate off-policy Q-values—FQE and MRDR [27, 11]; and (c) **Mixture methods:** DR [22] and MAGIC [53], which blend estimates from the methods in the aforementioned categories. We use implementations from the Caltech OPE Benchmarking Suite (COBS) [58] for all methods. The representative method for minmax style estimators, IH [31], is designed for the infinite horizon setting and performs poorly with  $\gamma = 1$  [63]. Due to the instability of IH estimates in our experiments, we excluded it from comparison.

**Results [Figure 3 and 4]:** The ranges of  $|\mathcal{Z}|$  and  $c$  for each domain, which induces a set of ARPs that we consider, are detailed in Appendix C.2. Figure 4 in Appendix C.2 demonstrates the performance of each ARP from that set as a heatmap, showing that the best configuration of  $(\phi, c)$  varies with the dataset size for each domain. In critical applications like sepsis treatment, as simulated in ICU-Sepsis, where incorrect policy performance estimates can be catastrophic, the best estimators in






Figure 3: Mean squared prediction errors of **best** and **median** ARPs from STAR compared against existing OPE methods. The empirically estimated bias-variance decomposition of the error is shown. The results are averaged over 200 trials, with error bars indicating standard error. Note: For ICU-Sepsis, regression-based methods (MRDR and Q-Reg) were computationally intractable due to the large state set, as the corresponding Weighted Least Squares methods for regression were too slow. In all domains and across all datasizes, the best ARP in STAR outperforms baselines in all cases, and the even the median estimator does so in 7 out of 12 cases.

STAR achieve prediction errors that are an order of magnitude lower than the best baseline. These results in Figure 3 emphasize that *highly compact ARPs that distill complex sequential processes are particularly effective for off-policy evaluation*. The **best** ARP in STAR for each domain abstracts: (a) the continuous state space of CartPole to  $|\mathcal{Z}| = 32$  abstract states, (b) the 747 states of ICU-Sepsis to  $|\mathcal{Z}| = 16$  abstract states, and (c) the 400 states of Asterix to  $|\mathcal{Z}| = 8$  abstract states, to construct compact finite ARPs for OPE. Furthermore, the horizon length of CartPole is 50, and episodes in ICU-Sepsis and Asterix go up to 120 and 75 timesteps respectively. Such relatively long horizons have proven to be challenging for prior OPE methods. STAR leverages ARPs to enable OPE at scales commonly seen in practice.

## 6 Related Work

The problem of off-policy evaluation (OPE) has been extensively studied due to its relevance for practical applications of reinforcement learning [38, 35]. Extensive surveys on the topic, both theoretical [56] and empirical [58, 14] delineate the numerous approaches to the problem. Model-based approaches for OPE have proven effective [32, 62] but are restricted by the model class used. In this work, we use state abstraction to define compact models called abstract reward processes, and demonstrate their effectiveness as off-policy estimators. Historically, state abstraction research has focused on grouping similar states in a way that does not change the essence of the underlying problem [28, 43, 1], to reduce the complexity of the problem. However, the use of state abstractions for OPE remains under-explored. Pavse and Hanna [41] show that the use of state abstraction with marginalized importance sampling achieves variance reduction in high-dimensional state spaces, but their approach does not use abstraction to construct models. Jiang et al. [23] study abstraction selection for model-based RL, balancing model complexity and policy value suboptimality. Abstraction discovery or

learning has focussed on discretization of continuous state spaces to reduce problem complexity [47], and distilling the Markov features [2] or reward relevant features [12, 59].

## 7 Discussion and Conclusion

In this paper, we have introduced a new framework for consistent model-based off-policy evaluation. This framework leverages state abstraction to prevent model class mismatch along with importance sampling to consistently learn models from off-policy data. Unlike traditional model-based methods, our approach eliminates the need for model class assumptions and provides theoretical guarantees for the obtained performance estimates. Moreover, using state abstraction increases the effective sample size [21], which is particularly beneficial in limited data regimes. Importantly, this work presents a framework with a new approach to OPE, rather than a specific new method. Estimators that lie within this framework significantly outperform existing OPE methods, with the best estimator consistently outperforming all baselines, as demonstrated in our empirical evaluation.

The framework has two main limitations: it requires knowledge of the probabilities of observed actions under the behavior policy, which may not always be available, and a principled method for selecting well-performing configurations of the abstraction function and the weight clipping factor remain elusive. Combining this work with regression IS [19] would be a practical extension that addresses the first limitation. Additionally, a data-driven approach to automated estimator selection based on characteristics of the domain, dataset sizes, and other factors, as suggested by Su et al. [48], would enhance its practical application.

Our findings indicate that even a simple class of abstraction functions can provide competitive OPE performance. We theoretically demonstrate the existence of certain abstraction functions that may offer better performance. Investigating the properties of abstraction functions and developing automated approaches to *abstraction discovery* for ARPs are promising directions for future work on creating high-performing OPE methods.

**Acknowledgements** We thank Yash Chandak, Mohammad Ghavamzadeh and Dhawal Gupta for their helpful discussions and feedback on this work. We would also like to thank Josiah Hanna, Bo Liu, Cameron Allen, and the anonymous reviewers for their feedback.

## References

- [1] David Abel. A theory of state abstraction for reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 9876–9877, 2019.
- [2] Cameron Allen, Neev Parikh, Omer Gottesman, and George Konidaris. Learning Markov state abstractions for deep reinforcement learning. *Advances in Neural Information Processing Systems*, 34:8229–8241, 2021.
- [3] Oliver Bembom and Mark J van der Laan. Data-adaptive selection of the truncation level for inverse-probability-of-treatment-weighted estimators. 2008.
- [4] Patrick Billingsley. *Convergence of Probability Measures*. John Wiley & Sons, 2013.
- [5] Léon Bottou, Jonas Peters, Joaquin Quiñero-Candela, Denis X Charles, D Max Chickering, Elon Portugaly, Dipankar Ray, Patrice Simard, and Ed Snelson. Counterfactual reasoning and learning systems: The example of computational advertising. *Journal of Machine Learning Research*, 14(11), 2013.
- [6] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. OpenAI gym. *arXiv Preprint arXiv:1606.01540*, 2016.
- [7] Olivier Chapelle and Lihong Li. An empirical evaluation of Thompson sampling. *Advances in Neural Information Processing Systems*, 24, 2011.
- [8] Kartik Choudhary, Dhawal Gupta, and Philip S. Thomas. ICU-Sepsis: A benchmark MDP built from real medical data, 2024.

- [9] Yonathan Efroni, Chi Jin, Akshay Krishnamurthy, and Sobhan Miryoosefi. Provable reinforcement learning with a short-term memory. In *Proceedings of the International Conference on Machine Learning*, pages 5832–5850. PMLR, 2022.
- [10] Amir-Massoud Farahmand and Csaba Szepesvári. Model selection in reinforcement learning. *Machine Learning*, 85(3):299–332, 2011.
- [11] Mehrdad Farajtabar, Yinlam Chow, and Mohammad Ghavamzadeh. More robust doubly robust off-policy evaluation. In *Proceedings of the International Conference on Machine Learning*, pages 1447–1456. PMLR, 2018.
- [12] Norman Ferns, Prakash Panangaden, and Doina Precup. Metrics for finite Markov decision processes. *arXiv Preprint arXiv:1207.4114*, 2012.
- [13] Chris Field and B Smith. Robust estimation: A weighted maximum likelihood approach. *International Statistical Review/Revue Internationale de Statistique*, pages 405–424, 1994.
- [14] Justin Fu, Mohammad Norouzi, Ofir Nachum, George Tucker, Ziyu Wang, Alexander Novikov, Mengjiao Yang, Michael R Zhang, Yutian Chen, Aviral Kumar, et al. Benchmarks for deep off-policy evaluation. *arXiv Preprint arXiv:2103.16596*, 2021.
- [15] Ge Gao, Song Ju, Markel Sanz Ausin, and Min Chi. Hope: Human-centric off-policy evaluation for e-learning and healthcare. *arXiv Preprint arXiv:2302.09212*, 2023.
- [16] LC Evans-RF Garipey and F Ronald. *Measure theory and fine properties of functions, Revised edition*. CRC Press, 2015.
- [17] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *International conference on machine learning*, pages 1321–1330. PMLR, 2017.
- [18] Zhaohan Guo, Philip S Thomas, and Emma Brunskill. Using options and covariance testing for long horizon off-policy policy evaluation. *Advances in Neural Information Processing Systems*, 30, 2017.
- [19] Josiah Hanna, Scott Niekum, and Peter Stone. Importance sampling policy evaluation with an estimated behavior policy. In *Proceedings of the International Conference on Machine Learning*, pages 2605–2613. PMLR, 2019.
- [20] Edward L Ionides. Truncated importance sampling. *Journal of Computational and Graphical Statistics*, 17(2):295–311, 2008.
- [21] Nan Jiang. Notes on state abstractions, 2018. URL <https://nanjiang.cs.illinois.edu/files/cs598/note4.pdf>.
- [22] Nan Jiang and Lihong Li. Doubly robust off-policy value evaluation for reinforcement learning. In *Proceedings of the International Conference on Machine Learning*, pages 652–661. PMLR, 2016.
- [23] Nan Jiang, Alex Kulesza, and Satinder Singh. Abstraction selection in model-based reinforcement learning. In *Proceedings of the International Conference on Machine Learning*, pages 179–188. PMLR, 2015.
- [24] Alistair EW Johnson, Tom J Pollard, Lu Shen, Li-wei H Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. MIMIC-III, a freely accessible critical care database. *Scientific Data*, 3(1):1–9, 2016.
- [25] Herman Kahn and Theodore E Harris. Estimation of particle transmission by random sampling. *National Bureau of Standards Applied Mathematics Series*, 12:27–30, 1951.
- [26] Matthieu Komorowski, Leo A Celi, Omar Badawi, Anthony C Gordon, and A Aldo Faisal. The artificial intelligence clinician learns optimal treatment strategies for sepsis in intensive care. *Nature Medicine*, 24(11):1716–1720, 2018.

- [27] Hoang Le, Cameron Voloshin, and Yisong Yue. Batch policy learning under constraints. In *Proceedings of the International Conference on Machine Learning*, pages 3703–3712. PMLR, 2019.
- [28] Lihong Li, Thomas J Walsh, and Michael L Littman. Towards a unified theory of state abstraction for MDPs. *Artificial Intelligence and Machine Learning (AI&M)*, 1(2):3, 2006.
- [29] Lihong Li, Rémi Munos, and Csaba Szepesvári. Toward minimax off-policy value estimation. In *Proceedings of the Artificial Intelligence and Statistics Conference*, pages 608–616. PMLR, 2015.
- [30] Michael Littman and Richard S. Sutton. Predictive representations of state. *Advances in Neural Information Processing Systems*, 14, 2001.
- [31] Qiang Liu, Lihong Li, Ziyang Tang, and Dengyong Zhou. Breaking the curse of horizon: Infinite-horizon off-policy estimation. *Advances in Neural Information Processing Systems*, 31, 2018.
- [32] Yao Liu, Omer Gottesman, Aniruddh Raghu, Matthieu Komorowski, Aldo A Faisal, Finale Doshi-Velez, and Emma Brunskill. Representation balancing MDPs for off-policy policy evaluation. *Advances in Neural Information Processing Systems*, 31, 2018.
- [33] Stuart Lloyd. Least Squares Quantization in PCM. *IEEE Transactions on Information Theory*, 28(2):129–137, 1982.
- [34] James MacQueen. Some methods for classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, volume 1, pages 281–297, Oakland, CA, USA, 1967. University of California Press.
- [35] Travis Mandel, Yun-En Liu, Sergey Levine, Emma Brunskill, and Zoran Popovic. Offline policy evaluation across representations with applications to educational games. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*, volume 1077, 2014.
- [36] Vukosi N Marivate. *Improved Empirical Methods in Reinforcement-Learning Evaluation*. Rutgers The State University of New Jersey, School of Graduate Studies, 2015.
- [37] Andrei Andreevich Markov. The theory of algorithms. *Trudy Matematicheskogo Instituta Imeni VA Steklova*, 42:3–375, 1954.
- [38] Susan A Murphy, Mark J van der Laan, James M Robins, and Conduct Problems Prevention Research Group. Marginal mean models for dynamic regimes. *Journal of the American Statistical Association*, 96(456):1410–1423, 2001.
- [39] Khanh Xuan Nguyen. Converting POMDPs into MDPs using history representation. *Engineering Archive*, 2021.
- [40] Cosmin Paduraru. *Off-Policy Evaluation in Markov Decision Processes*. PhD thesis, 2013.
- [41] Brahma S Pavse and Josiah P Hanna. Scaling marginalized importance sampling to high-dimensional state-spaces via state abstraction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 9417–9425, 2023.
- [42] Doina Precup. Eligibility traces for off-policy policy evaluation. *Computer Science Department Faculty Publication Series*, page 80, 2000.
- [43] Balaraman Ravindran. *An Algebraic Approach to Abstraction in Reinforcement Learning*. University of Massachusetts Amherst, 2004.
- [44] Reuven Y Rubinstein and Dirk P Kroese. *Simulation and the Monte Carlo Method*. John Wiley & Sons, 2016.
- [45] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *Proceedings of the International Conference on Machine Learning*, pages 1889–1897. PMLR, 2015.

- [46] Pranab K Sen and Julio M Singer. *Large Sample Methods in Statistics: An Introduction with Applications*, volume 25. CRC Press, 1994.
- [47] Sean R Sinclair, Siddhartha Banerjee, and Christina Lee Yu. Adaptive discretization for episodic reinforcement learning in metric spaces. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 3(3):1–44, 2019.
- [48] Yi Su, Pavithra Srinath, and Akshay Krishnamurthy. Adaptive estimator selection for off-policy evaluation. In *Proceedings of the International Conference on Machine Learning*, pages 9196–9205. PMLR, 2020.
- [49] Richard S Sutton and Andrew G Barto. *Reinforcement Learning: An Introduction*. MIT Press, 2018.
- [50] Richard S Sutton, Doina Precup, and Satinder Singh. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence*, 112(1-2): 181–211, 1999.
- [51] Erik N Talvitie. *Simple Partial Models for Complex Dynamical Systems*. PhD thesis, University of Michigan, 2010.
- [52] Shengpu Tang and Jenna Wiens. Model selection for offline reinforcement learning: Practical considerations for healthcare settings. In *Proceedings of the Machine Learning for Healthcare Conference*, pages 2–35. PMLR, 2021.
- [53] Philip Thomas and Emma Brunskill. Data-efficient off-policy policy evaluation for reinforcement learning. In *Proceedings of the International Conference on Machine Learning*, pages 2139–2148. PMLR, 2016.
- [54] Edward C Tolman. Cognitive maps in rats and men. *Psychological Review*, 55(4):189, 1948.
- [55] Takuma Udagawa, Haruka Kiyohara, Yusuke Narita, Yuta Saito, and Kei Tateno. Policy-adaptive estimator selection for off-policy evaluation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 10025–10033, 2023.
- [56] Masatoshi Uehara, Chengchun Shi, and Nathan Kallus. A review of off-policy evaluation in reinforcement learning. *arXiv Preprint arXiv:2212.06355*, 2022.
- [57] Aad W Van der Vaart. *Asymptotic Statistics*, volume 3. Cambridge University Press, 2000.
- [58] Cameron Voloshin, Hoang M Le, Nan Jiang, and Yisong Yue. Empirical study of off-policy policy evaluation for reinforcement learning. *arXiv Preprint arXiv:1911.06854*, 2019.
- [59] Tongzhou Wang, Simon S Du, Antonio Torralba, Phillip Isola, Amy Zhang, and Yuandong Tian. Denoised MDPs: Learning world models better than the world itself. *arXiv Preprint arXiv:2206.15477*, 2022.
- [60] Neil A. Weiss. *A Course in Probability*. Addison-Wesley, Boston, 2005. ISBN 0-321-18954-X.
- [61] Kenny Young and Tian Tian. MinAtar: An Atari-inspired testbed for thorough and reproducible reinforcement learning experiments. *arXiv Preprint arXiv:1903.03176*, 2019.
- [62] Michael R Zhang, Tom Le Paine, Ofir Nachum, Cosmin Paduraru, George Tucker, Ziyu Wang, and Mohammad Norouzi. Autoregressive dynamics models for offline policy evaluation and optimization. *arXiv Preprint arXiv:2104.13877*, 2021.
- [63] Ruiyi Zhang, Bo Dai, Lihong Li, and Dale Schuurmans. Gendice: Generalized offline estimation of stationary values. *arXiv Preprint arXiv:2002.09072*, 2020.
- [64] Daniel Zwillinger and Stephen Kokoska. *CRC Standard Probability and Statistics Tables and Formulae*. CRC Press, 1999.

## A Preliminaries

**Consistency and Almost Sure Convergence** Let  $\hat{\theta}_n$  denote the estimator of a statistic  $\theta$ , estimated from  $n$  data points. As the amount of data approaches infinity, i.e., as  $n \rightarrow \infty$ , the estimator is said to *converge almost surely* to  $\theta$  if and only if

$$\Pr\left(\lim_{n \rightarrow \infty} \hat{\theta}_n = \theta\right) = 1.$$

We write  $\theta_n \xrightarrow{\text{a.s.}} \theta$  to denote that the sequence  $\theta_n$  converges almost surely to  $\theta$ . An estimator that converges almost surely to the true value of the statistic is said to be (strongly) **consistent**.

**Continuous Mapping Theorem** The continuous mapping theorem [4, 57] states that if a sequence of random variables  $(X_i)_{i=1}^n$  converges almost surely to a random variable  $X$ , a function  $f$  has discontinuity points  $D_f$ , and  $\Pr(X \in D_f) = 0$ , then the sequence  $(f(X_i))_{i=1}^n$  converges almost surely to  $f(X)$ .

**Tower Rule** For random variables  $X$  and  $Y$  defined on the same probability space, the tower rule states that the expected value of the conditional expected value of  $X$  given  $Y$  is the same as the expected value of  $X$ , i.e.,

$$\mathbb{E}[X] = \mathbb{E}[\mathbb{E}[X | Y]].$$

### A.1 Additional Notes on MRPs and ARPs

Here we provide additional details regarding the definitions of MDPs and MRPs. Specifically, we show how the expression of the transition and reward functions of the MRP can be simplified when  $\mathcal{X} = \mathcal{S}$  and discuss how to handle termination in ARPs that are derived from finite-horizon MDPs.

#### A.1.1 Simplification of Equation 2

The expressions for the components of a Markov reward process (MRP) can be simplified by using the Markov property. The transition function simplifies as:

$$\begin{aligned} p^\pi(x, x') &= \frac{\sum_t \Pr(S_{t+1} = x', S_t = x; \pi)}{\sum_t \Pr(S_t = x; \pi)} \\ &= \frac{\sum_t \Pr(S_{t+1} = x' | S_t = x; \pi) \Pr(S_t = x; \pi)}{\sum_t \Pr(S_t = x; \pi)} \\ &= \frac{\sum_t \sum_{a \in \mathcal{A}} \Pr(S_{t+1} = x' | S_t = x, A_t = a; \pi) \Pr(A_t = a | S_t = x) \Pr(S_t = x; \pi)}{\sum_t \Pr(S_t = x; \pi)} \\ &= \frac{\sum_t \sum_{a \in \mathcal{A}} p(x, a, x') \pi(x, a) \Pr(S_t = x; \pi)}{\sum_t \Pr(S_t = x; \pi)} \\ &= \frac{\sum_{a \in \mathcal{A}} p(x, a, x') \pi(x, a) (\sum_t \Pr(S_t = x; \pi))}{\sum_t \Pr(S_t = x; \pi)} \\ &= \sum_{a \in \mathcal{A}} p(x, a, x') \pi(x, a). \end{aligned}$$

Similarly, the reward function simplifies as:

$$\begin{aligned} r^\pi(x) &= \frac{\sum_t \mathbb{E}[R_t | S_t = x; \pi] \Pr(S_t = x; \pi)}{\sum_t \Pr(S_t = x; \pi)} \\ &= \frac{\sum_t \sum_{a \in \mathcal{A}} r(x, a) \pi(x, a) \Pr(S_t = x; \pi)}{\sum_t \Pr(S_t = x; \pi)} \\ &= \frac{\sum_{a \in \mathcal{A}} r(x, a) \pi(x, a) (\sum_t \Pr(S_t = x; \pi))}{\sum_t \Pr(S_t = x; \pi)} \\ &= \sum_{a \in \mathcal{A}} r(x, a) \pi(x, a). \end{aligned}$$

### A.1.2 Handling termination in an MDPs, MRPs, and ARPs

Each episode in a finite-horizon MDP concludes by some timestep  $T \in \mathbb{N}$ , referred to as the *termination* of that episode. In practice, there are two common ways of modeling termination of episodes in finite-horizon MDPs: 1) including an absorbing state  $s_\infty$  in the state set, or 2) introducing a termination function  $\beta : \mathcal{S} \rightarrow [0, 1]$  representing the probability of termination of an episode from each state, i.e.,  $\beta(s) := \Pr(\text{terminate} | S_t = s)$ . In the first case, by timestep  $T$  the process transitions into the absorbing  $s_\infty$  after which it continually transitions back to  $s_\infty$  getting a reward of zero. In the second case, after timestep  $T$  the process stops and there are no subsequent samples. Expressions involving sums over samples in an episode are denoted correspondingly, for example, the expected return  $J(\pi) := \mathbb{E}[\sum_{t=1}^{\infty} R_t; \pi]$  in the first case, and  $J(\pi) := \mathbb{E}[\sum_{t=1}^T R_t; \pi]$  in the second. Both approaches are mathematically equivalent and correspondingly, MRPs induced by the combination of finite-horizon MDPs with a policy can model termination in either way.

However, termination in ARPs requires special attention. The introduction of an absorbing abstract state  $z_\infty$  places a condition on the abstraction function  $\phi$ —it requires the abstraction function to map *only*  $s_\infty$ , and no other state, to  $z_\infty$ . In this work, we *implement* termination in code by estimating abstract termination functions that represent the probability of termination from each abstract state. Denote the abstract termination function of an ARP by  $\beta_\phi^\pi : \mathcal{Z} \rightarrow [0, 1]$ . It must be noted that both approaches continue to be mathematically equivalent. The probability of transitioning from any abstract state  $z$  to  $z_\infty$  is the same as the the probability of termination from that state, i.e.,  $\beta_\phi^\pi(z) = P_\phi^\pi(z, z_\infty)$ . Note that

$$P_\phi^\pi(z_\infty, z) = \begin{cases} 0 & \text{if } z \neq z_\infty \\ 1 & \text{if } z = z_\infty. \end{cases}$$

Similarly, the reward function is zero for the absorbing abstract state:  $R_\phi^\pi(z_\infty) = 0$ . It is mathematically more succinct to model termination with  $z_\infty$ , and we use this (equivalent) approach in our theoretical analysis.

In our implementation,  $z_\infty$  is not defined and thus the components of  $\mathfrak{R}_\phi$  do not take  $z_\infty$  as input. The equivalent form of the expected return of an ARP, that models a termination function, is then given by:

$$\begin{aligned} J(\pi; \mathfrak{R}_\phi^\pi) &= R_\phi^\pi \eta_\phi + (\mathbf{I} - \text{diag}(\beta_\phi^\pi)) [P_\phi^\pi R_\phi^\pi \eta_\phi + (\mathbf{I} - \text{diag}(\beta_\phi^\pi)) [(P_\phi^\pi)^2 R_\phi^\pi \eta_\phi + \dots]] \\ &= R_\phi^\pi \eta_\phi + (\mathbf{I} - \text{diag}(\beta_\phi^\pi)) P_\phi^\pi R_\phi^\pi \eta_\phi + (\mathbf{I} - \text{diag}(\beta_\phi^\pi))^2 (P_\phi^\pi)^2 R_\phi^\pi \eta_\phi + \dots \\ &= \sum_{k=0}^{\infty} ((\mathbf{I} - \text{diag}(\beta_\phi^\pi)) P_\phi^\pi)^k R_\phi^\pi \eta_\phi. \end{aligned}$$

Correspondingly, the closed form expression for the expected return that accounts for the termination function is given by,

$$J(\pi; \mathfrak{R}_\phi^\pi) = (\mathbf{I} - (\mathbf{I} - \text{diag}(\beta_\phi^\pi)) P_\phi^\pi)^{-1} R_\phi^\pi(z) \eta_\phi. \quad (6)$$

Note that  $(\mathbf{I} - (\mathbf{I} - \text{diag}(\beta_\phi^\pi)) P_\phi^\pi)$  is an invertible matrix, as the left hand side of the expression,  $J(\pi; \mathfrak{R}_\phi^\pi)$ , is equal to the expected return  $J(\pi; M)$  of  $\pi$  by Theorem 3.1, which is bounded.



## B Proofs of Theoretical Results

In this section, we provide proofs of the theorems and lemmas in the main text. We start with Theorem 3.1 that states that ARPs are *performance-preserving*.

### B.1 ARPs are Performance Preserving

**Theorem 3.1.**  $\forall \phi \in \Phi$ , the performance of a policy  $\pi$  is equal to the expected return of the abstract reward process  $\mathfrak{R}_\phi^\pi$  defined from MDP  $M$ , i.e.,  $J(\pi; \mathfrak{R}_\phi^\pi) = J(\pi; M)$ .

*Proof.* This result states that despite the use of a discrete state abstraction, the expected return of the ARP  $\mathfrak{R}_\phi^\pi$  is equal to the performance of policy  $\pi$ . To prove this, we leverage two equivalent forms of defining the expected return of a policy. The first form:

$$J(\pi; \mathfrak{R}_\phi^\pi) = \sum_t \mathbb{E}[R_t; \pi], \quad (7)$$

defines the expected return as a sum of the expected value of the reward at each timestep, where the distribution of rewards at each timestep is governed by the components of the ARP, namely,  $\eta_\phi$ ,  $P_\phi^\pi$  and  $R_\phi^\pi$ . The second, equivalent, form of expressing the expected return is given by:

$$J(\pi; \mathfrak{R}_\phi^\pi) = \sum_t \sum_{z \in \mathcal{Z}} \Pr(Z_t = z; \pi) \mathbb{E}[R_t | Z_t = z; \pi] = \sum_t \sum_{z \in \mathcal{Z}} \Pr(Z_t = z; \pi) R_\phi^\pi(z), \quad (8)$$

where the expected return is denoted as a sum of the reward at each abstract state multiplied by the visitation frequency of that abstract state. Both Equations 7 and 8 are equivalent, i.e.,

$$J(\pi; \mathfrak{R}_\phi^\pi) = \sum_t \mathbb{E}[R_t; \pi] = \sum_t \sum_{z \in \mathcal{Z}} \Pr(Z_t = z; \pi) R_\phi^\pi(z). \quad (9)$$

Next, we denote the terms in Equation 8 in terms of the states of the underlying MDP to show the final result. The reward function  $R_\phi^\pi$  can be expressed as,

$$\begin{aligned} R_\phi^\pi(z) &:= \frac{\sum_t \mathbb{E}[R_t | \phi(S_t) = z; \pi] \Pr(\phi(S_t) = z; \pi)}{\sum_t \Pr(\phi(S_t) = z; \pi)} \\ &= \frac{\sum_t \sum_{s \in \mathcal{S}} \mathbf{1}\{\phi(s) = z\} \mathbb{E}[R_t | S_t = s; \pi] \Pr(S_t = s; \pi)}{\sum_t \sum_{s \in \mathcal{S}} \mathbf{1}\{\phi(s) = z\} \Pr(S_t = s; \pi)} \\ &= \frac{\sum_t \sum_{s \in \mathcal{S}, a \in \mathcal{A}} \mathbf{1}\{\phi(s) = z\} r(s, a) \pi(s, a) \Pr(S_t = s; \pi)}{\sum_t \sum_{s \in \mathcal{S}} \mathbf{1}\{\phi(s) = z\} \Pr(S_t = s; \pi)} \\ &= \frac{\sum_{s \in \mathcal{S}, a \in \mathcal{A}} \mathbf{1}\{\phi(s) = z\} r(s, a) \pi(s, a) (\sum_t \Pr(S_t = s; \pi))}{\sum_{s \in \mathcal{S}} \mathbf{1}\{\phi(s) = z\} (\sum_t \Pr(S_t = s; \pi))} \\ &= \frac{\sum_{s \in \mathcal{S}, a \in \mathcal{A}} \psi^\pi(s) \pi(s, a) r(s, a) \mathbf{1}\{\phi(s) = z\}}{\sum_{s \in \mathcal{S}} \psi^\pi(s) \mathbf{1}\{\phi(s) = z\}}, \end{aligned}$$

where the undiscounted state distribution [49, Section 9.2] under policy  $\pi$  is denoted by  $\psi^\pi(s) \propto \sum_t \Pr(S_t = s; \pi)$ , where  $\psi^\pi(s_\infty)$  is set to be proportional to any constant value. The normalization constant,  $\kappa$ , that makes  $\psi^\pi(s)$  a valid distribution cancels out from both the numerator and the denominator of the above expression. The term  $\Pr(Z_t = z; \pi)$  can be expressed in terms of states as,

$$\Pr(Z_t = z; \pi) = \sum_{s \in \mathcal{S}} \Pr(S_t = s; \pi) \mathbf{1}\{\phi(s) = z\}.$$

Substituting these expressions into Equation 8 gives,

$$\begin{aligned}
J(\pi; \mathfrak{R}_\phi^\pi) &= \sum_t \sum_{z \in \mathcal{Z}} \Pr(Z_t = z; \pi) R_\phi^\pi(z) \\
&= \sum_t \sum_{z \in \mathcal{Z}} \left( \sum_{s \in \mathcal{S}} \Pr(S_t = s; \pi) \mathbf{1}\{\phi(s) = z\} \right) \left( \frac{\sum_{s \in \mathcal{S}, a \in \mathcal{A}} \psi^\pi(s) \pi(s, a) r(s, a) \mathbf{1}\{\phi(s) = z\}}{\sum_{s \in \mathcal{S}} \psi^\pi(s) \mathbf{1}\{\phi(s) = z\}} \right) \\
&= \sum_{z \in \mathcal{Z}} \left( \sum_{s \in \mathcal{S}} \sum_t \Pr(S_t = s; \pi) \mathbf{1}\{\phi(s) = z\} \right) \left( \frac{\sum_{s \in \mathcal{S}, a \in \mathcal{A}} \psi^\pi(s) \pi(s, a) r(s, a) \mathbf{1}\{\phi(s) = z\}}{\sum_{s \in \mathcal{S}} \psi^\pi(s) \mathbf{1}\{\phi(s) = z\}} \right) \\
&= \sum_{z \in \mathcal{Z}} \left( \sum_{s \in \mathcal{S}} \kappa \psi^\pi(s) \mathbf{1}\{\phi(s) = z\} \right) \left( \frac{\sum_{s \in \mathcal{S}, a \in \mathcal{A}} \psi^\pi(s) \pi(s, a) r(s, a) \mathbf{1}\{\phi(s) = z\}}{\sum_{s \in \mathcal{S}} \psi^\pi(s) \mathbf{1}\{\phi(s) = z\}} \right) \\
&= \sum_{z \in \mathcal{Z}} \sum_{s \in \mathcal{S}, a \in \mathcal{A}} \kappa \psi^\pi(s) \pi(s, a) r(s, a) \mathbf{1}\{\phi(s) = z\} \\
&\stackrel{(a)}{=} \sum_{s \in \mathcal{S}, a \in \mathcal{A}} \sum_t \Pr(S_t = s; \pi) \pi(s, a) r(s, a) \\
&= J(\pi; M),
\end{aligned}$$

where (a) follows from the law of total probability [64]. □

We have established that ARPs are performance-preserving. Next, we prove Lemma 3.2, which states that the performance estimate obtained from an estimated model of the ARP converges almost surely to the performance of the policy that induces it. In order to do so, we first introduce properties that will be useful in the proof of Lemma 3.2.

## B.2 Estimation of ARPs from On-Policy Data

At a high level, we prove Lemma 3.2 by first studying the almost sure convergence of the estimated transition function, reward function, and initial state distributions. Once the almost sure convergence of these components has been established, we reason about the implications for the convergence of the policy performance predictions that result from these terms. We begin by studying the almost sure convergence of the transition function.

**Property B.2.** For all abstract states  $z \in \mathcal{Z}$  and  $z' \in \mathcal{Z}$ , if

$$\sum_t \Pr(Z_t = z; \pi) \neq 0,$$

then

$$\widehat{\mathbb{P}}_{n, \phi}^\pi(z, z') \xrightarrow{a.s.} \mathbb{P}_\phi^\pi(z, z'). \quad (10)$$

*Proof.* Recall that

$$\widehat{\mathbb{P}}_{n, \phi}^\pi(z, z') = \frac{\sum_{i=1}^n \sum_t \mathbf{1}\{Z_t^i = z, Z_{t+1}^i = z'\}}{\sum_{i=1}^n \sum_t \mathbf{1}\{Z_t^i = z\}}.$$

Let  $X_n := \frac{1}{n} \sum_{i=1}^n \sum_t \mathbf{1}\{Z_t^i = z, Z_{t+1}^i = z'\}$  and  $Y_n := \frac{1}{n} \sum_{i=1}^n \sum_t \mathbf{1}\{Z_t^i = z\}$ . We can then rewrite  $\widehat{\mathbb{P}}_{n, \phi}^\pi(z, z')$  as

$$\widehat{\mathbb{P}}_{n, \phi}^\pi(z, z') = \frac{X_n}{Y_n},$$

which is a continuous function of  $X_n$  and  $Y_n$  when  $Y_n > 0$ . At a high level, we will show what  $X_n$  and  $Y_n$  each converge to almost surely, and will then apply the continuous mapping theorem to reason about the almost sure convergence of  $\widehat{\mathbb{P}}_{n, \phi}^\pi(z, z')$ .

However, there may be some abstract states  $\tilde{\mathcal{Z}} \subset \mathcal{Z}$  that are not reached and thus are not observed in the data. Such abstract states pose a problem: If  $\tilde{z} \in \tilde{\mathcal{Z}}$ , then  $Y_n$  for  $\widehat{\mathbb{P}}_{n, \phi}^\pi(\tilde{z}, z')$  will be zero,

and so  $\widehat{P}_{n,\phi}^\pi(\tilde{z}, z')$  will be undefined and not a continuous function of  $X_n$  and  $Y_n$  (which must be handled appropriately in the proof of consistency). First, to ensure that the ARP is well-defined even when  $\tilde{Z}$  is not empty, for abstract states  $\tilde{z} \in \tilde{Z}$  (abstract states that were not observed in the data), special values can be hardcoded into the transition function, reward functions, and initial state distribution so that the components of the ARP continue to be well-defined.<sup>4</sup> In particular, for  $\tilde{z} \in \tilde{Z}$ , the transition function can be set to self-transition back to  $\tilde{z}$  with probability 1, i.e.,  $\widehat{P}_{n,\phi}^\pi(\tilde{z}, \tilde{z}) = 1$  and  $\widehat{P}_{n,\phi}^\pi(\tilde{z}, \bar{z}) = 0$  for all  $\bar{z} \notin \tilde{Z}$ , the reward function is set to  $\widehat{R}_{n,\phi}^\pi(\tilde{z}) = 0$  and the initial abstract state distribution is  $\widehat{\eta}_{n,\phi}^\pi(\tilde{z}) = 0$  for all  $\tilde{z} \in \tilde{Z}$ .<sup>5</sup>

To reason about the almost sure convergence of  $\widehat{P}_{n,\phi}^\pi(z, z')$ , first consider  $\lim_{n \rightarrow \infty} X_n$ ,

$$\begin{aligned} \lim_{n \rightarrow \infty} X_n &= \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \sum_t \mathbf{1}\{Z_t^i = z, Z_{t+1}^i = z'\} \\ &= \lim_{n \rightarrow \infty} \sum_t \left( \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{Z_t^i = z, Z_{t+1}^i = z'\} \right). \end{aligned}$$

Note that  $\mathbf{1}\{Z_t^i = z, Z_{t+1}^i = z'\}$  is independent across episodes (for each  $i$ ), has bounded variance, and has the same mean for each episode. By Kolmogorov's strong law of large numbers [46],

$$\frac{1}{n} \sum_{i=1}^n \mathbf{1}\{Z_t^i = z, Z_{t+1}^i = z'\} \xrightarrow{\text{a.s.}} \mathbb{E}[\mathbf{1}\{Z_t = z, Z_{t+1} = z'\}; \pi],$$

This convergence guarantee holds for all  $t$ . From the definition of almost sure convergence, this means that

$$\forall t \in \mathbb{N}, \Pr \left( \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{Z_t^i = z, Z_{t+1}^i = z'\} = \mathbb{E}[\mathbf{1}\{Z_t = z, Z_{t+1} = z'\}; \pi] \right) = 1, \quad (11)$$

where the expectation results from the use of the policy  $\pi$  that generated dataset  $\mathcal{D}_n^{(\pi)}$ . By the countable additivity property of probability measures (and the fact that the number of time steps is countable), this implies the following (notice that  $\forall t \in \mathbb{N}$  is now inside the statement of probability):

$$\Pr \left( \forall t \in \mathbb{N}, \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{Z_t^i = z, Z_{t+1}^i = z'\} = \mathbb{E}[\mathbf{1}\{Z_t = z, Z_{t+1} = z'\}; \pi] \right) = 1. \quad (12)$$

Hence,

$$\Pr \left( \sum_t \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{Z_t^i = z, Z_{t+1}^i = z'\} = \sum_t \mathbb{E}[\mathbf{1}\{Z_t = z, Z_{t+1} = z'\}; \pi] \right) = 1. \quad (13)$$

Next, by the dominated convergence theorem [16, Theorem 1.19], the summation over time and limit over  $n$  can be interchanged, giving:

$$\Pr \left( \lim_{n \rightarrow \infty} \sum_t \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{Z_t^i = z, Z_{t+1}^i = z'\} = \sum_t \mathbb{E}[\mathbf{1}\{Z_t = z, Z_{t+1} = z'\}; \pi] \right) = 1. \quad (14)$$

Returning to  $\xrightarrow{\text{a.s.}}$  notation, this means that

$$\sum_t \left( \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{Z_t^i = z, Z_{t+1}^i = z'\} \right) \xrightarrow{\text{a.s.}} \sum_t \mathbb{E}[\mathbf{1}\{Z_t = z, Z_{t+1} = z'\}; \pi]. \quad (15)$$

<sup>4</sup>The particular behavior of the ARP in these unobserved abstract states is not of consequence. Since they are unreached, the behavior in these abstract states does not impact the expected return of the ARP. However, we wish to ensure that the transition function, reward function, and initial state distribution still represent a well-defined ARP.

<sup>5</sup>Note that the initial state distribution is not actually a problem like the transition and reward functions—the previous definitions already ensured that  $\widehat{\eta}_{n,\phi}^\pi(\tilde{z}) = 0$ .

The expected value of the indicator is equal to the probability of the event, and so

$$\sum_t \mathbb{E} [\mathbf{1}\{Z_t = z, Z_{t+1} = z'\}; \pi] = \sum_t \Pr(Z_t = z, Z_{t+1} = z'; \pi). \quad (16)$$

This, combined with the fact that the left hand side of Equation 15 is  $X_n$  means that

$$X_n \xrightarrow{\text{a.s.}} \sum_t \Pr(Z_t = z, Z_{t+1} = z'; \pi). \quad (17)$$

Similarly, these same exact steps can be followed to show that

$$Y_n \xrightarrow{\text{a.s.}} \sum_t \Pr(Z_t = z; \pi). \quad (18)$$

Let  $\Psi_i = (X_i, Y_i)$ , so that  $\Psi_n = (X_i, Y_i)_{i=1}^n$  is a sequence of vector-valued random variables. Since  $X_n \xrightarrow{\text{a.s.}} \sum_t \Pr(Z_t = z, Z_{t+1} = z'; \pi)$  and  $Y_n \xrightarrow{\text{a.s.}} \sum_t \Pr(Z_t = z; \pi)$ , we have that

$$\Psi_n \xrightarrow{\text{a.s.}} \left( \sum_t \Pr(Z_t = z, Z_{t+1} = z'; \pi), \sum_t \Pr(Z_t = z; \pi) \right). \quad (19)$$

Consider the function  $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ , defined as:

$$f(x, y) = \begin{cases} \frac{x}{y} & \text{if } y \neq 0 \\ \widehat{P}_{n,\phi}^\pi(z, z') & \text{otherwise.} \end{cases} \quad (20)$$

Note that discontinuities of  $f$  may occur when  $y = 0$ . Applying the continuous mapping theorem, we have that

$$f(\Psi_n) \xrightarrow{\text{a.s.}} \frac{\sum_t \Pr(Z_t = z, Z_{t+1} = z'; \pi)}{\sum_t \Pr(Z_t = z; \pi)} \quad (21)$$

if<sup>6</sup>

$$\Pr \left( \sum_t \Pr(Z_t = z; \pi) = 0 \right) = 0. \quad (22)$$

In our case,  $\sum_t \Pr(Z_t = z; \pi)$  is a constant, and so the condition in Equation 22 is simply:

$$\sum_t \Pr(Z_t = z; \pi) \neq 0. \quad (23)$$

Also, Equation 21 can be simplified to:

$$\widehat{P}_{n,\phi}^\pi(z, z') \xrightarrow{\text{a.s.}} P_\phi^\pi(z, z'), \quad (24)$$

since  $f(\Psi_n) = X_n/Y_n = \widehat{P}_{n,\phi}^\pi(z, z')$ ,<sup>7</sup> and  $P_\phi^\pi(z, z') = \frac{\sum_t \Pr(Z_t = z, Z_{t+1} = z'; \pi)}{\sum_t \Pr(Z_t = z; \pi)}$  by Equation 3. Restating this conclusion, we have that if

$$\sum_t \Pr(Z_t = z; \pi) \neq 0, \quad (25)$$

then

$$\widehat{P}_{n,\phi}^\pi(z, z') \xrightarrow{\text{a.s.}} P_\phi^\pi(z, z'), \quad (26)$$

which establishes the property.  $\square$

Next, we present a property that establishes the almost sure convergence of the reward function estimate.

<sup>6</sup>This condition arises due to the discontinuity of  $f$  when the second argument is zero. See the statement of the continuous mapping theorem in Appendix A for details.

<sup>7</sup>To be more precise  $f(\Psi_n) = X_n/Y_n$  when  $Y_n \neq 0$ . However, even when  $Y_n = 0$ , the conclusion that  $f(\Psi_n) = \widehat{P}_{n,\phi}^\pi(z, z')$  holds from the definition of  $f$  in Equation 20.

**Property B.3.** For all abstract states  $z \in \mathcal{Z}$ , if

$$\sum_t \Pr(Z_t = z; \pi) \neq 0, \quad (27)$$

then

$$\widehat{R}_{n,\phi}^\pi(z) \xrightarrow{\text{a.s.}} R_\phi^\pi(z). \quad (28)$$

*Proof.* Recall that

$$\widehat{R}_{n,\phi}^\pi(z) = \frac{\sum_{i=1}^n \sum_t \mathbf{1}\{Z_t^i = z\} R_t^i}{\sum_{i=1}^n \sum_t \mathbf{1}\{Z_t^i = z\}}.$$

Let  $X_n := \frac{1}{n} \sum_{i=1}^n \sum_t \mathbf{1}\{Z_t^i = z\} R_t^i$  and  $Y_n := \frac{1}{n} \sum_{i=1}^n \sum_t \mathbf{1}\{Z_t^i = z\}$ . We will show that  $X_n$  and  $Y_n$  converge almost surely, and then apply the continuous mapping theorem to reason about the almost sure convergence of  $\widehat{R}_{n,\phi}^\pi(z)$ .

To ensure that the ARP is well-defined even when  $\tilde{\mathcal{Z}}$  is not empty, for abstract states  $\tilde{z} \in \tilde{\mathcal{Z}}$  (abstract states that were not observed in the data), the reward function can be set to  $\widehat{R}_{n,\phi}^\pi(\tilde{z}) = 0$ , as elaborated in the proof of Property B.2.

Consider  $\lim_{n \rightarrow \infty} X_n$ ,

$$\begin{aligned} \lim_{n \rightarrow \infty} X_n &= \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \sum_t \mathbf{1}\{Z_t^i = z\} R_t^i \\ &= \lim_{n \rightarrow \infty} \sum_t \left( \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{Z_t^i = z\} R_t^i \right). \end{aligned}$$

Note that  $\mathbf{1}\{Z_t^i = z\} R_t^i$  is independent across episodes (for each  $i$ ), has bounded variance, and has the same mean for each episode. By Kolmogorov's strong law of large numbers [46],

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{Z_t^i = z\} R_t^i &\xrightarrow{\text{a.s.}} \mathbb{E}[\mathbf{1}\{Z_t = z\} R_t; \pi] \\ &= \sum_{r \in \mathbb{R}, \tilde{z} \in \mathcal{Z}} \Pr(R_t = r, Z_t = \tilde{z}; \pi) \mathbf{1}\{Z_t = z\} r \\ &= \sum_{r \in \mathbb{R}} \Pr(R_t = r, Z_t = z; \pi) r \\ &= \left( \sum_{r \in \mathbb{R}} r \Pr(R_t = r \mid Z_t = z; \pi) \right) \Pr(Z_t = z; \pi) \\ &= \mathbb{E}[R_t \mid Z_t = z; \pi] \Pr(Z_t = z; \pi). \end{aligned}$$

This convergence guarantee holds for all  $t$ . As in the proof of Property B.2, by the countable additivity property of probability measures, in conjunction with the dominated convergence theorem [16, Theorem 1.19], this implies that,

$$\Pr \left( \lim_{n \rightarrow \infty} \sum_t \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{Z_t^i = z\} R_t^i = \sum_t \mathbb{E}[R_t \mid Z_t = z; \pi] \Pr(Z_t = z; \pi) \right) = 1. \quad (29)$$

Returning to  $\xrightarrow{\text{a.s.}}$  notation, this means that

$$X_n \xrightarrow{\text{a.s.}} \sum_t \mathbb{E}[R_t \mid Z_t = z; \pi] \Pr(Z_t = z; \pi). \quad (30)$$

Similarly, these same exact steps can be followed to show that

$$Y_n \xrightarrow{\text{a.s.}} \sum_t \Pr(Z_t = z; \pi). \quad (31)$$

Let  $\Psi_i = (X_i, Y_i)$ , so that  $\Psi_n = (X_i, Y_i)_{i=1}^n$  is a sequence of vector-valued random variables. Considering the function  $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ , defined as:

$$f(x, y) = \begin{cases} \frac{x}{y} & \text{if } y \neq 0 \\ \widehat{R}_{n,\phi}^\pi(z) & \text{otherwise.} \end{cases} \quad (32)$$

Note that discontinuities of  $f$  may occur when  $y = 0$ . Applying the continuous mapping theorem, we have that

$$f(\Psi_n) \xrightarrow{\text{a.s.}} \frac{\sum_t \mathbb{E}[R_t | Z_t = z; \pi] \Pr(Z_t = z; \pi)}{\sum_t \Pr(Z_t = z; \pi)} \quad (33)$$

if

$$\sum_t \Pr(Z_t = z; \pi) \neq 0. \quad (34)$$

Thus, when  $\sum_t \Pr(Z_t = z; \pi) \neq 0$

$$\widehat{R}_{n,\phi}^\pi(z) \xrightarrow{\text{a.s.}} R_\phi^\pi(z), \quad (35)$$

which establishes the property.  $\square$

Next, we present a property that establishes the almost sure convergence of the initial state distribution.

**Property B.4.** For all abstract states  $z \in \mathcal{Z}$ ,  $\widehat{\eta}_{n,\phi}(z) \xrightarrow{\text{a.s.}} \eta_\phi(z)$ .

*Proof.* Recall that

$$\widehat{\eta}_{n,\phi}(z) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{Z_0^i = z\}.$$

Let  $X_n := \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{Z_0^i = z\}$ . We will show that  $X_n$  converges almost surely, and then apply the continuous mapping theorem to reason about the almost sure convergence of  $\widehat{\eta}_{n,\phi}(z)$ . Note that even when  $\tilde{Z}$  is not empty, the definition of  $\widehat{\eta}_{n,\phi}^\pi$  ensures that it is well-defined.

Consider  $\lim_{n \rightarrow \infty} X_n$ ,

$$\lim_{n \rightarrow \infty} X_n = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{Z_0^i = z\}.$$

Note that  $\mathbf{1}\{Z_0^i = z\}$  is independent across episodes (for each  $i$ ), has bounded variance, and has the same mean for each episode. By Kolmogorov's strong law of large numbers [46],

$$\frac{1}{n} \mathbf{1}\{Z_0^i = z\} \xrightarrow{\text{a.s.}} \mathbb{E}[\mathbf{1}\{Z_0 = z\}] = \Pr(Z_0 = z).$$

In contrast to the proofs of the previous two properties, this proof holds for all  $z \in \mathcal{Z}$ , not just the abstract states that have non-zero probability of occurring. This implies that

$$\widehat{\eta}_{n,\phi}(z) \xrightarrow{\text{a.s.}} \eta_\phi(z), \quad (36)$$

which establishes the property.  $\square$

Having established properties that will be useful in the proof of Lemma 3.2, we now turn to proving the lemma.

**Lemma 3.2.**  $\forall \phi \in \Phi$ , the expected return of the maximum likelihood estimate  $\widehat{\mathfrak{R}}_{n,\phi}^\pi$  converges almost surely to the expected return of the policy  $\pi$ , i.e.,  $J(\pi; \widehat{\mathfrak{R}}_{n,\phi}^\pi) \xrightarrow{\text{a.s.}} J(\pi; M)$ .

*Proof.* From Theorem 3.1, we have that  $J(\pi; \mathfrak{R}_\phi^\pi) = J(\pi; M)$ . Thus, to prove this result, we only need to show that  $J(\pi; \widehat{\mathfrak{R}}_{n,\phi}^\pi) \xrightarrow{\text{a.s.}} J(\pi; \mathfrak{R}_\phi^\pi)$ . Several of the preliminary results required to establish this result were provided in Properties B.2, B.3, and B.4. Specifically, Properties B.2 and B.3 establish that for all abstract states  $z \in \mathcal{Z}$  and  $z' \in \mathcal{Z}$ , if

$$\sum_t \Pr(Z_t = z; \pi) \neq 0, \quad (37)$$

then

$$\widehat{\mathbb{P}}_{n,\phi}^\pi(z, z') \xrightarrow{\text{a.s.}} \mathbb{P}_\phi^\pi(z, z'), \quad (38)$$

and

$$\widehat{\mathbb{R}}_{n,\phi}^\pi(z) \xrightarrow{\text{a.s.}} \mathbb{R}_\phi^\pi(z). \quad (39)$$

Similarly, Property B.4 establishes that for all abstract states  $z \in \mathcal{Z}$ ,

$$\widehat{\eta}_{n,\phi}(z) \xrightarrow{\text{a.s.}} \eta_\phi(z). \quad (40)$$

The expected return  $J(\pi; \widehat{\mathfrak{R}}_{n,\phi}^\pi)$  is a continuous function of  $\widehat{\mathfrak{R}}_{n,\phi}^\pi := (\widehat{\mathbb{P}}_{n,\phi}^\pi, \widehat{\mathbb{R}}_{n,\phi}^\pi, \widehat{\eta}_{n,\phi})$ , given by  $J(\pi; \widehat{\mathfrak{R}}_{n,\phi}^\pi) = \left( \mathbb{I} - \widehat{\mathbb{P}}_{n,\phi}^\pi \right)^{-1} \widehat{\mathbb{R}}_{n,\phi}^\pi \widehat{\eta}_{n,\phi} = \sum_z \sum_t \widehat{\text{Pr}}(Z_t = z; \pi) \widehat{\mathbb{R}}_{n,\phi}^\pi(z)$ <sup>8</sup> where  $\widehat{\text{Pr}}$  denotes the empirical estimate of the corresponding probability.

Note that the term  $\sum_t \widehat{\text{Pr}}(Z_t = z; \pi) = \left[ \left( \mathbb{I} - \widehat{\mathbb{P}}_{n,\phi}^\pi \right)^{-1} \widehat{\eta}_{n,\phi} \right]_z$ , i.e., the empirical estimate of the sum of probabilities of encountering the abstract state  $z$  over all timesteps is equal to value of the vector  $\left( \mathbb{I} - \widehat{\mathbb{P}}_{n,\phi}^\pi \right)^{-1} \widehat{\eta}_{n,\phi}$  at  $z$ . By the continuous mapping theorem, if  $\sum_t \text{Pr}(Z_t = z; \pi) \neq 0$ ,

$$\sum_t \widehat{\text{Pr}}(Z_t = z; \pi) \xrightarrow{\text{a.s.}} \sum_t \text{Pr}(Z_t = z; \pi). \quad (41)$$

Let  $\bar{\mathcal{Z}} = \{z : \sum_t \text{Pr}(Z_t = z; \pi) = 0\}$ , where  $\bar{\mathcal{Z}} \subset \mathcal{Z}$ , be the set of abstract states that will never be observed empirically as they have no probability of being visited under  $\pi$ .<sup>9</sup> The expression for the expected return can be divided into two terms: (1) the first term sums over abstract states for which  $\sum_t \text{Pr}(Z_t = z; \pi) = 0$ , and the values of the components of the ARP—in particular,  $\widehat{\mathbb{R}}_{n,\phi}^\pi$ —have been specially hardcoded as detailed in Property B.2, and (2) the second term sums over the abstract states for which  $\sum_t \text{Pr}(Z_t = z; \pi) \neq 0$  and thus Equation 41 holds.

$$J(\pi; \widehat{\mathfrak{R}}_{n,\phi}^\pi) = \sum_{z \in \bar{\mathcal{Z}}} \sum_t \widehat{\text{Pr}}(Z_t = z; \pi) \underbrace{\widehat{\mathbb{R}}_{n,\phi}^\pi(z)}_{=0} + \sum_{z \in \mathcal{Z} \setminus \bar{\mathcal{Z}}} \sum_t \widehat{\text{Pr}}(Z_t = z; \pi) \widehat{\mathbb{R}}_{n,\phi}^\pi(z) \quad (42)$$

$$= \sum_{z \in \mathcal{Z} \setminus \bar{\mathcal{Z}}} \sum_t \widehat{\text{Pr}}(Z_t = z; \pi) \widehat{\mathbb{R}}_{n,\phi}^\pi(z) \quad (43)$$

As detailed in Property B.2,  $\widehat{\mathbb{R}}_{n,\phi}^\pi(z) = 0$  for  $z \in \bar{\mathcal{Z}}$  and thus  $\forall z \in \bar{\mathcal{Z}}$ , making the first term equal to zero. The constituents of the second term converge almost surely:

$$\sum_{z \in \mathcal{Z} \setminus \bar{\mathcal{Z}}} \sum_t \widehat{\text{Pr}}(Z_t = z; \pi) \widehat{\mathbb{R}}_{n,\phi}^\pi(z) \xrightarrow{\text{a.s.}} \sum_{z \in \mathcal{Z} \setminus \bar{\mathcal{Z}}} \sum_t \text{Pr}(Z_t = z; \pi) \mathbb{R}_\phi^\pi(z) \quad (44)$$

$$= \sum_{z \in \mathcal{Z} \setminus \bar{\mathcal{Z}}} \sum_t \text{Pr}(Z_t = z; \pi) \mathbb{R}_\phi^\pi(z) + 0 \quad (45)$$

$$= \sum_{z \in \mathcal{Z} \setminus \bar{\mathcal{Z}}} \sum_t \text{Pr}(Z_t = z; \pi) \mathbb{R}_\phi^\pi(z) + \sum_{z \in \bar{\mathcal{Z}}} \underbrace{\left( \sum_t \text{Pr}(Z_t = z; \pi) \right)}_{=0} \mathbb{R}_\phi^\pi(z) \quad (46)$$

$$= \sum_{z \in \mathcal{Z}} \sum_t \text{Pr}(Z_t = z; \pi) \mathbb{R}_\phi^\pi(z) \quad (47)$$

$$= J(\pi; \mathfrak{R}_\phi^\pi). \quad (48)$$

We then obtain the final result:

$$J(\pi; \widehat{\mathfrak{R}}_{n,\phi}^\pi) \xrightarrow{\text{a.s.}} J(\pi; \mathfrak{R}_\phi^\pi) = J(\pi; M). \quad (49)$$

□

<sup>8</sup>Termination is handled as detailed in Appendix A.1.

<sup>9</sup>In relation to  $\bar{\mathcal{Z}}$  defined in Property B.2, note  $\bar{\mathcal{Z}} \subseteq \bar{\mathcal{Z}}$ . As  $n \rightarrow \infty$  the two sets, with probability 1, become equal.



### B.3 Estimation of ARPs from Off-Policy Data

So far, we have shown that the expected return of the estimated ARP converges almost surely to the expected return of the policy that induced it, i.e., consistent *on-policy* evaluation. Next, we show that a model of the ARP can be consistently estimated from *off-policy* data. This requires us to prove that transition functions estimated using off-policy data, incorporating importance weights in their estimation, converge almost surely to the true transition functions induced by the evaluation policy, and that the same holds for the reward functions and initial state distributions.

**Property B.6.** For all abstract states  $z \in \mathcal{Z}$  and  $z' \in \mathcal{Z}$ , if

$$\sum_t \Pr(Z_t = z; \pi_b) \neq 0, \quad (50)$$

then

$$\hat{\mathbb{P}}_{n,\phi}^{\pi_b \rightarrow \pi_e}(z, z') \xrightarrow{\text{a.s.}} \mathbb{P}_{\phi}^{\pi_e}(z, z'). \quad (51)$$

*Proof.* Recall that

$$\hat{\mathbb{P}}_{n,\phi}^{\pi_b \rightarrow \pi_e}(z, z') = \frac{\sum_{i,t} \mathbf{1}_t^i\{z, z'\} \rho_{0:t}}{\sum_{i,t} \mathbf{1}_t^i\{z\} \rho_{0:t}}.$$

Similar to the proof structure of Property B.2, let  $X_n := \frac{1}{n} \sum_{i,t} \mathbf{1}_t^i\{z, z'\} \rho_{0:t}$  and  $Y_n := \frac{1}{n} \sum_{i,t} \mathbf{1}_t^i\{z\} \rho_{0:t}$ . For abstract states  $\tilde{\mathcal{Z}} \subset \mathcal{Z}$  that are never reached and thus are not observed in the data, special values can be hardcoded into the transition function, reward functions, and initial state distribution so that the components of the ARP continue to be well-defined. In particular, for  $\tilde{z} \in \tilde{\mathcal{Z}}$ , the transition function can be set to self-transition back to  $\tilde{z}$  with probability 1, i.e.,  $\hat{\mathbb{P}}_{n,\phi}^{\pi_b \rightarrow \pi_e}(\tilde{z}, \tilde{z}) = 1$  and  $\hat{\mathbb{P}}_{n,\phi}^{\pi_b \rightarrow \pi_e}(\tilde{z}, \tilde{z}') = 0$  for all  $\tilde{z}' \notin \tilde{\mathcal{Z}}$ , the reward function is set to  $\hat{\mathbb{R}}_{n,\phi}^{\pi_b \rightarrow \pi_e}(\tilde{z}) = 0$  and the initial abstract state distribution is  $\hat{\eta}_{n,\phi}^{\pi_b \rightarrow \pi_e}(\tilde{z}) = 0$  for all  $\tilde{z} \in \tilde{\mathcal{Z}}$ .

Consider  $\lim_{n \rightarrow \infty} X_n$ ,

$$\begin{aligned} \lim_{n \rightarrow \infty} X_n &= \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i,t} \mathbf{1}_t^i\{z, z'\} \rho_{0:t} \\ &= \lim_{n \rightarrow \infty} \sum_t \left( \frac{1}{n} \sum_{i=1}^n \mathbf{1}_t^i\{z, z'\} \rho_{0:t} \right). \end{aligned}$$

Note that  $\mathbf{1}_t^i\{z, z'\} \rho_{0:t}$  is independent across episodes (for each  $i$ ), has bounded variance, and has the same mean for each episode. By Kolmogorov's strong law of large numbers [46],

$$\frac{1}{n} \sum_{i=1}^n \mathbf{1}_t^i\{z, z'\} \rho_{0:t} \xrightarrow{\text{a.s.}} \mathbb{E}[\mathbf{1}\{Z_t = z, Z_{t+1} = z'\} \rho_{0:t}; \pi_b].$$

The importance weights change the distribution over which the expectation is computed, as below:

$$\begin{aligned} &\mathbb{E}[\mathbf{1}\{Z_t = z, Z_{t+1} = z'\} \rho_{0:t}; \pi_b] \\ &= \sum_t \sum_{\substack{S_{0:t+1}, \\ A_{0:t}}} \left( \eta(S_0) \prod_{l=0}^t \pi_b(S_l, A_l) p(S_l, A_l, S_{l+1}) \right) \mathbf{1}\{\phi(S_t) = z, \phi(S_{t+1}) = z'\} \left( \prod_{j=0}^t \frac{\pi_e(S_j, A_j)}{\pi_b(S_j, A_j)} \right) \\ &= \sum_t \sum_{\substack{S_{0:t+1}, \\ A_{0:t}}} \left( \eta(S_0) \prod_{l=0}^t \pi_e(S_l, A_l) p(S_l, A_l, S_{l+1}) \right) \mathbf{1}\{\phi(S_t) = z, \phi(S_{t+1}) = z'\} \\ &= \sum_t \mathbb{E}[\mathbf{1}\{Z_t = z, Z_{t+1} = z'\}; \pi_e] \end{aligned}$$

This convergence guarantee holds for all  $t$ . By the countable additivity property of probability measures in conjunction with the dominated convergence theorem [16], this implies that,

$$X_n \xrightarrow{\text{a.s.}} \sum_t \mathbb{E}[\mathbf{1}\{Z_t = z, Z_{t+1} = z'\}; \pi_e]. \quad (52)$$

Similarly, these same exact steps can be followed to show that

$$Y_n \xrightarrow{\text{a.s.}} \sum_t \mathbb{E}[\mathbf{1}\{Z_t = z\}; \pi_e]. \quad (53)$$

Let  $\Psi_i = (X_i, Y_i)$ , so that  $\Psi_n = (X_i, Y_i)_{i=1}^n$  is a sequence of vector-valued random variables. Considering the function  $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ , defined as:

$$f(x, y) = \begin{cases} \frac{x}{y} & \text{if } y \neq 0 \\ \widehat{\mathbf{P}}_{n, \phi}^{\pi_b \rightarrow \pi_e}(z, z') & \text{otherwise.} \end{cases} \quad (54)$$

Note that discontinuities of  $f$  may occur when  $y = 0$ . Applying the continuous mapping theorem, we have that

$$f(\Psi_n) \xrightarrow{\text{a.s.}} \frac{\sum_t \mathbb{E}[\mathbf{1}\{Z_t = z, Z_{t+1} = z'\}; \pi_e]}{\sum_t \mathbb{E}[\mathbf{1}\{Z_t = z\}; \pi_e]} \quad (55)$$

if  $\sum_t \mathbb{E}[\mathbf{1}\{Z_t = z\}; \pi_e] \neq 0$ . This implies that

$$\widehat{\mathbf{P}}_{n, \phi}^{\pi_b \rightarrow \pi_e}(z, z') \xrightarrow{\text{a.s.}} \mathbf{P}_{\phi}^{\pi_e}(z, z'), \quad (56)$$

when  $\sum_t \Pr(Z_t = z; \pi_b) \neq 0$ , which establishes the property.  $\square$

Next, we present the properties that establish almost sure convergence of the reward function estimate and the initial state distribution estimated from off-policy data. This follows a similar proof structure to their corresponding on-policy properties, with the key difference being the use of importance weights in the estimation.

**Property B.7.** For all abstract states  $z \in \mathcal{Z}$ , if

$$\sum_t \Pr(Z_t = z; \pi_b) \neq 0, \quad (57)$$

then

$$\widehat{\mathbf{R}}_{n, \phi}^{\pi_b \rightarrow \pi_e}(z) \xrightarrow{\text{a.s.}} \mathbf{R}_{\phi}^{\pi_e}(z). \quad (58)$$

*Proof.* Recall that

$$\widehat{\mathbf{R}}_{n, \phi}^{\pi_b \rightarrow \pi_e}(z) = \frac{\sum_{i=1}^n \sum_t \mathbf{1}\{Z_t^i = z\} R_t^i \rho_{0:t}}{\sum_{i=1}^n \sum_t \mathbf{1}\{Z_t^i = z\} \rho_{0:t}}.$$

Let  $X_n := \frac{1}{n} \sum_{i=1}^n \sum_t \mathbf{1}\{Z_t^i = z\} R_t^i \rho_{0:t}$  and  $Y_n := \frac{1}{n} \sum_{i=1}^n \sum_t \mathbf{1}\{Z_t^i = z\} \rho_{0:t}$ . Following the exact steps from Property B.3, we can show that

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{Z_t^i = z\} R_t^i &\xrightarrow{\text{a.s.}} \mathbb{E}[\mathbf{1}\{Z_t = z\} R_t \rho_{0:t}; \pi_b] \\ &= \mathbb{E}[\mathbf{1}\{Z_t = z\} R_t; \pi_e] \\ &= \sum_{r \in \mathbb{R}, z \in \mathcal{Z}} \Pr(R_t = r, Z_t = z; \pi_e) \mathbf{1}\{Z_t = z\} r \\ &= \sum_{r \in \mathbb{R}} \Pr(R_t = r, Z_t = z; \pi_e) r \\ &= \left( \sum_{r \in \mathbb{R}} r \Pr(R_t = r \mid Z_t = z; \pi_e) \right) \Pr(Z_t = z; \pi) \\ &= \mathbb{E}[R_t \mid Z_t = z; \pi] \Pr(Z_t = z; \pi_e). \end{aligned}$$

This convergence guarantee holds for all  $t$ . By the countable additivity property of probability measures, in conjunction with the dominated convergence theorem, this implies that

$$X_n \xrightarrow{\text{a.s.}} \sum_t \mathbb{E}[R_t \mid Z_t = z; \pi_e] \Pr(Z_t = z; \pi_e). \quad (59)$$

Similarly, these same exact steps can be followed to show that

$$Y_n \xrightarrow{\text{a.s.}} \sum_t \Pr(Z_t = z; \pi_e). \quad (60)$$

Noting that  $\widehat{R}_{n,\phi}^{\pi_b \rightarrow \pi_e} = \frac{X_n}{Y_n}$ , we can apply the continuous mapping theorem, as done in the proof of Property B.3, to show that

$$\widehat{R}_{n,\phi}^{\pi_b \rightarrow \pi_e}(z) \xrightarrow{\text{a.s.}} R_\phi^{\pi_e}(z), \quad (61)$$

when  $\sum_t \Pr(Z_t = z; \pi_b) \neq 0$ .  $\square$

**Property B.8.** For all abstract states  $z \in \mathcal{Z}$ ,  $\widehat{\eta}_{n,\phi}^{\pi_b \rightarrow \pi_e}(z) \xrightarrow{\text{a.s.}} \eta_\phi^{\pi_e}(z)$ .

*Proof.* Recall that

$$\widehat{\eta}_{n,\phi}^{\pi_b \rightarrow \pi_e}(z) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{Z_0^i = z\} \rho_0.$$

Let  $X_n := \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{Z_0^i = z\} \rho_0$ . Following the exact steps from Property B.4, we can show that

$$\frac{1}{n} \mathbf{1}\{Z_t^i = z\} \xrightarrow{\text{a.s.}} \mathbb{E}[\mathbf{1}\{Z_0 = z\} \rho_0] = \Pr(Z_0 = z).$$

This proof holds for all  $z \in \mathcal{Z}$ , not just the abstract states that have non-zero probability of occurring. This implies that

$$\widehat{\eta}_{n,\phi}^{\pi_b \rightarrow \pi_e}(z) \xrightarrow{\text{a.s.}} \eta_\phi^{\pi_e}(z). \quad (62)$$

$\square$

**Lemma 3.3.** Under Assumption 2.1, the weighted maximum likelihood estimate  $\widehat{\mathfrak{R}}_{n,\phi}^{\pi_b \rightarrow \pi_e}$  converges almost surely to the ground-truth ARP  $\mathfrak{R}_\phi^{\pi_e}$ , i.e.,  $\widehat{\mathfrak{R}}_{n,\phi}^{\pi_b \rightarrow \pi_e} \xrightarrow{\text{a.s.}} \mathfrak{R}_\phi^{\pi_e}$ .

*Proof.* The results required to establish this result are provided in Properties B.6, B.7, and B.8. Specifically, the properties establish that if

$$\sum_t \Pr(Z_t = z; \pi_b) \neq 0,$$

then

$$\widehat{P}_{n,\phi}^{\pi_b \rightarrow \pi_e}(z, z') \xrightarrow{\text{a.s.}} P_\phi^{\pi_e}(z, z'); \quad \widehat{R}_{n,\phi}^{\pi_b \rightarrow \pi_e}(z) \xrightarrow{\text{a.s.}} R_\phi^{\pi_e}(z); \quad \widehat{\eta}_{n,\phi}^{\pi_b \rightarrow \pi_e}(z) \xrightarrow{\text{a.s.}} \eta_\phi(z),$$

giving the final result  $\widehat{\mathfrak{R}}_{n,\phi}^{\pi_b \rightarrow \pi_e} \xrightarrow{\text{a.s.}} \mathfrak{R}_\phi^{\pi_e}$ , while the values of the components for  $z \in \tilde{\mathcal{Z}}$  are hardcoded as detailed in Property B.6.  $\square$

#### B.4 Off-Policy Evaluation with ARPs

**Theorem 4.1.** The expected return of the ARP  $\widehat{\mathfrak{R}}_{n,\phi}^{\pi_b \rightarrow \pi_e}$  (built from off-policy data) converges almost surely to the expected return of  $\pi_e$ , i.e.,  $J(\pi_e; \widehat{\mathfrak{R}}_{n,\phi}^{\pi_b \rightarrow \pi_e}) \xrightarrow{\text{a.s.}} J(\pi_e; M)$ .

*Proof.* The expected return  $J(\pi_e; \widehat{\mathfrak{R}}_{n,\phi}^{\pi_b \rightarrow \pi_e})$  is a continuous function of  $\widehat{\mathfrak{R}}_{n,\phi}^{\pi_b \rightarrow \pi_e} := (\widehat{P}_{n,\phi}^{\pi_b \rightarrow \pi_e}, \widehat{R}_{n,\phi}^{\pi_b \rightarrow \pi_e}, \widehat{\eta}_{n,\phi}^{\pi_b \rightarrow \pi_e})$ , given by  $J(\pi_e; \widehat{\mathfrak{R}}_{n,\phi}^{\pi_b \rightarrow \pi_e}) = \left( I - \widehat{P}_{n,\phi}^{\pi_b \rightarrow \pi_e} \right)^{-1} \widehat{R}_{n,\phi}^{\pi_b \rightarrow \pi_e} \widehat{\eta}_{n,\phi}^{\pi_b \rightarrow \pi_e}$ . Similar to the proof of Lemma 3.2, by the continuous mapping theorem and Lemma 3.3, we obtain the final result that enables off-policy evaluation with ARPs:

$$J(\pi_e; \widehat{\mathfrak{R}}_{n,\phi}^{\pi_b \rightarrow \pi_e}) \xrightarrow{\text{a.s.}} J(\pi_e; \mathfrak{R}_\phi^{\pi_e}) = J(\pi_e; M).$$

$\square$

## B.5 Variance Reduction: Leveraging Markovness of the State Abstraction

Variance in estimation of the model of the ARP from off-policy can be reduced by clipping the importance weights. We now show that when the state abstraction function is  $c$ -th order Markov, clipping the importance weights to the  $c$  most recent terms continues to provide a consistent off-policy expected return estimate.

**Theorem 4.3.** *Given a  $c$ -th order Markov  $\phi$ , the expected return of the abstract reward process  $\widehat{\mathfrak{R}}_{\phi,c}^{\pi_b \rightarrow \pi_e}$  converges almost surely to the expected return of  $\pi_e$ , i.e.,  $J(\pi_e; \widehat{\mathfrak{R}}_{\phi,c}^{\pi_b \rightarrow \pi_e}) \xrightarrow{\text{a.s.}} J(\pi_e; M)$ .*

*Proof.* We only need to show that for a state abstraction function  $\phi \in \Phi$  that is  $c$ -th order Markov, i.e.,

$$\Pr(\phi(S_{t+1}) | \phi(S_t), \phi(S_{t-1}), \dots, \phi(S_{(t-c+1)^+}); \pi) = \Pr(\phi(S_{t+1}) | \phi(S_t), \phi(S_{t-1}), \dots, \phi(S_1), \phi(S_0); \pi),$$

for  $\pi \in \{\pi_b, \pi_e\}$  the following holds:

$$\widehat{\mathfrak{R}}_{\phi,c}^{\pi_b \rightarrow \pi_e} \xrightarrow{\text{a.s.}} \mathfrak{R}_{\phi}^{\pi_e}.$$

The remaining steps to show that the return estimates are consistent follow from Theorem 4.1. For brevity, we denote by  $Z_t := \phi(S_t)$  for the rest of the proof. Consider the expression for the transition function without any weight clipping,

$$\widehat{\mathbb{P}}_{n,\phi}^{\pi_b \rightarrow \pi_e}(z, z') := \frac{\sum_{i,t} \mathbf{1}_t^i\{z, z'\} \rho_{0:t}}{\sum_{i,t} \mathbf{1}_t^i\{z\} \rho_{0:t}}.$$

As shown in Property B.6, if  $\sum_t \Pr(Z_t = z; \pi_b) \neq 0$ , as  $n \rightarrow \infty$  the numerator of  $\widehat{\mathbb{P}}_{n,\phi}^{\pi_b \rightarrow \pi_e}(z, z')$  converges almost surely to  $\sum_t \mathbb{E}[\mathbf{1}\{Z_t = z, Z_{t+1} = z'\} \rho_{0:t}; \pi_b]$ . When the abstract states are  $c$ -th order Markov,

$$\rho_{0:(t-c)^+} \perp \rho_{(t-c+1)^+:t}, \text{ conditioned on } (Z_i)_{i=t-c+1}^t. \quad (63)$$

Using the tower rule [60],  $\mathbb{E}[\mathbf{1}\{Z_t = z, Z_{t+1} = z'\} \rho_{0:t}; \pi_b]$  can be re-written as:

$$\begin{aligned} & \mathbb{E}[\mathbf{1}\{Z_t = z, Z_{t+1} = z'\} \rho_{0:t}; \pi_b] \\ &= \mathbb{E}[\mathbf{1}\{Z_t = z, Z_{t+1} = z'\} \rho_{0:(t-c)^+} \rho_{(t-c+1)^+:t}; \pi_b] \\ &= \mathbb{E}\left[\mathbb{E}\left[\mathbf{1}\{Z_t = z, Z_{t+1} = z'\} \rho_{0:(t-c)^+} \rho_{(t-c+1)^+:t} \mid (Z_i)_{i=(t-c+1)^+}^t; \pi_b\right]; \pi_b\right] \\ &\stackrel{(a)}{=} \mathbb{E}\left[\mathbb{E}\left[\mathbf{1}\{Z_t = z, Z_{t+1} = z'\} \rho_{(t-c+1)^+:t} \mid (Z_i)_{i=(t-c+1)^+}^t; \pi_b\right] \mathbb{E}\left[\rho_{0:(t-c)^+} \mid (Z_i)_{i=(t-c+1)^+}^t; \pi_b\right]; \pi_b\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[\mathbf{1}\{Z_t = z, Z_{t+1} = z'\} \rho_{(t-c+1)^+:t} \mid (Z_i)_{i=(t-c+1)^+}^t; \pi_b\right]; \pi_b\right] \\ & \quad \mathbb{E}\left[\mathbb{E}\left[\rho_{0:(t-c)^+} \mid (Z_i)_{i=(t-c+1)^+}^t; \pi_b\right]; \pi_b\right] \\ &= \mathbb{E}[\mathbf{1}\{Z_t = z, Z_{t+1} = z'\} \rho_{(t-c+1)^+:t}; \pi_b] \underbrace{\mathbb{E}[\rho_{0:(t-c)^+}; \pi_b]}_{=1} \\ &\stackrel{(b)}{=} \Pr(Z_t = z, Z_{t+1} = z'; \pi_e), \end{aligned}$$

where (a) follows from Equation 63 and (b) follows from Lemma 3.3. Thus when the abstract states are  $c$ -th order Markov, the numerator estimated with  $c$ -clipped importance weights almost surely converges to the same value as if the importance weights were not clipped. The value to which the denominator converges similarly undergoing a change of measure by the use of clipped importance weights, i.e.,

$$\mathbb{E}[\mathbf{1}\{Z_t = z\} \rho_{0:t}; \pi_b] = \mathbb{E}[\mathbf{1}\{Z_t = z\} \rho_{(t-c+1)^+:t}; \pi_b] = \Pr(Z_t = z; \pi_e), \quad (64)$$

when  $\phi$  is  $c$ -th order Markov. Note that the expression for  $\widehat{\mathbb{P}}_{\phi,c}^{\pi_b \rightarrow \pi_e}$  is:

$$\widehat{\mathbb{P}}_{\phi,c}^{\pi_b \rightarrow \pi_e}(z, z') := \frac{\sum_{i,t} \mathbf{1}_t^i\{z, z'\} \rho_{(t-c+1)^+:t}}{\sum_{i,t} \mathbf{1}_t^i\{z\} \rho_{(t-c+1)^+:t}}.$$

Let  $X_n := \frac{1}{n} \sum_{i,t} \mathbf{1}_t^i\{z, z'\} \rho_{(t-c+1)^+:t}$  and  $Y_n := \frac{1}{n} \sum_{i,t} \mathbf{1}_t^i\{z\} \rho_{(t-c+1)^+:t}$ . We can write  $\widehat{P}_{\phi,c}^{\pi_b \rightarrow \pi_e}(z, z') = \frac{X_n}{Y_n}$ . Following the exact steps from Lemma 3.3, that entail a careful application of the continuous mapping theorem, we can show that

$$\widehat{P}_{\phi,c}^{\pi_b \rightarrow \pi_e}(z, z') \xrightarrow{\text{a.s.}} P_{\phi}^{\pi_e}(z, z'),$$

if  $\sum_t \Pr(Z_t = z; \pi_b) \neq 0$ . Similar derivations can be followed for the reward function and the initial state distribution, leading to the required result:

$$\widehat{\mathfrak{R}}_{\phi,c}^{\pi_b \rightarrow \pi_e} \xrightarrow{\text{a.s.}} \mathfrak{R}_{\phi}^{\pi_e}$$

when  $\sum_t \Pr(Z_t = z; \pi_b) \neq 0$ , that enables almost sure convergence of the expected return estimate:

$$J(\pi_e; \widehat{\mathfrak{R}}_{\phi,c}^{\pi_b \rightarrow \pi_e}) \xrightarrow{\text{a.s.}} J(\pi_e; \mathfrak{R}_{\phi}^{\pi_e}) = J(\pi_e; M).$$

□

## C Empirical Details

In this section we provide additional empirical details for the experiments presented in the main paper. An overall step-by-step algorithm for STAR is as follows:

---

### Algorithm 1 Overview of STAR( $\phi, c$ )

---

**Input:**  $\pi_e, \pi_b, \mathcal{D}_n^{(\pi_b)}$

1. Apply state abstraction to  $\mathcal{D}_n^{(\pi_b)}$  and compute importance weights:

$$\forall i, t : \text{Store} \left( Z_t^{(i)} = \phi \left( S_t^{(i)} \right), A_t^{(i)}, R_t^{(i)}, \rho_{(t-c+1)^+:t}^{(i)} = \prod_{j=(t-c+1)^+}^t \frac{\pi_e(S_j^{(i)}, A_j^{(i)})}{\pi_b(S_j^{(i)}, A_j^{(i)})} \right)$$

2. Estimate the components of the ARP  $\hat{\mathfrak{R}}_{\phi,c}^{\pi_b \rightarrow \pi_e}$ :

$$\hat{\mathfrak{R}}_{\phi,c}^{\pi_b \rightarrow \pi_e} = \left( \hat{\mathbb{P}}_{\phi,c}^{\pi_b \rightarrow \pi_e}, \hat{\mathbb{R}}_{\phi,c}^{\pi_b \rightarrow \pi_e}, \hat{\eta}_{\phi,c}^{\pi_b \rightarrow \pi_e} \right) \text{ where,}$$

$$\hat{\mathbb{P}}_{\phi,c}^{\pi_b \rightarrow \pi_e} = \frac{\sum_{i,t} \mathbf{1}\{\phi(S_t^{(i)})=z, \phi(S_{t+1}^{(i)})=z'\} \rho_{(t-c+1)^+:t}^{(i)}}{\sum_{i,t} \mathbf{1}\{\phi(S_t^{(i)})=z\} \rho_{(t-c+1)^+:t}^{(i)}},$$

$$\hat{\mathbb{R}}_{\phi,c}^{\pi_b \rightarrow \pi_e} = \frac{\sum_{i,t} \mathbf{1}\{\phi(S_t^{(i)})=z\} \rho_{(t-c+1)^+:t}^{(i)} R_t^{(i)}}{\sum_{i,t} \mathbf{1}\{\phi(S_t^{(i)})=z\} \rho_{(t-c+1)^+:t}^{(i)}},$$

$$\hat{\eta}_{\phi,c}^{\pi_b \rightarrow \pi_e}(z) = \frac{\sum_{i=1}^n \mathbf{1}\{\phi(S_0^{(i)})=z\}}{n}$$

3. Compute the expected return  $J(\pi_e; \hat{\mathfrak{R}}_{\phi,c}^{\pi_b \rightarrow \pi_e})$  from the ARP.

$$J(\pi_e; \hat{\mathfrak{R}}_{\phi,c}^{\pi_b \rightarrow \pi_e}) := (\mathbf{I} - \hat{\mathbb{P}}_{\phi,c}^{\pi_b \rightarrow \pi_e})^{-1} \hat{\mathbb{R}}_{\phi,c}^{\pi_b \rightarrow \pi_e} \hat{\eta}_{\phi,c}^{\pi_b \rightarrow \pi_e}$$

**Output:**  $J(\pi_e; \hat{\mathfrak{R}}_{\phi,c}^{\pi_b \rightarrow \pi_e})$

---

### C.1 Domains

We perform empirical evaluations on a range of domains that consist of continuous domains and domains with large state spaces with long horizons. The domains are as follows:

**CartPole** The CartPole domain is a classic control problem from OpenAI Gym [6]. The task is to balance a pole on a cart by moving the cart left or right. The state space is continuous, and the action space is discrete. We use the standard CartPole environment from OpenAI Gym. In the experiments,  $\pi_b$  is uniformly random, i.e., the left and right actions are each taken with probability 0.5. The evaluation policy  $\pi_e$  is a policy that takes the action right with probability 0.9 when the pole is leaning left, and right with probability 0.1 when the pole is leaning right. This results in a policy that is not optimal, but is somewhat successful at balancing the pole.

**ICU-Sepsis** The ICU-Sepsis domain simulates the treatment of sepsis in the ICU. Built from the MIMIC-III database [24] and drawing from the analysis of Komorowski et al. [26], it consists of 747 states that denote the status of a patient and 25 possible actions that denote possible medical interventions. At the end of each episode, if the patient survives, a reward of +1 is given, while death corresponds to a reward of 0, with all intermediate rewards also being 0. This results in the expected return of a policy corresponding to the probability that a randomly selected patient will survive. In the experiments,  $\pi_e$  is set to an *expert policy*, provided with the domain and included in the submitted codebase, while  $\pi_b$  is a policy that is a more stochastic version of the expert policy constructed by temperature scaling [17] the action probabilities of expert policy with a temperature parameter  $\tau = 2$ .

**Asterix** The Asterix domain is a miniaturized version of the Atari game Asterix where the task is to collect items while avoiding enemies. We use the implementation of the game from the MinAtar testbed [61], where the dimension of each state is  $10 \times 10 \times 4$ , and the action set consists of six actions. The data collecting policy  $\pi_b$  is uniformly random, while the evaluation policy  $\pi_e$  picks actions with non-uniform skewed probabilities.

## C.2 Additional Results for Estimator Selection

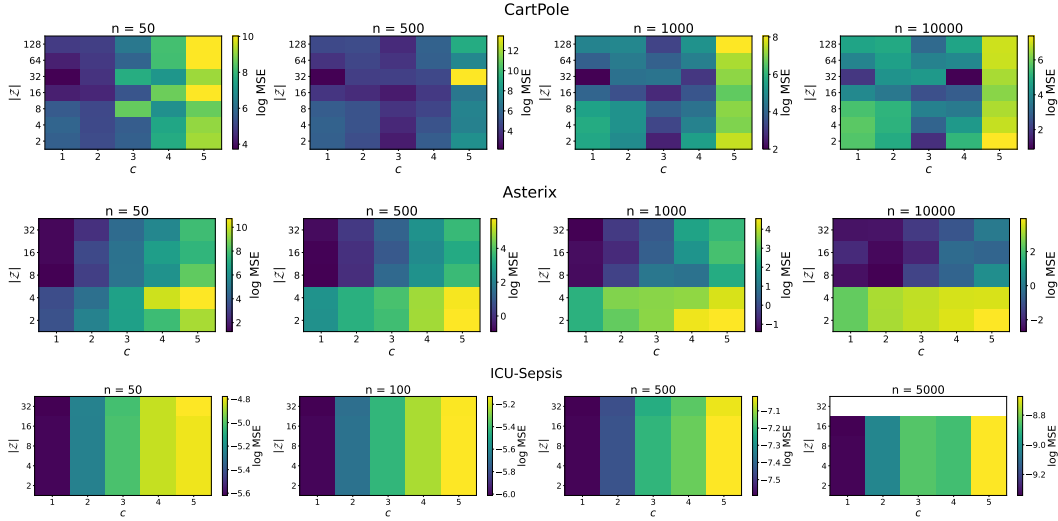


Figure 4: Heatmap of the log mean squared error (MSE) of the OPE estimators for the CartPole, Asterix, and ICU-Sepsis domains. The vertical axis represents the number of abstract states  $|\mathcal{Z}|$ , and the horizontal axis represents the value of the hyperparameter  $c$ . The color intensity indicates the log MSE, with lower values denoting better performance. Note that the variation with  $|\mathcal{Z}|$  is strongly influenced by the class of abstraction functions used, which in this work is CluSTAR.

For each domain, we evaluate the OPE performance of the estimated ARP induced by varying configurations of  $(\phi, c)$ . For the class of abstraction function, we observe that the simple method CluSTAR performs well across all domains, and hence we use it for all experiments. CluSTAR takes an input a single hyperparameter, the number of centroids initialized, denoted by  $|\mathcal{Z}|$ . We evaluate the following configurations of  $\mathcal{Z}$  and  $c$  for each domain:

1. CartPole: 35 estimators -  $|\mathcal{Z}| \in \{2, 4, 8, 16, 32, 64, 128\}$ ,  $c \in \{1, 2, 3, 4, 5\}$ .
2. ICU-Sepsis: 25 estimators -  $|\mathcal{Z}| \in \{2, 4, 8, 16, 32\}$ ,  $c \in \{1, 2, 3, 4, 5\}$ . ICU-Sepsis with  $|\mathcal{Z}| = 32$  and  $n = 5000$  is excluded for computational reasons.
3. Asterix: 25 estimators -  $|\mathcal{Z}| \in \{2, 4, 8, 16, 32\}$ ,  $c \in \{1, 2, 3, 4, 5\}$ .

In Figure 4 we report the log MSE for each estimator.

**Effect of  $|\mathcal{Z}|$ :** For the class of abstraction functions induced by CluSTAR, the effect of  $|\mathcal{Z}|$  varies across domains. In some domains, such as ICU-Sepsis, the estimators obtained by varying this hyperparameter show similar performance. However, in domains like Asterix and CartPole, performance is more sensitive to the choice of  $|\mathcal{Z}|$ , with larger values performing better.

**Effect of  $c$ :** It is expected that large values of  $c$  will lead to higher variance of the estimates. In low data regimes, small values of  $c$  result in relatively lower MSE as the variance is lowered at the cost of increased bias. As the amount of data increases, the best value of  $c$  also increases. This effect is most pronounced in the Asterix domain, and to a lesser extent in the CartPole domain.

The key takeaway from Figure 4 is that the optimal  $|\mathcal{Z}|$  and  $c$  are a function of the amount of data ( $n$ ), the characteristics of the domain and the class of abstraction functions used.

## C.3 Compute

All experiments were run for 200 seeds each, on 3 domains in total. Each run took between 3 hours to 3 days (depending on the domain) and this duration includes offline data collection. The experiments



were run using 32 threads on Xeon E5-2680 CPUs on a computing cluster, bringing the total compute time to roughly 45000 compute hours.

## NeurIPS Paper Checklist

### 1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: All claims are contributions are supported by corresponding theorems and empirical results.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

### 2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: The limitations are highlighted both in the concluding discussion and as a separate "Limitations" section in the supplementary material.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

### 3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: All proofs are provided in the supplementary material, and assumptions states clearly in the main text.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

#### 4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: The code has been included with the submission. Random seeds are set and fixed across multiple trials, along with the source code of the domain used.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

#### 5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: Anonymized code is submitted as a .zip file with the submission. The codebase will be made public upon acceptance.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

## 6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: There is a specific section, in the main text and supplementary material, dedicated to choice of hyperparameters along with necessary experimental details.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

## 7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: Standard error is reported for main experiments, with the number of trials specified.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

## 8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Details about the compute resources are provided in the supplementary material.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

## 9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification: Naive application of the framework to safety-critical applications could lead to unintended consequences. A discussion on the limitations and potential impacts of the proposed framework is provided in the supplementary material.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

## 10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: Both positive impacts, in terms of strong results for OPE and potential negative impacts, in terms of naive application of the framework to safety-critical applications, are discussed in the paper.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

#### 11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

#### 12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: Original creators of all domains used and all prior work are duly cited in the paper.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.

- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, [paperswithcode.com/datasets](https://paperswithcode.com/datasets) has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

### 13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: The code and related artifacts, such as pre-trained policies, are well documented and provided alongside the submission.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

### 14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

### 15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.



- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.