

---

# Performance Guarantees for Homomorphisms Beyond Markov Decision Processes

---

Sultan Javed Majeed<sup>1</sup>, Marcus Hutter<sup>2</sup>

<sup>1,2</sup> Research School of Computer Science, ANU, Australia \*

<sup>1</sup> <http://www.sultan.pk>, <sup>2</sup> <http://www.hutter1.net>

10 November 2018

## Abstract

Most *real-world* problems have huge state and/or action spaces. Therefore, a naive application of existing *tabular* solution methods is not tractable on such problems. Nonetheless, these solution methods are quite useful if an agent has access to a relatively small state-action space homomorphism of the true environment and near-optimal performance is guaranteed by the map. A plethora of research is focused on the case when the homomorphism is a Markovian representation of the underlying process. However, we show that near-optimal performance is sometimes guaranteed even if the homomorphism is non-Markovian. Moreover, we can aggregate significantly more states by lifting the Markovian requirement without compromising on performance. In this work, we expand Extreme State Aggregation (ESA) framework to joint state-action aggregations. We also lift the policy uniformity condition for aggregation in ESA that allows even coarser modeling of the true environment.

## Contents

<b>1</b>	<b>Introduction</b>	<b>2</b>
<b>2</b>	<b>Preliminaries</b>	<b>4</b>
<b>3</b>	<b>Motivation for Non-MDP Homomorphisms</b>	<b>7</b>
<b>4</b>	<b>Key Elements to Go Beyond MDPs</b>	<b>10</b>
<b>5</b>	<b>Main Results</b>	<b>12</b>
<b>6</b>	<b>Discussion, Outlook and Conclusion</b>	<b>16</b>
<b>A</b>	<b>Omitted Proofs</b>	<b>21</b>
<b>B</b>	<b>Lemmas</b>	<b>26</b>
<b>C</b>	<b>List of Notation</b>	<b>29</b>
<b>D</b>	<b>Example Aggregations</b>	<b>30</b>

## Keywords

homomorphism, state aggregation, non-MDP, action-value aggregation, reinforcement learning.

---

\*A shorter version appeared in the proceedings of the AAAI 2019 conference.

# 1 Introduction

The task of learning a near-optimal behavior from a sequence of experiences can naturally be formulated as a Reinforcement Learning (RL) problem [18]. In a typical RL framework, an agent interacts with an environment by taking an action and receiving a feedback. The agent strives to maximize the expected positive feedback from the environment over time.

It is typically assumed that the agent is facing a small state-action space<sup>1</sup> Markov Decision Process (MDP) so the agent can advise a stationary policy as a function of state [13]. Unfortunately, the number of state-action pairs in most of real-world problems is prohibitively large, e.g. driving a car, playing Go, personal assistance, controlling a plant with real-valued inputs, and so forth. The agent can neither simply visit each state-action pair nor can it keep record of these visits to learn a near-optimal behavior. This explosion of state-action space is known as the *curse of dimensionality* [18]. Therefore, it is essential for the agent to generalize over its experiences in such a huge state-action space problem.

The curse of dimensionality is not a mere artifact of limited experience and computation constraints if the problem has an infinite state-action space. In a typical General Reinforcement Learning<sup>2</sup> (GRL) framework the agent is faced with an environment without any known structure. The GRL setup is arguably the most general setup: it can represent MDP,  $k$ -MDP, partially observed MDP (POMDP) and other typical environment models [6, 8]. But this generality comes at the cost of an infinite state-action space: every agent-environment interaction generates a *unique* history. Hence, there is no other option but to consider every history as a unique state of the environment. Therefore, GRL suffers, inevitably, by the curse of dimensionality.

Solving the curse of dimensionality is an active field of research. A typical solution is to provide (or learn) a finite state-action space model such that the agent can perform near-optimal in the true environment. State aggregation, linear function approximation [12], neural networks — a type of non-linear function approximation — [18] are some well-celebrated solutions to the aforementioned problem. These methods are collectively called generalization, aggregation or dimensionality reduction methods. We are going to refer to them loosely as abstraction/aggregation procedures in this work.

Although most of the abstraction proposals concentrate on the state space reduction [1, 9], there is another equally important dimension of action space that hinders the application of traditional RL methods to real-world problems. The problem with a small state but huge — sometimes continuous — action space is equally challenging for learning and planning, *cf.* continuous bandit problem [3].

A *homomorphism* framework originated by Whitt [20] is a well-studied solution to handle the state-action space curse of dimensionality. In the homomorphism

---

<sup>1</sup>We refer a state and action space/pair jointly as a *state-action* space/pair.

<sup>2</sup>We formally define a typical GRL setup in Section 2.1.

framework a problem of a large state-action space is *solved* by using an *abstract* problem with a relatively small state-action space. The (near-)optimal policy of the abstract problem is a *solution* if it is also a (near-)optimal policy in the true environment.

It is important to highlight that homomorphism is not the only technique for abstracting actions. The *options* framework is a competing method for temporal action abstractions [17]. In the option/macro-action framework, the original action space is augmented with long term/built-in policies [11]. The agent using an option/marco-action commits to execute a fixed set of actions for a fixed (expected) time duration. This temporal action abstraction framework is arguably more powerful but beyond the scope of this work. Because, to the best of our knowledge, there are no theoretical performance guarantees available for such methods, and most probably such bounds might not exist.

In the homomorphism framework, it is typically assumed that the abstract problem is an MDP [14, 15, 19]. However, the size of the abstract state-action space can be significantly reduced if non-MDP abstractions are possible [1, 9]. Moreover, the reduction of abstract state-action space roughly<sup>3</sup> translates into faster learning and planning [16, 7].

The model of the environment can be a crude approximation of the true environment: mapping every history to a single state, a fixed discretization of continuous action space and so on. Therefore, a coarse reduction is not a direct measure of quality of the abstraction: the agent can be arbitrarily worse off in the true environment. Nevertheless, we prefer short maps, i.e. small aggregated state-action space, since the agent can learn faster as compared to a non-aggregation based agent: it has more data available to learn with the same amount of experience. But, it is only possible if the abstraction has a bounded performance loss guarantee, i.e. the agent does not loose much if it considers the aggregated model as the true environment.

It has recently been shown that the MDP restriction is not a necessary condition for near-optimal performance guarantees in state-only abstractions of GRL [6]. In this work, we use similar notation and techniques of Hutter [6] but investigate and prove optimality bounds for non-MDP state-action homomorphisms in GRL. Since state abstraction is a special case of homomorphism (the action space is not reduced/mapped), our work is a generalization of Extreme State Aggregation (ESA) [6].

The homomorphism framework has been extended beyond MDPs to finite-state POMDPs [21]. As mentioned earlier, GRL has an infinite set of histories and no two histories are alike. We can represent a finite-state POMDP environment as a history-based process by imposing a structure that there is an internal MDP that generates the observations and rewards. The GRL framework, by design, is more powerful and expressive than a finite-state POMDP [21, 8]. Therefore, our results are more general than finite-state POMDP homomorphisms.

---

<sup>3</sup>Although reduction of state-action space is necessary for faster learning/planning but not sufficient [10].

**Contributions.** We expand ESA performance guarantees to joint state-action aggregations that scale beyond MDPs. We make another important technical contribution by relaxing the policy uniformity condition in ESA. In ESA, the states are aggregated together if they have the same policy. We show that this requirement can be relaxed and states with approximately similar policies can also be aggregated together with little performance loss. It enables us to have near-optimal maps with, considerably, coarsely aggregated state-action pairs.

The rest of the paper is structured as follows. Section 2 lays the foundations of our homomorphism setup. In Section 3, we motivate the importance of non-MDP homomorphism by an example. Section 4 contains the key elements required to prove the main results. Section 5 contains the main results of this work. In Section 6, we discuss and conclude the paper. We provide all the omitted proofs and lemmas in Appendix A and B, respectively. Appendix C provides a detailed list of notation. Section D provides another example non-MDP homomorphism.

## 2 Preliminaries

This section provides the required notation, a typical GRL framework and our homomorphism setup. We consider a simple agent-environment setup [18]. The agent has a finite set of actions  $\mathcal{A}$ . The environment receives an action from the agent and gives a standard observation from a finite set of observations  $\mathcal{O}$  and a real-valued reward from a finite set  $\mathcal{R} \subseteq \mathbb{R}$ . The agent interacts with the environment in cycles and in each cycle the agent performs an action and receives an observation and reward from the environment. This agent-environment interaction generates a possibly infinite history from an infinite set of histories  $\mathcal{H} := \bigcup_{t=0}^{\infty} (\mathcal{A} \times \mathcal{O} \times \mathcal{R})^t$ . Hence, the original state-action space is the history-action space<sup>4</sup>, i.e.  $\mathcal{H} \times \mathcal{A}$ . Similarly, we define an abstract finite state space  $\mathcal{S}$  and action space  $\mathcal{B}$  to form the abstract state-action space, i.e.  $\mathcal{S} \times \mathcal{B}$ .

We use a consistent notation throughout the paper unless otherwise stated. We use  $\Delta(\cdot)$  to denote a probability distribution over its argument,  $\|\cdot\|_1$  expresses a first norm,  $\tilde{x}$  is a local variable and  $x'$  is a different member of the same set. We use a shorthand notation  $\forall f(x) = y$  to imply  $\forall x, y: f(x) = y$ . A detailed list of notation is provided in Appendix C. We often make references to the results presented later in the paper. The reader is not encouraged to follow these *justifying* references in the first reading.

### 2.1 General Reinforcement Learning Framework

This section provides a formal layout of a typical GRL framework and some assumptions we make about the setup. We start our setup by defining two center

---

<sup>4</sup>In general, histories are considered as the states of the environment, so we interchangeably call the history-action space the *original* state-action space.

pieces of any RL setup: the environment and the agent/policy<sup>5</sup>. The environment, also referred as the *original process*  $P$ , is defined as a stochastic mapping from a history-action pair to a distribution over the observation-reward pairs, i.e.  $P: \mathcal{H} \times \mathcal{A} \rightarrow \Delta(\mathcal{O} \times \mathcal{R})$ . The *history-based* agent/policy  $\Pi$  is defined to be a function that stochastically maps a history to the actions as  $\Pi: \mathcal{H} \rightarrow \Delta(\mathcal{A})$ .

The agent has to maximize the expected sum of rewards it gets from the environment. But this sum can diverge if the agent keeps on adding infinitely many positive rewards. There are, at least, two options to rectify this problem by either assuming a finite sum of rewards or letting the agent discount its future rewards.

**Assumption 1.** (Geometric discounting) *We assume a geometric discounting over the rewards — i.e. the agent discounts its future rewards by a constant discount factor  $\gamma \in [0,1)$ .*

This discounting of rewards serves a dual purpose for the agent, first, it eliminates the problem of infinite sum and second it serves as a parameter for the effective future rewards the agent should care about. A small discount factor makes the agent short sighted and a large discount factor lets the agent be more concerned about future rewards. The goal of the agent is to maximize the expected discounted sum of rewards which is generally expressed with Bellman equations of (action-)value functions [18]. The agent tries to maximize this value function and strives to reach the most valuable states. We define the action-value function  $Q^\Pi$  for any history  $h \in \mathcal{H}$  and action  $a \in \mathcal{A}$  as

$$Q^\Pi(h, a) := \sum_{\tilde{o} \in \mathcal{O}, \tilde{r} \in \mathcal{R}} P(\tilde{o}\tilde{r}|ha) \left( \tilde{r} + \gamma V^\Pi(\tilde{h}) \right) \quad (1)$$

where  $\tilde{h} := ha\tilde{o}\tilde{r}$  is an extended history and the corresponding value function  $V^\Pi$  is defined as

$$V^\Pi(h) := \sum_{\tilde{a} \in \mathcal{A}} Q^\Pi(h, \tilde{a}) \Pi(\tilde{a}|h). \quad (2)$$

The (action-)value functions are maximized if the agent is following an *optimal policy*  $\Pi^* \in \operatorname{argmax}_{\tilde{\Pi}} V^{\tilde{\Pi}}$ . In general the environment can dispense any real-valued reward. But, for a simplified analysis we assume that rewards are bounded and positive.

**Assumption 2.** (Bounded positive reward) *We assume bounded and positive rewards and without loss of generality we assume  $\mathcal{R} : \subseteq [0,1]$ .*

It is easy to see that the bounded rewards bound the (action-)value functions between zero and  $1/(1-\gamma)$ .

---

<sup>5</sup>While it can be argued that an agent and a policy are two separate entities, in this work we use them interchangeably.

## 2.2 Homomorphism Setup

We define a homomorphism as a surjective mapping  $\psi$  from the original state-action space  $\mathcal{H} \times \mathcal{A}$  to the abstract state-action space  $\mathcal{S} \times \mathcal{B}$ .

For a succinct exposition, we also define a few marginalized mapping functions. These marginalized maps do not have any special significance other than making the notation a bit simpler.

**Histories mapped to an  $sb$ -pair.** For a given abstract action  $b \in \mathcal{B}$ , we define a marginalized abstract state map as

$$\psi_b^{-1}(s) := \{h \in \mathcal{H} \mid \exists a \in \mathcal{A} : \psi(h, a) = (s, b)\}. \quad (3)$$

**Actions mapped to an  $sb$ -pair.** Similarly, we also define a marginalized abstract action map for any abstract state  $s \in \mathcal{S}$  and history  $h \in \mathcal{H}$  as

$$\psi_s^{-1}(b) := \{a \in \mathcal{A} \mid \psi(h, a) = (s, b)\}. \quad (4)$$

It is important to note that  $\psi_s^{-1}(b)$  is also a function of history. This dependence is always clear from the context, so we suppress it in the notation.

**Abstract states mapped by a history.** By a slight abuse of notation we overload  $\psi$ , and define a history to abstract state marginalized map as

$$\psi(h) := \{s \in \mathcal{S} \mid \exists a \in \mathcal{A}, b \in \mathcal{B} : \psi(h, a) = (s, b)\}. \quad (5)$$

**Histories mapped to an abstract state.** Finally, an abstract state to history marginalized map is defined as

$$\psi^{-1}(s) := \{h \in \mathcal{H} \mid \exists a \in \mathcal{A}, b \in \mathcal{B} : \psi(h, a) = (s, b)\}. \quad (6)$$

In this work, we assume a structure for the aggregation map  $\psi$ . The general unstructured case is left for future research.

**Assumption 3.** ( $\psi(h)=s$ ) *We assume that an abstract state is determined only by the history — i.e.  $\psi(h,a)=(s=f(h),b)$ , where  $f$  is any fixed surjective function of history and is independent of actions  $a$  and  $b$ .*

The above assumption implies that  $\psi(h)$  is *singleton*. This is not only a technical necessity but a requirement to make the mapping *causal*, i.e. the current history  $h$  corresponds to a unique state  $s$  independent of the next action taken by the agent. If we drop this assumption then the current history might resolve to a different state based on the next (future) action taken by the agent.

A homomorphic map  $\psi$  lets the agent merge the experiences from  $P$  and induces a *history-based* abstract process  $P_\psi$ . Formally, for all  $\psi(h,a)=(s,b)$  and any next abstract state  $s'$ , we express  $P_\psi$  as

$$P_\psi(s'r|ha) := \sum_{\tilde{o}:\psi(ha\tilde{o}r)=s'} P(\tilde{o}r|ha). \quad (7)$$

The map  $\psi$  also induces a *history-based* abstract policy  $\Pi_\psi$  as

$$\Pi_\psi(b|h) := \sum_{\tilde{a} \in \psi_s^{-1}(b)} \Pi(\tilde{a}|h). \quad (8)$$

It is clear from (7) and (8) that the induced abstract process and policy are in general non-Markovian, i.e. both are functions of the history  $h$  and not only the abstract state  $s$ .

**Non-MDP homomorphisms.** In this work we consider two types of non-Markovian homomorphisms: **a)** *Q-uniform* homomorphisms, where the state-action pairs are merged if they have close Q-values, i.e.  $Q^\Pi(h,a) \approx Q^\Pi(h',a')$  for all  $\psi(h,a) = \psi(h',a')$ , and **b)** *V-uniform* homomorphisms, when the merged state-action pairs have close values, i.e.  $V^\Pi(h) \approx V^\Pi(h')$  for all  $\psi(h) = \psi(h')$ . A formal treatment of these non-MDP homomorphisms is provided in the main results section. In both Q and V-uniform homomorphisms,  $P_\psi$  can be history-dependent, in result, the abstract process is non-MDP.

### 3 Motivation for Non-MDP Homomorphisms

In this section we motivate the importance of non-MDP homomorphisms by an example. We show that a non-MDP homomorphism can cater to a large set of domains and allows more compact representations.

**Navigational Grid-world.** Let us consider a simplified version of the asymmetric grid-world example by Ravindran and Barto [15] in Figure 1. In this navigational domain, the goal of an agent  $\Pi$  is to navigate the grid to reach the target cell  $T$ . The unreachable cells are grayed-out. The agent receives a large positive reward if it enters the cell  $T$ , otherwise a small negative reward is given to the agent at each time-step. The agent is capable of moving in the four directions, i.e. up, down, left and right. This domain has an *almost* similar transition and reward structure across a diagonal axis. We call this an approximate MDP axis and denote it by  $\approx$ MDP. This axis of symmetry enables us to create a homomorphism of the domain using approximately half of the original state-space (see Figure 2).

This grid-world example has primarily been studied in the context of either exact, approximate or Bounded parameter MDP (BMDP) homomorphisms [15]: the abstract model *approximately* preserves the one-step dynamics of the original environment. However, as we later prove in this paper (see Theorem 8ii), some non-MDP homomorphisms can also be used to find a near-optimal policy in the original process. We motivate the need of non-MDP homomorphisms, first, by highlighting the fact<sup>6</sup> that in the grid-world domain, the states with similar dynamics have similar optimal action-values. Afterwards, we modify the grid-world domain such that the

---

<sup>6</sup>This section is an informal motivation, we formally deal with this fact in the main results section (Theorem 6i).




Figure 1: The original navigational grid-world with the axis of approximate symmetry. The gray cells are not reachable. The target cell is at the top right corner. The figure shows two possible positions of the agent and corresponding optimal actions.

modified grid-world does not have an approximate MDP symmetry axis, but still has the same approximate optimal action-values symmetry.

We apply Value Iteration (VI) [2] with some fixed but irrelevant parameters on the grid world (see Figure 3). The grid world has the same approximate symmetry axis for the optimal values, denoted by  $\approx Q$ -uniform axis. It is easy to see that each merged state in Figure 2 has the same action-values. Hence, the  $\approx MDP$  axis is also an  $\approx Q$ -uniform axis in the grid-world.

**Modified Navigational Grid-world.** Now we modify the grid world such that it does not have an  $\approx MDP$  axis (Figure 1) but it still has the same  $\approx Q$ -uniform axis (Figure 3). The idea is to take a pair of merged states from Figure 1 and change the reward and transition probabilities such that the states no longer have similar one-step dynamics but still have similar action-values. For example, let us consider the cells highlighted with dashed borders in Figure 3 and denote the cell in the bottom half with  $s_{23}$ . Let  $u, d, p_u$  and  $p_d$  denote the actions up and down, and the probabilities to reach the desired cell by taking the corresponding action, respectively. Let  $r_u$  and  $r_d$  be the expected rewards for each action in the state  $s_{23}$ . In general, we get an *under-determined* set of equations for the action-value function at state  $s_{23}$  as

$$Q^*(s_{23}, a) = \begin{cases} r_u + 0.73\gamma p_u + 3.01\gamma & \text{if } a = u \\ r_d - 0.73\gamma p_d + 3.74\gamma & \text{if } a = d. \end{cases} \quad (9)$$

In the original navigational grid-world problem  $p_u = p_d = 1$ , i.e. each action leads deterministically to the indented reachable cell, and  $r_u = r_d = r_n$ , where  $r_n$  is a fixed






Figure 2: A possible MDP homomorphism by merging the mirror state-action pairs together. The presence of a hashed cell indicates that it is not an exact homomorphism. The agent  $\pi$  solves the problem in this abstract domain.

small negative reward. We can break the  $\approx$ MDP similarity by setting<sup>7</sup>  $p_u = p_d := 0$ , i.e. the actions behave in the opposite way in the lower half,  $r_u := r_n + 0.73\gamma$  and  $r_d := r_n - 0.73\gamma$ , without changing the  $\approx$ Q-uniform similarity. In fact, we can have infinite combinations of rewards and transitions to get a set of modified domains since the set of equations (9) is under-determined.

This set of *modified* domains, by design, no longer allows the approximate MDP homomorphism of Figure 2. Every state is different in terms of reward and transition structure across the  $\approx$ MDP axis of Figure 1. Any one-step model similarity abstraction would be approximately of the same size as the original problem. However, if we consider Q-uniform homomorphisms, i.e. state-action pairs are merged if the action-values are close, then the set of modified domains has a same Q-uniform homomorphism.

In GRL, it is natural to assume that the (expected) rewards are function of realized history. The above modification argument is more likely to hold in a GRL setting: the reward and transition similarity might be hard to satisfy. Therefore, a GRL agent is better to consider such non-MDP homomorphisms to cover more domains with a single abstract model. Now we ask the main question, does such a non-MDP homomorphism, e.g. Q-uniform homomorphism, have a guaranteed solution for the original problem? In the next section, we answer this question in affirmative for Q-uniform homomorphisms (Theorem 8ii), but in negative for V-uniform homomorphisms with a weaker positive result (Theorem 10ii).

---

<sup>7</sup> $p_u = 0$  implies that the action  $u$  now takes the agent to the down cell and vice versa for the action  $d$ .

<b>2.42</b>	<b>2.17</b>		4.16	4.64	5.16
2.70	3.01	3.35	3.74		4.64
2.42	2.70		<b>3.35</b>	3.74	4.16
	2.42	2.17		<b>3.35</b>	
1.94	2.17	2.42	2.70	3.01	<b>2.70</b>
1.74	1.94		2.42	2.70	

Figure 3: The optimal values at each approachable cell. The bold-faced values are not exactly matched across the symmetry axis.

## 4 Key Elements to Go Beyond MDPs

This section introduces the key elements of the paper that enables us to prove performance bounds for non-MDP homomorphisms.

### 4.1 A Stochastic Inverse and Surrogate MDP

The key idea to get a near-optimal policy of the true environment  $P$  is to transform  $P_\psi$  into a surrogate MDP on the abstract state-action space. Afterwards, the optimal policy of this surrogate MDP is uplifted to  $P$ . This technique of casting a non-MDP process as an MDP has been used in ESA [6]. To get this surrogate MDP, we define a stochastic inverse  $B$  of the homomorphism  $\psi$  as a probability measure over the history-action space given an abstract state-action pair, formally,  $B : \mathcal{S} \times \mathcal{B} \rightarrow \Delta(\mathcal{H} \times \mathcal{A})$ . Moreover, we require  $B(ha|sb) := 0$  for any  $\psi(h,a) \neq (s,b)$ . The surrogate MDP is defined as

$$p_B(s'r'|sb) := \sum_{\tilde{h} \in \mathcal{H}, \tilde{a} \in \mathcal{A}} P_\psi(s'r'|\tilde{h}\tilde{a})B(\tilde{h}\tilde{a}|sb). \quad (10)$$

It might seem like a paradoxical idea to solve a non-Markovian  $P_\psi$  using an MDP  $p_B$ , but the paradox is superficial. It is the stochastic inverse that complements the non-Markovianness of  $P_\psi$ . Finding such an inverse algorithmically, hence the surrogate MDP, is not a trivial task in general [6].

This action-dependent stochastic inverse separates our work from the action-independent weighting function considered by Abel et al. [1]. Although, learning of such weighting function is beyond the scope of this paper, an action-independent

weighting function is not learnable. Because, when this weighting function is built from the true sampling distribution, it becomes action-dependent. Hutter [6] constructs such a learnable action-dependent inverse for the state abstraction case. Fortunately, the choice of  $B$  becomes irrelevant in Q-uniform homomorphisms (see Theorem 8), however this is not the case in V-uniform homomorphisms (see Theorem 10).

Similar to the original process, we also define the (action-)value functions for the surrogate MDP  $p_B$  on the abstract state-action space  $\mathcal{S} \times \mathcal{B}$  with an abstract state-based policy  $\pi$ . The action-value function is given as

$$q^\pi(s, b) := \sum_{\tilde{s} \in \mathcal{S}, \tilde{r} \in \mathcal{R}} p_B(\tilde{s}\tilde{r}|sb) (\tilde{r} + \gamma v^\pi(\tilde{s})) \quad (11)$$

where the value function is

$$v^\pi(s) := \sum_{\tilde{b} \in \mathcal{B}} q^\pi(s, \tilde{b}) \pi(\tilde{b}|s). \quad (12)$$

An abstract state-based *optimal* policy  $\pi^*$  is a value maximizer, i.e.  $\pi^* \in \operatorname{argmax}_{\tilde{\pi}} v^{\tilde{\pi}}$ .

## 4.2 Representative Abstract Policy

As discussed earlier, we are primarily interested in the optimal policies of the surrogate MDP. However, it is also interesting to consider a general policy case (e.g. Theorems 4, 5, 7 and 9) akin to an on-policy result where we uplift a representative policy. We use any arbitrary member as a representative policy  $\pi_R$  on the abstract state  $s$ .

$$\pi_R(\cdot|s) := \Pi_\psi(\cdot|\tilde{h}), \quad \text{for some } \tilde{h} \in \psi^{-1}(s). \quad (13)$$

This arbitrary choice of representative introduces a policy representation error  $\varepsilon_\Pi$  for each abstract state  $s$ , expressed as

$$\varepsilon_\Pi(s) := \sup_{\tilde{h} \in \psi^{-1}(s)} \left\| \pi_R(\cdot|s) - \Pi_\psi(\cdot|\tilde{h}) \right\|_1. \quad (14)$$

This representation error is small/zero when the induced abstract policy  $\Pi_\psi$  is approximately/piecewise constant, i.e.  $\Pi_\psi(\cdot|h) = \Pi_\psi(\cdot|h')$  for all  $\psi(h) = \psi(h')$ .

It is easy to see that the state aggregation mapping function  $\phi$  of ESA setup<sup>8</sup> is a special case of our generalized mapping function  $\psi$ . We can mimic any  $\phi$  by using a  $\psi_\phi$  with identity transformation over action spaces — i.e.  $\psi_\phi(h, a) := (\phi(h), a)$ .

In the next section, we provide the main results of this work. We construct a near-optimal policy for the original process from the surrogate MDP even if the homomorphism is non-MDP.

---

<sup>8</sup>The reader is encouraged to see [6] for more details about  $\phi$ .

## 5 Main Results

In this section, we prove that under some conditions an optimal policy of the abstract process losses only a fraction of the value when uplifted in the original process. Also, these results hold even when the marginalized process is not an MDP. We analyze three types of homomorphisms in this work: MDP, Q-Uniform and V-Uniform homomorphisms. Both Q and V-Uniform homomorphisms are non-Markovian by definition. In general, MDP and Q-Uniform homomorphisms admit a *deterministic* near-optimal policy of the original process, while V-Uniform homomorphisms do not.

### 5.1 Markov Decision Process Homomorphisms

A homomorphism is an MDP homomorphism if the induced abstract process  $P_\psi$  is an MDP. Then, there exists a process  $p_{\text{MDP}}$  such that for all  $\psi(h,a) = (s,b)$  and for all  $\tilde{s}$  and  $\tilde{r}$ , it holds:

$$P_\psi(\tilde{s}\tilde{r}|ha) = p_{\text{MDP}}(\tilde{s}\tilde{r}|sb). \quad (15)$$

Using the above condition in (10), renders  $p_B = p_{\text{MDP}}$  and independent of  $B$ . The condition (15) is a *stronger* version of the bisimulation condition [4] that is generalized to joint history-action pairs. This condition is strong enough to preserve the optimal (action-)value functions of the original process (see Theorem 6). But, it is not strong enough to preserve arbitrary policy (action-)value functions (see Theorem 4). Unless we define a notion of action-value function representative and a corresponding representation error. For an abstract state-action pair, the representative action-value is defined as

$$Q^\Pi(\psi^{-1}(s,b)) := Q^\Pi(\tilde{h}, \tilde{a}), \text{ for some } \psi(\tilde{h}, \tilde{a}) = (s,b) \quad (16)$$

and the representation error of the action-value function is expressed as

$$\varepsilon_Q(s) := \sup_{\tilde{h}, \tilde{a}, \tilde{b}: \psi(\tilde{h}, \tilde{a}) = (s, \tilde{b})} \left| Q^\Pi(\psi^{-1}(s, \tilde{b})) - Q^\Pi(\tilde{h}, \tilde{a}) \right|. \quad (17)$$

Similar to  $\varepsilon_\Pi$ , this representation error is small/zero if the action-value function is approximately/piecewise constant. At this point, we have all the required components properly defined to state the first theorem of the paper.

**Theorem 4.** ( $\psi_{\text{MDP}\Pi}$ ) *Let  $\psi$  be a homomorphism such that  $P_\psi$  is an MDP, then for any policy  $\Pi$  and all  $\psi(h,a) = (s,b)$  it holds:*

$$\begin{aligned} |q^{\pi_R}(s,b) - Q^\Pi(h,a)| &\leq \frac{\gamma \varepsilon_{\max}}{1-\gamma} \quad \text{and} \\ |v^{\pi_R}(s) - V^\Pi(h)| &\leq \frac{\varepsilon_{\max}}{1-\gamma} \end{aligned}$$

where  $\varepsilon_{\max} := \max_{\tilde{s} \in \mathcal{S}} \left( \varepsilon_Q(\tilde{s}) + \frac{\varepsilon_\Pi(\tilde{s})}{1-\gamma} \right)$ .

*Proof.* See Appendix A. ■

The above theorem shows that the surrogate MDP *approximately* preserves the (action-)value functions of the original process for any arbitrary policy. However, these (action-)value functions are preserved *exactly* if we further impose a policy uniformity condition in addition to an MDP assumption.

**Theorem 5.** ( $\psi_{\text{MDP}\Pi=}$ ) *Let  $\psi$  be a homomorphism such that  $P_\psi$  is an MDP and  $\Pi(\cdot|h) = \Pi(\cdot|h')$  (i.e. the policy similarity condition holds) for some policy  $\Pi$  and for all  $\psi(h) = \psi(h')$ . Then for all  $\psi(h,a) = (s,b)$  it holds:*

$$q^{\pi_R}(s, b) = Q^\Pi(h, a) \quad \text{and} \quad v^{\pi_R}(s) = V^\Pi(h).$$

*Proof.* See Appendix A. ■

Theorems 4 and 5 are important but not very useful results. As already discussed, we are interested in the (near-)optimal policies of the original process. And, we want to find the abstract policies that can be lifted with a performance guarantee from the abstract state-action space to the original history-action space.

**Theorem 6.** ( $\psi_{\text{MDP}^*}$ ) *Let  $\psi$  be a homomorphism such that  $P_\psi$  is an MDP, then for all  $\psi(h,a) = (s,b)$  it holds:*

(i)  $q^*(s,b) = Q^*(h,a)$  and  $v^*(s) = V^*(h)$ .

(ii)  $V^*(h) = V^{\check{\Pi}}(h)$

where  $\check{\Pi}(h) := \psi_s^{-1}(\pi^*(s))$  for any  $\psi(h) = s$ .

*Proof.* See Appendix A. ■

For any MDP homomorphism, the performance guarantee is provided by Theorem 6ii. The abstract optimal policy  $\pi^*$  is also an optimal policy for the original process when lifted to the original history-action space.

## 5.2 Q-Uniform Homomorphisms

In this section, we relax the MDP condition (see Equation 15) on the abstract-process provided by the homomorphism. We show that there still exists an abstract policy that is near-optimal in the original process (see Theorem 8ii). We start with proving a value loss bound for an arbitrary policy when the action-value function of the original process is *approximately*  $\psi$ -uniform.

**Theorem 7.** ( $\psi_{Q^\Pi}$ ) Assume  $|Q^\Pi(h,a) - Q^\Pi(h',a')| \leq \varepsilon$  for some policy  $\Pi$  and for all  $\psi(h,a) = \psi(h',a')$ . Then for all  $\psi(h,a) = (s,b)$  it holds:

$$\begin{aligned} |Q^\Pi(h,a) - q^{\pi_R}(s,b)| &\leq \varepsilon + \frac{\gamma\varepsilon(s)}{1-\gamma} \quad \text{and} \\ |V^\Pi(h) - v^{\pi_R}(s)| &\leq \frac{\varepsilon(s)}{1-\gamma} \end{aligned}$$

where  $\varepsilon(s) := 2\varepsilon + \frac{\varepsilon_\Pi(s)}{1-\gamma}$ .

*Proof.* See Appendix A. ■

The following theorem *improves* the optimal policy value loss bounds, cf. Theorem 7, and establishes the existence of a near-optimal policy of the original history-action space in the abstract state-action space.

**Theorem 8.** ( $\psi_{Q^*}$ ) Let  $|Q^*(h,a) - Q^*(h',a')| \leq \varepsilon$  for all  $\psi(h,a) = \psi(h',a')$ , then for all  $\psi(h,a) = (s,b)$  it holds:

$$(i) \quad |Q^*(h,a) - q^*(s,b)| \leq \frac{2\varepsilon}{1-\gamma} \quad \text{and} \\ |V^*(h) - v^*(s)| \leq \frac{2\varepsilon}{1-\gamma}.$$

$$(ii) \quad 0 \leq V^*(h) - V^{\check{\Pi}}(h) \leq \frac{4\varepsilon}{(1-\gamma)^2}$$

where  $\check{\Pi}(h) := \psi_s^{-1}(\pi^*(s))$  for any  $\psi(h) = s$ .

*Proof.* See Appendix A. ■

It is important to note that Theorem 8 holds for any stochastic inverse  $B$ . Every choice of  $B$  gives a different surrogate MDP  $p_B$ , so the theorem provides a *near-optimal* performance guarantee for the *uplifted* abstract optimal policies of *any* possible surrogate MDP. Therefore, for any non-MDP Q-uniform homomorphism and a fixed  $B$  there exists an uplifted near-optimal policy ( $\check{\Pi}$  from Theorem 8ii).

### 5.3 V-Uniform Homomorphisms

All the previous results are valid for any choice of the stochastic inverse  $B$ . However, for V-uniform homomorphisms, the results are explicitly dependent on  $B$  (see Theorem 9 and 10). We need a couple of more entities to express the results of this section. We denote the  $B$ -average of the action-value function of the original process as

$$\langle Q^\Pi(\psi^{-1}(s,b)) \rangle_B := \sum_{\tilde{h} \in \mathcal{H}, \tilde{a} \in \mathcal{A}} Q^\Pi(\tilde{h}, \tilde{a}) B(\tilde{h}\tilde{a}|sb). \quad (18)$$

Furthermore, we can decompose  $B$  into two distinct parts: action dependent and independent. With an abuse of notation, assume an arbitrary joint distribution  $B$  over  $\mathcal{H}, \mathcal{A}, \mathcal{S}$  and  $\mathcal{B}$ . By using the chain rule of probability distributions on  $B$ ,

$$\begin{aligned}
B(ha|sb) &= B(h|sb)B(a|bhs) \\
&= \frac{B(hs)B(b|hs)}{B(sb)}B(a|bhs) \\
&\stackrel{(a)}{=} \frac{B(hs)B(b|h)}{B(sb)}B(a|bh) \\
&= B(h|s)\frac{B(b|h)}{B(b|s)}B(a|bh) \\
&= \underbrace{B(h|s)}_{\text{action-independent}} \cdot \overbrace{\left(\frac{B(ab|h)}{B(b|s)}\right)}^{\text{action-dependent}}
\end{aligned} \tag{19}$$

(a) follows from Assumption 3, the state is determined only by the history.

Using the action-dependent part from (19), we define a history and state based induced measure on the original action space for any  $B$  and an abstract state based policy  $\pi$  as

$$B^\pi(a|hs) := \sum_{\tilde{b} \in \mathcal{B}} \left( \frac{B(a\tilde{b}|h)}{B(\tilde{b}|s)} \right) \pi(\tilde{b}|s). \tag{20}$$

This seemingly complex and arbitrary relationship has a well-structured explanation. If *approximately*, the  $B$  distribution is linked to the actual dynamics of an agent  $\pi$  acting in the abstract state-action space, i.e.  $B(b|s) \approx \pi(b|s)$ , then  $B^\pi(a|hs) \approx B(a|h)$ , which is effectively a *shadow* agent induced by the agent  $\pi$  on the original history-action space.

To prove a result analogous to Theorem 7 for a  $V$ -uniform homomorphism, we need to impose an extra condition on  $B$ , cf. Theorem 7, which requires a structure on  $B$  and/or on the underlying original process. For general  $B$ , there exist some known counter examples [6].

**Theorem 9.** ( $\psi_{V^\Pi}$ ) Let  $|V^\Pi(h) - V^\Pi(h')| \leq \varepsilon$  for some policy  $\Pi$  and for all  $\psi(h) = \psi(h')$ , and  $|\sum_{\tilde{a} \in \mathcal{A}} Q^\Pi(h, \tilde{a}) B^{\pi_R}(\tilde{a}|hs) - V^\Pi(h)| \leq \varepsilon_B$  for all  $s = \psi(h)$ , then it holds:

$$\begin{aligned}
|\langle Q^\Pi(\psi^{-1}(s, b)) \rangle_B - q^{\pi_R}(s, b)| &\leq \frac{\gamma(\varepsilon + \varepsilon_B)}{1 - \gamma} \quad \text{and} \\
|V^\Pi(h) - v^{\pi_R}(s)| &\leq \frac{\varepsilon + \varepsilon_B}{1 - \gamma}.
\end{aligned}$$

*Proof.* See Appendix A. ■

In Theorem 7, we had an absolute loss-bound for action-value functions but in Theorem 9 we only have a  $B$ -average relationship. So far, we were able to get a near-optimal performance guarantee when the optimal policy of a surrogate MDP is uplifted to the original process (see Theorems 6ii and 8ii). However, there does not exist such a near-optimal performance guarantee for  $V$ -uniform homomorphisms. A *deterministic* abstract policy could be arbitrarily worse off when uplifted to the original process [6, Theorem 10] in  $V$ -uniform state-abstractions, which are a special case of  $V$ -uniform homomorphisms. However, a relatively weak result is still possible.

**Theorem 10.** ( $\psi_{V^*}$ ) *Let  $|V^*(h) - V^*(h')| \leq \varepsilon$  for all  $\psi(h) = \psi(h')$  and  $|\sum_{\tilde{a} \in \mathcal{A}} Q^*(h, \tilde{a}) B^{\pi^*}(\tilde{a}|hs) - V^*(h)| \leq \varepsilon_B$  for all  $s = \psi(h)$ , then for all  $\psi(h, a) = (s, b)$  it holds:*

$$(i) \quad | \langle Q^*(\psi^{-1}(s, b)) \rangle_B - q^*(s, b) | \leq \frac{3\gamma(\varepsilon + \varepsilon_B)}{(1-\gamma)^2} \text{ and}$$

$$|V^*(h) - v^*(s)| \leq \frac{3(\varepsilon + \varepsilon_B)}{(1-\gamma)^2}.$$

(ii) *If  $\varepsilon + \varepsilon_B = 0$  then  $\psi(h, \Pi^*(h)) = (s, \pi^*(s))$  for all  $\psi(h) = s$ .*

*Proof.* See Appendix A. ■

In the *approximate* case, i.e.  $\varepsilon + \varepsilon_B > 0$ , Theorem 10 is not as useful as Theorem 8 because of the missing performance guarantee, *cf.* Theorem 8ii. However, it is still an important theorem for the *exact*  $V$ -uniform homomorphisms, i.e.  $\varepsilon + \varepsilon_B = 0$ . In that case, it is guaranteed that the optimal actions of *all* member histories are mapped to the *same* abstract optimal action (see Theorem 10ii).

## 6 Discussion, Outlook and Conclusion

In this paper we analyzed *approximate* homomorphisms of a general history-based environment. The main idea was to find a *deterministic* policy in the abstract state-action space such that, when uplifted, it is a near-optimal policy in the original problem. Using the surrogate MDP technique, we proved near-optimal performance bounds for both MDP (Theorem 6ii) and  $Q$ -uniform homomorphisms (Theorem 8ii). In general, there does not exist a near-optimal *deterministic* uplifted policy for  $V$ -uniform homomorphisms. However, we proved a weaker result (Theorem 10ii) for the *exact*  $V$ -uniform homomorphisms: the optimal actions of the member histories are mapped to the same abstract optimal action at the corresponding state of the surrogate MDP.

**Versus ESA.** We borrow some notation and techniques from Hutter [6]. But this work is crucially different from ESA. Apart from the obvious difference of being a generalization to homomorphisms, there are also some other key differences. In ESA, the policy  $\Pi$  is required to be state uniform for various of the main results [6,



Theorems 1,5,6 and 9], whereas we do not make any such assumption. Moreover, at the first instance our results might look *almost* similar to ESA but the important difference is in the definition of  $\varepsilon(s)$  which is *not* a simple addition of both state and action representation errors. It is a non-trivial weighted average of representation errors. The extra conditions on Theorems 9 and 10 are *weaker* than the policy-uniformity condition, *cf.* [6, Theorems 6 and 9], and do not have direct counterparts in ESA.

**Versus Options.** As briefly addressed in the introduction section, the options framework does not have any provable performance guarantees, yet. Whereas our restriction of uplifting a state-based policy and using a deceptively “spatial-looking” abstraction of actions have such guarantees. Since we allow the action mapping part of  $\psi$  to be a function of history, which is arguably a function of time, our framework also admits temporal dependencies. It enables  $\psi$  to model much more than mere renaming of the original action space distributions. A thorough comparison between these two approaches is left for future work.

## 6.1 Outlook

The results in this work are quite general but there are various open questions left for future research.

**Reinforcement Learning (Learning Problem).** For a given homomorphism  $\psi$ , the most obvious question we left open is the choice of  $B$ . We call this the *learning problem*. Two of the three main results in this work (Theorems 6ii and 8ii) are valid for any choice of  $B$ , so any fixed  $B$  would suffice. But the third main result (Theorem 10ii) is very much involved with the choice of  $B$ . However, it is not a strong result in itself. Nevertheless, in a state-abstraction context,  $B$  facilitates learning of the surrogate MDP from the induced abstract process [6]. Therefore, it is an intriguing direction to explore for homomorphisms.

**Feature Reinforcement Learning (Discovery Problem).** The focus of this paper is to provide performance guarantees for a given homomorphism. While in practice, the agent has to learn such a reduction/model from experience. It is known as the *discovery problem* [9] in RL and *Feature Reinforcement Learning* (FRL) [5] in the GRL context. It is non-trivial to solve this problem even in a state-abstraction framework [6]. Our result can help to build such an FRL algorithm for homomorphisms, e.g. during the model learning/building, the algorithm may use the bounds from this work to select/discard a candidate model.

**Special Environment Classes.** In general, we do not use/exploit structure of the underlying original process. However, effects of a specific model class can be expressed in terms of the (action-)value functions. For example, if the original process is a finite state POMDP then our results provide the performance-loss guarantee by representing a belief-state based value function of the POMDP by a state-based

value function. A similar argument can be rendered for various other types of model classes. Since the results in this work are general, they are not expected to gracefully scale down to some class specific tight performance bounds. Nevertheless, it is an important agenda to get the scaled-down variants of these results for some specific model classes.

**Continuous state-action space.** The results in this paper easily extend to the continuous state-action space homomorphisms for the *measurable* maps. The summations change to integrals and the measurability constraint make sure that these integrals are well-defined. In this case, a homomorphism map has a natural interpretation of being a discretization of the underlying space. However, it is sometimes desirable to use a restricted continuity condition, e.g. Lipschitz or Holder continuity, rather than the weak measurability constraint. A continuous state-action homomorphism under some restricted continuity constraints would be a nice generalization of our results.

**Fully Generalized Homomorphism.** In a sense our results are not *fully* general since we assumed a structure on the homomorphism. A fully generalized homomorphism formulation with no  $\psi(h,a) = (f(h),b)$  assumption would be an interesting extension of this work. However, lifting this condition may lead to some bizarre non-causal effects, e.g. the *current* abstract state would be decided by the *next* action!

## 6.2 Conclusion

In conclusion, our results show that in GRL the near-optimal performance guarantee is not limited only to MDP homomorphisms. It is sometimes possible to have non-MDP models, i.e. Q-uniform homomorphisms, with bounded performance loss guarantees. We also relax the *strong* policy-uniformity condition in ESA to allow stochastic policies.

**Acknowledgements.** This work has in parts been supported by Australian Research Council grant DP150104590.

## References

- [1] David Abel, D. Ellis Hershkowitz, and Michael L. Littman. Near Optimal Behavior via Approximate State Abstraction. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning*, pages 2915–2923, jan 2016. ISBN 9781510829008.
- [2] Richard Bellman. A Markovian Decision Process. *Indiana University Mathematics Journal*, 6(4):679–684, 1957. ISSN 01650114. doi: 10.1007/BF02935461.

- [3] Sébastien Bubeck and Nicolò Cesa-Bianchi. Regret Analysis of Stochastic and Non-stochastic Multi-armed Bandit Problems. *Foundations and Trends® in*, 2012. ISSN 1935-8237. doi: 10.1561/22000000024.
- [4] Robert Givan, Thomas Dean, and Matthew Greig. Equivalence Notions and Model Minimization in Markov Decision Processes. *Artificial Intelligence*, 147(1-2):163–223, 2003. ISSN 00043702. doi: 10.1016/S0004-3702(02)00376-4.
- [5] Marcus Hutter. Feature Reinforcement Learning: Part I. Unstructured MDPs. *Journal of Artificial General Intelligence*, 1(1):3–24, jan 2009. ISSN 1946-0163. doi: 10.2478/v10229-011-0002-8.
- [6] Marcus Hutter. Extreme State Aggregation Beyond Markov Decision Processes. *Theoretical Computer Science*, 650:73–91, 2016. ISSN 03043975. doi: 10.1016/j.tcs.2016.07.032.
- [7] Tor Lattimore and Marcus Hutter. Near-optimal PAC bounds for discounted MDPs. *Theoretical Computer Science*, 558(C):125–143, 2014. ISSN 03043975. doi: 10.1016/j.tcs.2014.09.029.
- [8] Jan Leike. *Nonparametric General Reinforcement Learning*. PhD thesis, Australian National University, nov 2016.
- [9] Lihong Li, Thomas J Walsh, and Michael L Littman. Towards a Unified Theory of State Abstraction for MDPs. In *Proceedings of the Ninth International Symposium on Artificial Intelligence and Mathematics*, pages 531–539, 2006.
- [10] Michael L. Littman, Thomas L. Dean, and Leslie Pack Kaelbling. On the Complexity of Solving Markov Decision Problems. In *Proceedings of the Eleventh conference on Uncertainty in artificial intelligence*, pages 394–402, feb 1995. ISBN 1-55860-385-9. doi: 10.1007/11871842.
- [11] Amy McGovern, Richard S. Sutton, and Andrew H. Fagg. Roles of Macro-Actions in Accelerating Reinforcement Learning. In *incompleteideas.net*, pages 13–18, 1997. ISBN 9788578110796. doi: 10.1017/CBO9781107415324.004.
- [12] R Parr, L Li, G Taylor, C Painter-Wakefield, and M L Littman. An analysis of linear models, linear value-function approximation, and feature selection for reinforcement learning. In *Icml 2008*, pages 752–759, 2008. ISBN 9781605582054. doi: 10.1145/1390156.1390251.
- [13] Ml Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, 2014. ISBN 9780470316887. doi: 10.1002/9780470316887.
- [14] B Ravindran and a G Barto. Relativized Options: Choosing the Right Transformation. *Proceedings of the Twentieth International Conference on Machine Learning*, pages 608–615, 2003.

- [15] B Ravindran and A G Barto. Approximate Homomorphisms: A Framework for Non-exact Minimization in Markov Decision Processes. *Proceedings of the Fifth International Conference on Knowledge Based Computer Systems (KBCS 04)*, pages 19–22, 2004.
- [16] Alexander L Strehl, Lihong Li, and Michael L Littman. Reinforcement Learning in Finite MDPs : PAC Analysis. *Journal of Machine Learning Research*, 10:2413–2444, 2009. ISSN 1532-4435. doi: 10.1.1.153.9841.
- [17] R S Sutton, D Precup, and S Singh. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence*, 112(1): 181–211, 1999.
- [18] R.S. Sutton and A.G. Barto. *Reinforcement Learning: An Introduction*, volume 135. MIT press Cambridge, 1st edition, 1998. ISBN 0262193981. doi: 10.1109/TNN.1998.712192.
- [19] Jonathan Taylor, Doina Precup, and Prakash Panangaden. Bounding Performance Loss in Approximate MDP Homomorphisms. *Advances in Neural Information Processing Systems (NIPS) 21*, pages 1649–1656, 2008.
- [20] Ward Whitt. Approximations of Dynamic Programs, I. *Mathematics of Operations Research*, 3(3):231–243, 1978. ISSN 0364-765X. doi: 10.1287/moor.4.2.179.
- [21] A.P. Wolfe. Paying Attention To What Matters : Observation Abstraction In Partially Observable Environments. *Open Access Dissertations*, (February 2010):188, 2010.

## A Omitted Proofs

We provide all omitted proofs in this appendix. We use  $\leq$  to denote a two side inequality, e.g. if  $|x-y| \leq \varepsilon$ , which in effect implies  $-\varepsilon \leq x-y \leq \varepsilon$ , we denote it as  $x \leq y \pm \varepsilon$  to express both inequalities at the same time.

### A.1 Proof of Theorem 4

*Proof.* Let  $\delta := \sup_{\tilde{h}, \tilde{a}, \tilde{s}, \tilde{b}: \psi(\tilde{h}, \tilde{a}) = (\tilde{s}, \tilde{b})} |q^\pi(\tilde{s}, \tilde{b}) - Q^\Pi(\tilde{h}, \tilde{a})|$ , and for any  $\psi(h) = s$  we have,

$$\begin{aligned}
 |v^\pi(s) - V^\Pi(h)| &\stackrel{(a)}{\leq} \left| \sum_{\tilde{b} \in \mathcal{B}} q^\pi(s, \tilde{b}) \pi(\tilde{b}|s) - \sum_{\tilde{b} \in \mathcal{B}} Q^\Pi(\psi^{-1}(s, \tilde{b})) \pi(\tilde{b}|s) \right| + \varepsilon_Q(s) + \frac{\varepsilon_\Pi(s)}{1-\gamma} \\
 &\leq \sum_{\tilde{b} \in \mathcal{B}} |q^\pi(s, \tilde{b}) - Q^\Pi(\psi^{-1}(s, \tilde{b}))| \pi(\tilde{b}|s) + \varepsilon_Q(s) + \frac{\varepsilon_\Pi(s)}{1-\gamma} \\
 &\leq \delta + \varepsilon_Q(s) + \frac{\varepsilon_\Pi(s)}{1-\gamma} \tag{21}
 \end{aligned}$$

(a) follows from the definition of  $v^\pi(s)$  and Lemma 11. Now for any  $\psi(h, a) = (s, b)$ , we have,

$$\begin{aligned}
 Q^\Pi(h, a) &\equiv \sum_{\tilde{o} \in \mathcal{O}, \tilde{r} \in \mathcal{R}} P(\tilde{o}\tilde{r}|ha) (\tilde{r} + \gamma V^\Pi(\tilde{h})) \quad [\tilde{h} = ha\tilde{o}\tilde{r}] \\
 &\stackrel{(b)}{\leq} \sum_{\tilde{s} \in \mathcal{S}, \tilde{r} \in \mathcal{R}} P_\psi(\tilde{s}\tilde{r}|ha) (\tilde{r} + \gamma v^\pi(\tilde{s})) \pm \gamma \left( \delta + \varepsilon_Q(\tilde{s}) + \frac{\varepsilon_\Pi(\tilde{s})}{1-\gamma} \right) \\
 &\stackrel{(15)}{=} \sum_{\tilde{s} \in \mathcal{S}, \tilde{r} \in \mathcal{R}} p(\tilde{s}\tilde{r}|sb) (\tilde{r} + \gamma v^\pi(\tilde{s})) \pm \gamma \left( \delta + \varepsilon_Q(\tilde{s}) + \frac{\varepsilon_\Pi(\tilde{s})}{1-\gamma} \right) \\
 &\leq q^\pi(s, b) \pm \gamma(\delta + \varepsilon_{\max})
 \end{aligned}$$

(b) follows from value function error bound (21) and definition of  $P_\psi$  given by (7). Since,  $\delta \equiv \sup |q^\pi(s, b) - Q^\Pi(h, a)| \leq \gamma(\delta + \varepsilon_{\max})$  therefore,  $\delta \leq \frac{\gamma \varepsilon_{\max}}{1-\gamma}$ , hence completes the proof.  $\blacksquare$

### A.2 Proof of Theorem 6

*Proof.* (i) Let  $\delta := \sup_{h, a, s, b: \psi(h, a) = (s, b)} |q^*(s, b) - Q^*(h, a)|$ . Now for any  $\psi(h) = s$  we have,

$$\begin{aligned}
 |v^*(s) - V^*(h)| &\stackrel{(a)}{\leq} \left| \max_{\tilde{b} \in \mathcal{B}} q^*(s, \tilde{b}) - \max_{\tilde{b} \in \mathcal{B}} Q^*(\psi^{-1}(s, \tilde{b})) \right| + \varepsilon_Q(s) \\
 &\stackrel{(b)}{\leq} \delta \tag{22}
 \end{aligned}$$

(a) follows from the definitions of  $v^*(s)$  and Lemma 12, and (b) is due to Lemma 13.

$$\begin{aligned}
Q^*(h, a) &\equiv \sum_{\tilde{o} \in \mathcal{O}, \tilde{r} \in \mathcal{R}} P(\tilde{o}\tilde{r}|ha)(\tilde{r} + \gamma V^*(\tilde{h})) \quad [\tilde{h} = ha\tilde{o}\tilde{r}] \\
&\stackrel{(22)}{\leq} \sum_{\tilde{s} \in \mathcal{S}, \tilde{r} \in \mathcal{R}} P_\psi(\tilde{s}\tilde{r}|ha)(\tilde{r} + \gamma v^*(\tilde{s})) \pm \gamma\delta \\
&\stackrel{(15)}{=} \sum_{\tilde{s} \in \mathcal{S}, \tilde{r} \in \mathcal{R}} p(\tilde{s}\tilde{r}|sb)(\tilde{r} + \gamma v^*(\tilde{s})) \pm \gamma\delta \\
&\equiv q^*(s, b) \pm \gamma\delta
\end{aligned}$$

Therefore,  $\delta \leq \gamma\delta$ , therefore  $\delta = 0$  which completes the proof.

(ii) For  $\psi(h) = s$  and  $\check{\Pi}(h) \in \psi_s^{-1}(\pi^*(s))$ ,

$$V^*(h) \stackrel{(i)}{=} v^*(s) \equiv q^*(s, \pi^*(s)) \stackrel{(i)}{=} Q^*(h, \check{\Pi}(h))$$

which implies  $Q^*(h, \check{\Pi}(h)) = V^*(h)$  and Lemma 16 concludes the proof. ■

### A.3 Proof of Theorem 5

*Proof.* Let  $\delta := \sup_{h, h': \psi(h) = \psi(h'), a, a'} |Q^\Pi(h, a) - Q^\Pi(h', a')|$ , then for all  $\psi(h) = \psi(h')$ ,

$$\begin{aligned}
|V^\Pi(h) - V^\Pi(h')| &\equiv \left| \sum_{\tilde{a} \in \mathcal{A}} Q^\Pi(h, \tilde{a}) \Pi(\tilde{a}|h) - \sum_{\tilde{a} \in \mathcal{A}} Q^\Pi(h', \tilde{a}) \Pi(\tilde{a}|h') \right| \\
&\stackrel{(a)}{=} \left| \sum_{\tilde{a} \in \mathcal{A}} (Q^\Pi(h, \tilde{a}) - Q^\Pi(h', \tilde{a})) \Pi(\tilde{a}|h) \right| \leq \delta \tag{23}
\end{aligned}$$

(a) follows from the assumption. Now for all  $\psi(h, a) = \psi(h', a') = (s, b)$ ,

$$\begin{aligned}
|Q^\Pi(h, a) - Q^\Pi(h', a')| &\stackrel{(b)}{=} \left| \sum_{\tilde{s} \in \mathcal{S}, \tilde{r} \in \mathcal{R}} P_\psi(\tilde{s}\tilde{r}|ha) (\tilde{r} + \gamma V^\Pi(\tilde{h})) \right. \\
&\quad \left. - \sum_{\tilde{s} \in \mathcal{S}, \tilde{r} \in \mathcal{R}} P_\psi(\tilde{s}\tilde{r}|h'a') (\tilde{r} + \gamma V^\Pi(\tilde{h}')) \right| \\
&\stackrel{(15)}{=} \gamma \left| \sum_{\tilde{s} \in \mathcal{S}, \tilde{r} \in \mathcal{R}} p(\tilde{s}\tilde{r}|sb) (V^\Pi(\tilde{h}) - V^\Pi(\tilde{h}')) \right| \\
&\stackrel{(23)}{\leq} \gamma\delta \tag{24}
\end{aligned}$$

(b) follows from the definition and  $\tilde{h} = ha\tilde{o}\tilde{r}$  and  $\tilde{h}' = h'a'\tilde{o}\tilde{r}$ , and  $\psi(\tilde{h}) = \psi(\tilde{h}') = \tilde{s}$ . From the inequality (24), we have  $\delta \leq \gamma\delta \Rightarrow \delta = 0$ . Therefore,  $Q^\Pi(h, a) = Q^\Pi(h', a')$  for

all  $\psi(h,a) = \psi(h',a')$ . Note that this also implies,  $\varepsilon_Q = 0$  and  $\varepsilon_\Pi = 0$  by assumption.  $\blacksquare$

## A.4 Proof of Theorem 7

*Proof.* Let  $\delta := \sup_{\tilde{h}, \tilde{a}, \tilde{s}, \tilde{b}: \psi(\tilde{h}, \tilde{a}) = (\tilde{s}, \tilde{b})} \left| Q^\Pi(\tilde{h}, \tilde{a}) - q^\pi(\tilde{s}, \tilde{b}) \right|$ , and for any  $\psi(h) = s$  we have,

$$\begin{aligned}
V^\Pi(h) - v^\pi(s) &\stackrel{Lem.11}{\leq} \sum_{\tilde{b} \in \mathcal{B}} \left( Q^\Pi(\psi^{-1}(s, \tilde{b})) \pi(\tilde{b}|s) - q^\pi(s, \tilde{b}) \pi(\tilde{b}|s) \right) \pm \left( \varepsilon_Q(s) + \frac{\varepsilon_\Pi(s)}{1-\gamma} \right) \\
&= \sum_{\tilde{b} \in \mathcal{B}} \left( Q^\Pi(\psi^{-1}(s, \tilde{b})) - q^\pi(s, \tilde{b}) \right) \pi(\tilde{b}|s) \pm \left( \varepsilon_Q(s) + \frac{\varepsilon_\Pi(s)}{1-\gamma} \right) \\
&\stackrel{(a)}{\leq} \pm \left( \delta + \varepsilon_Q(s) + \frac{\varepsilon_\Pi(s)}{1-\gamma} \right) \tag{25}
\end{aligned}$$

(a) follows from the definition of  $\delta$  and the fact that  $\pi(\tilde{b}|s)$ -average is smaller than the  $\tilde{b}$ -supremum. Using the above inequality (25) and Lemma 14 we get,

$$\left| \langle Q^\Pi(\psi^{-1}(s, b)) \rangle_B - q^\pi(s, b) \right| \leq \gamma \left( \delta + \varepsilon_Q(s) + \frac{\varepsilon_\Pi(s)}{1-\gamma} \right). \tag{26}$$

We exploit the theorem's assumption and derive a key relationship between the  $B$  average and any instance of action value.

$$\begin{aligned}
\langle Q^\Pi(\psi^{-1}(s, b)) \rangle_B &\equiv \sum_{\tilde{h} \in \mathcal{H}, \tilde{a} \in \mathcal{A}} Q^\Pi(\tilde{h}, \tilde{a}) B(\tilde{h}\tilde{a}|sb) \\
&\stackrel{(a)}{\leq} \sum_{\tilde{h} \in \mathcal{H}, \tilde{a} \in \mathcal{A}} \left( Q^\Pi(\psi^{-1}(s, b)) \pm \varepsilon \right) B(\tilde{h}\tilde{a}|sb) \\
&= Q^\Pi(\psi^{-1}(s, b)) \pm \varepsilon \tag{27}
\end{aligned}$$

(a) follows from the theorem's assumption. Since,  $Q^\Pi(\psi^{-1}(s, b))$  is a representative member in the pre-image set of  $(s, b)$ ; it is equivalent to say  $Q^\Pi(\psi^{-1}(s, b)) = Q^\Pi(h, a)$  for any  $\psi(h, a) = (s, b)$ . Therefore, combining (26) and (26) we get  $|Q^\Pi(h, a) - q^\pi(s, b)| \leq \gamma(\delta + \varepsilon_Q(s) + \frac{\varepsilon_\Pi(s)}{1-\gamma}) + \varepsilon$ , hence  $\delta \leq \frac{\varepsilon + \gamma\varepsilon_Q(s) + \frac{\gamma\varepsilon_\Pi(s)}{1-\gamma}}{1-\gamma}$ .  $\blacksquare$

## A.5 Proof of Theorem 9

*Proof.* Let  $\delta := \sup_{\tilde{s}, \tilde{h}: \psi(\tilde{h})=\tilde{s}} \left| V^\Pi(\tilde{h}) - v^\pi(\tilde{s}) \right|$  then for any  $\psi(h) = s$ , we have,

$$\begin{aligned}
\sum_{\tilde{b} \in \mathcal{B}} \langle Q^\Pi(\psi^{-1}(s, \tilde{b})) \rangle_B \pi(\tilde{b}|s) &\stackrel{(a)}{=} \sum_{\tilde{h} \in \mathcal{H}} B_{\mathcal{H}}(\tilde{h}|s) \sum_{\tilde{a} \in \mathcal{A}} Q^\Pi(\tilde{h}, \tilde{a}) \sum_{\tilde{b} \in \mathcal{B}} B_{\mathcal{A}}(\tilde{a}|\tilde{h}) \pi(\tilde{b}|s) \\
&\stackrel{(b)}{\leq} \sum_{\tilde{h}: \psi(\tilde{h})=s} B_{\mathcal{H}}(\tilde{h}|s) \left( V^\Pi(\tilde{h}) \pm \varepsilon_B \right) \\
&\stackrel{(c)}{\leq} \sum_{\tilde{h}: \psi(\tilde{h})=s} B_{\mathcal{H}}(\tilde{h}|s) \left( V^\Pi(h) \pm (\varepsilon + \varepsilon_B) \right) \\
&= V^\Pi(h) \pm (\varepsilon + \varepsilon_B) \tag{28}
\end{aligned}$$

(a) follows from the chain rule of joint distributions and since the history is action independently mapped; (b) and (c) follow from the theorem's assumptions. Further, we have,

$$\begin{aligned}
\left| \sum_{\tilde{b} \in \mathcal{B}} \langle Q^\Pi(\psi^{-1}(s, \tilde{b})) \rangle_B \pi(\tilde{b}|s) - v^\pi(s) \right| &\equiv \left| \sum_{\tilde{b} \in \mathcal{B}} \left( \langle Q^\Pi(\psi^{-1}(s, \tilde{b})) \rangle_B - q^\pi(s, \tilde{b}) \right) \pi(\tilde{b}|s) \right| \\
&\stackrel{(d)}{\leq} \sum_{\tilde{b} \in \mathcal{B}} \left| \langle Q^\Pi(\psi^{-1}(s, \tilde{b})) \rangle_B - q^\pi(s, \tilde{b}) \right| \pi(\tilde{b}|s) \\
&\stackrel{(e)}{\leq} \gamma \delta \tag{29}
\end{aligned}$$

(d) follows from the simple mathematical fact that  $|\sum_x f(x)| \leq \sum_x |f(x)|$ ; (e) uses Lemma 14. Now we prove the main result. Together with (28) and (29) we have,

$$\begin{aligned}
|V^\Pi(h) - v^\pi(s)| &\leq \left| V^\Pi(h) - \sum_{\tilde{b} \in \mathcal{B}} \langle Q^\Pi(\psi^{-1}(s, \tilde{b})) \rangle_B \pi(\tilde{b}|s) \right| \\
&\quad + \left| \sum_{\tilde{b} \in \mathcal{B}} \langle Q^\Pi(\psi^{-1}(s, \tilde{b})) \rangle_B \pi(\tilde{b}|s) - v^\pi(s) \right| \\
&\leq \gamma \delta + \varepsilon + \varepsilon_B
\end{aligned}$$

Hence,  $\delta \leq \frac{\varepsilon + \varepsilon_B}{1 - \gamma}$  and completes the proof. ■

## A.6 Proof of Theorem 8

*Proof.* (i) The proof follows the same steps as the proof of Theorem 7, replacing  $\Pi$  with  $\Pi^*$  and  $\pi$  with  $\pi^*$  and using Lemma 12 instead of Lemma 11. In the end we use Lemma 15 to conclude the proof.



(ii) For  $\psi(h) = s$  and  $\check{\Pi}(h) := \psi_s^{-1}(\pi^*(s))$ ,

$$V^*(h) \pm \frac{2\varepsilon}{1-\gamma} \stackrel{(i)}{\leq} v^*(s) \equiv q^*(s, \pi^*(s)) \stackrel{(i)}{\leq} Q^*\left(h, \check{\Pi}(h)\right) \pm \frac{2\varepsilon}{1-\gamma}$$

which implies  $\left|Q^*\left(h, \check{\Pi}(h)\right) - V^*(h)\right| \leq \frac{4\varepsilon}{1-\gamma}$  and Lemma 16 concludes the proof. ■

## A.7 Proof of Theorem 10

*Proof.* Let us define  $\pi_h(s)$  such that  $(s, \pi_h(s)) := \psi(h, \Pi^*(h))$  for  $\psi(h) = s$ . Then,

$$q^{\pi_h}(s, \pi_h(s)) = v^{\pi_h}(s) \stackrel{(a)}{\leq} V^*(h) \pm \frac{\varepsilon + \varepsilon_B}{1-\gamma} \quad (30)$$

(a) follows from Theorem 9 applied to  $\Pi = \Pi^*$  (with  $\pi = \pi_h$ ). Now we derive a bound for any  $b \in \mathcal{B}$ .

$$\begin{aligned} q^{\pi_h}(s, b) - \frac{\gamma(\varepsilon + \varepsilon_B)}{1-\gamma} &\stackrel{Thm.9}{\leq} \langle Q^*(\psi^{-1}(s, b)) \rangle_B \equiv \sum_{\tilde{h} \in \mathcal{H}, \tilde{a} \in \mathcal{A}} Q^*(\tilde{h}, \tilde{a}) B(\tilde{h}\tilde{a}|sb) \\ &= \sum_{\tilde{h} \in \mathcal{H}} B_{\mathcal{H}}(\tilde{h}|s) \sum_{\tilde{a} \in \mathcal{A}} Q^*(\tilde{h}, \tilde{a}) B_{\mathcal{A}}(\tilde{a}|sb\tilde{h}) \\ &\stackrel{(b)}{\leq} \sum_{\tilde{h} \in \mathcal{H}} B_{\mathcal{H}}(\tilde{h}|s) \sum_{\tilde{a} \in \mathcal{A}} Q^*(\tilde{h}, \Pi^*(\tilde{h})) B_{\mathcal{A}}(\tilde{a}|sb\tilde{h}) \\ &= \sum_{\tilde{h}: \psi(\tilde{h})=s} B_{\mathcal{H}}(\tilde{h}|s) V^*(\tilde{h}) \\ &\stackrel{(c)}{\leq} V^*(h) + \varepsilon \end{aligned} \quad (31)$$

(b) is due to the definition of optimal value and (c) follows from the theorem's assumptions. Together (30) and (31) imply,

$$v^{\pi_h}(s) = q^{\pi_h}(s, \pi_h(s)) \leq \max_{\tilde{b} \in \mathcal{B}} q^{\pi_h}(s, \tilde{b}) \stackrel{(31)}{\leq} V^*(h) + \frac{\varepsilon + \gamma\varepsilon_B}{1-\gamma} \stackrel{(30)}{\leq} v^{\pi_h}(s) + \frac{2(\varepsilon + \varepsilon_B)}{1-\gamma} \quad (32)$$

(ii) For  $\varepsilon = \varepsilon_B = 0$ , the previous equation implies  $v^{\pi_h}(s) = \max_{\tilde{b} \in \mathcal{B}} q^{\pi_h}(s, \tilde{b})$ . It shows that  $\pi_h(s) = \pi^*(s)$  for all  $\psi(h) = s$ .

(i) Now for general  $\varepsilon + \varepsilon_B > 0$  case. For all  $s \in \mathcal{S}$  and  $b \in \mathcal{B}$  we have,

$$0 \leq q^*(s, b) - q^{\pi_h}(s, b) \equiv \sum_{\tilde{s} \in \mathcal{S}, \tilde{r} \in \mathcal{R}} p(\tilde{s}\tilde{r}|sb) \gamma (v^*(\tilde{s}) - v^{\pi_h}(\tilde{s})) \stackrel{(d)}{\leq} \gamma \max_{\tilde{s} \in \mathcal{S}} (v^*(\tilde{s}) - v^{\pi_h}(\tilde{s}))$$

$$\begin{aligned}
0 \leq v^*(s) - v^{\pi_h}(s) &\stackrel{(e)}{\leq} \max_{\tilde{b} \in \mathcal{B}} q^*(s, \tilde{b}) - \max_{\tilde{b} \in \mathcal{B}} q^{\pi_h}(s, \tilde{b}) + \frac{2(\varepsilon + \varepsilon_B)}{1 - \gamma} \\
&\stackrel{(f)}{\leq} \max_{\tilde{b} \in \mathcal{B}} \left( q^*(s, \tilde{b}) - q^{\pi_h}(s, \tilde{b}) \right) + \frac{2(\varepsilon + \varepsilon_B)}{1 - \gamma}
\end{aligned}$$

(d) expectation is replaced by maximum operation; (e) follows from the definition of  $v^*(s)$  and (32); (f) is a simple mathematical fact of maximization operation. Together this implies,

$$\begin{aligned}
\max_{\tilde{s} \in \mathcal{S}} (v^*(\tilde{s}) - v^{\pi_h}(\tilde{s})) &\leq \gamma \max_{\tilde{s} \in \mathcal{S}} (v^*(\tilde{s}) - v^{\pi_h}(\tilde{s})) + \frac{2(\varepsilon + \varepsilon_B)}{1 - \gamma} \\
\Rightarrow \max_{\tilde{s} \in \mathcal{S}} (v^*(\tilde{s}) - v^{\pi_h}(\tilde{s})) &\leq \frac{2(\varepsilon + \varepsilon_B)}{(1 - \gamma)^2} \tag{33}
\end{aligned}$$

Hence for any  $\psi(h) = s$ , we have,

$$\begin{aligned}
V^*(h) - \frac{\varepsilon + \varepsilon_B}{1 - \gamma} &\stackrel{(30)}{\leq} v^{\pi_h}(s) \\
&\stackrel{(g)}{\leq} v^*(s) \\
&\stackrel{(33)}{\leq} v^{\pi_h}(s) + \frac{2(\varepsilon + \varepsilon_B)}{(1 - \gamma)^2} \\
&\stackrel{(30)}{\leq} V^*(h) + \frac{\varepsilon + \varepsilon_B}{1 - \gamma} + \frac{2(\varepsilon + \varepsilon_B)}{(1 - \gamma)^2} \leq V^*(h) + \frac{3(\varepsilon + \varepsilon_B)}{(1 - \gamma)^2}
\end{aligned}$$

(g) holds by definition. Hence, the main results follows by using Lemma 14 with  $\Pi$  replaced by  $\Pi^*$  and  $\pi$  by  $\pi^*$ .  $\blacksquare$

## B Lemmas

In this section, we establish a couple of important lemmas that bound the (optimal) value loss when we evaluate the action-value function only at the representatives.

**Lemma 11.** *For any policy  $\Pi$ , and  $\psi(h) = s$  the following holds,*

$$\left| V^\Pi(h) - \sum_{\tilde{b} \in \mathcal{B}} Q^\Pi(\psi^{-1}(s, \tilde{b})) \pi_R(\tilde{b}|s) \right| \leq \varepsilon_Q(s) + \frac{\varepsilon_\Pi(s)}{1 - \gamma}.$$

*Proof.* For any  $\psi(h) = s$ , we start from the value function of the original process,

$$\begin{aligned}
V^\Pi(h) &\equiv \sum_{\tilde{a} \in \mathcal{A}} Q^\Pi(h, \tilde{a}) \Pi(\tilde{a}|h) \\
&\stackrel{(a)}{=} \sum_{\tilde{b} \in \mathcal{B}} \sum_{\tilde{a} \in \psi_s^{-1}(\tilde{b})} Q^\Pi(h, \tilde{a}) \Pi(\tilde{a}|h) \\
&= \sum_{\tilde{b} \in \mathcal{B}} \sum_{\tilde{a} \in \psi_s^{-1}(\tilde{b})} \left( Q^\Pi(h, \tilde{a}) + Q^\Pi(\psi^{-1}(s, \tilde{b})) - Q^\Pi(\psi^{-1}(s, \tilde{b})) \right) \Pi(\tilde{a}|h) \\
&= \sum_{\tilde{b} \in \mathcal{B}} Q^\Pi(\psi^{-1}(s, \tilde{b})) \sum_{\tilde{a} \in \psi_s^{-1}(\tilde{b})} \Pi(\tilde{a}|h) + \sum_{\tilde{b} \in \mathcal{B}} \sum_{\tilde{a} \in \psi_s^{-1}(\tilde{b})} \left( Q^\Pi(h, \tilde{a}) - Q^\Pi(\psi^{-1}(s, \tilde{b})) \right) \Pi(\tilde{a}|h) \\
&\stackrel{(b)}{\leq} \sum_{\tilde{b} \in \mathcal{B}} Q^\Pi(\psi^{-1}(s, \tilde{b})) \Pi_\psi(\tilde{b}|h) + \varepsilon_Q(s) \\
&= \sum_{\tilde{b} \in \mathcal{B}} Q^\Pi(\psi^{-1}(s, \tilde{b})) \left( \Pi_\psi(\tilde{b}|h) + \pi_R(\tilde{b}|s) - \pi_R(\tilde{b}|s) \right) + \varepsilon_Q(s) \\
&= \sum_{\tilde{b} \in \mathcal{B}} Q^\Pi(\psi^{-1}(s, \tilde{b})) \pi_R(\tilde{b}|s) + \sum_{\tilde{b} \in \mathcal{B}} Q^\Pi(\psi^{-1}(s, \tilde{b})) \left( \Pi_\psi(\tilde{b}|h) - \pi_R(\tilde{b}|s) \right) + \varepsilon_Q(s) \\
&\stackrel{(c)}{\leq} \sum_{\tilde{b} \in \mathcal{B}} Q^\Pi(\psi^{-1}(s, \tilde{b})) \pi_R(\tilde{b}|s) + \frac{\varepsilon_\Pi(s)}{1 - \gamma} + \varepsilon_Q(s)
\end{aligned}$$

(a) follows from the fact that the mapping is defined to be surjective; (b) from (8) and (17), (c) uses (14) and the fact that action-value function is bounded, so is the representative. It is easy to get the other side of the inequality from similar steps.  $\blacksquare$

Now, we bound the error of evaluating the action-value function only at the representatives of the mapping when the agent is following the optimal policy in the original process.

**Lemma 12.** *For the optimal policy  $\Pi^*$ , and for all  $\psi(h) = s$  the following holds,*

$$\left| V^*(h) - \max_{\tilde{b} \in \mathcal{B}} Q^*(\psi^{-1}(s, \tilde{b})) \right| \leq \varepsilon_Q(s).$$

*Proof.* For any  $\psi(h) = s$  we have,

$$\begin{aligned}
V^*(h) &\equiv \max_{\tilde{a} \in \mathcal{A}} Q^*(h, \tilde{a}) \stackrel{(a)}{=} \max_{\tilde{b} \in \mathcal{B}} \max_{\tilde{a} \in \psi_s^{-1}(\tilde{b})} Q^*(h, \tilde{a}) \\
&\stackrel{(17)}{\leq} \max_{\tilde{b} \in \mathcal{B}} Q^*(\psi^{-1}(s, \tilde{b})) \pm \varepsilon_Q(s)
\end{aligned}$$

(a) follows from the fact that mapping is defined to be surjective.  $\blacksquare$

**Lemma 13.** Let  $P_\psi$  be an MDP then  $Q^*(h,a) = Q^*(h',a')$  for all  $\psi(h,a) = \psi(h',a')$ .

*Proof.* Let  $\delta := \sup_{h,h':\psi(h)=\psi(h'),a,a'} |Q^*(h,a) - Q^*(h',a')|$ , then for all  $\psi(h) = \psi(h')$ ,

$$\begin{aligned} |V^*(h) - V^*(h')| &\equiv \left| \max_{\tilde{a} \in \mathcal{A}} Q^*(h, \tilde{a}) - \max_{\tilde{a} \in \mathcal{A}} Q^*(h', \tilde{a}) \right| \\ &\stackrel{(a)}{\leq} \max_{\tilde{a} \in \mathcal{A}} |Q^*(h, \tilde{a}) - Q^*(h', \tilde{a})| \leq \delta \end{aligned} \quad (34)$$

(a) follows from simple mathematical fact of maximum value. Now for all  $\psi(h,a) = \psi(h',a') = (s,b)$ ,

$$\begin{aligned} |Q^*(h,a) - Q^*(h',a')| &\stackrel{(b)}{=} \left| \sum_{\tilde{s} \in \mathcal{S}, \tilde{r} \in \mathcal{R}} P_\psi(\tilde{s}\tilde{r}|ha)(\tilde{r} + \gamma V^*(\tilde{h})) \right. \\ &\quad \left. - \sum_{\tilde{s} \in \mathcal{S}, \tilde{r} \in \mathcal{R}} P_\psi(\tilde{s}\tilde{r}|h'a')(\tilde{r} + \gamma V^*(\tilde{h}')) \right| \\ &\stackrel{(15)}{=} \gamma \left| \sum_{\tilde{s} \in \mathcal{S}, \tilde{r} \in \mathcal{R}} p(\tilde{s}\tilde{r}|sb) (V^*(\tilde{h}) - V^*(\tilde{h}')) \right| \\ &\stackrel{(34)}{\leq} \gamma \delta \end{aligned} \quad (35)$$

(b) follows from the definition and  $\tilde{h} = ha\tilde{o}\tilde{r}$  and  $\tilde{h}' = h'a'\tilde{o}\tilde{r}$ , and  $\psi(\tilde{h}) = \psi(\tilde{h}') = \tilde{s}$ . By (35) we have  $\delta \leq \gamma\delta \Rightarrow \delta = 0$ .  $\blacksquare$

The following lemma is a stepping stone for the main results. It establishes a link between the value functions loss and the action-value functions loss.

**Lemma 14.** ( $\langle Q \rangle_B - q$ ) Let  $|V^\Pi(h) - v^\pi(s)| \leq \varepsilon$  for all  $\psi(h) = s$ . Then for all  $s \in \mathcal{S}$  and  $b \in \mathcal{B}$  it holds:

$$|\langle Q^\Pi(\psi^{-1}(s,b)) \rangle_B - q^\pi(s,b)| \leq \gamma\varepsilon.$$

*Proof.* We begin with the action value of the representative of any  $(s,b)$  as

$$\begin{aligned} \langle Q^\Pi(\psi^{-1}(s,b)) \rangle_B &\equiv \sum_{\tilde{h} \in \mathcal{H}, \tilde{a} \in \mathcal{A}} B(\tilde{h}\tilde{a}|sb) \sum_{\tilde{o} \in \mathcal{O}, \tilde{r} \in \mathcal{R}} P(\tilde{o}\tilde{r}|\tilde{h}\tilde{a}) (\tilde{r} + \gamma V^\Pi(\tilde{h}\tilde{a}\tilde{o}\tilde{r})) \\ &\stackrel{(a)}{=} \sum_{\tilde{h} \in \mathcal{H}, \tilde{a} \in \mathcal{A}} B(\tilde{h}\tilde{a}|sb) \sum_{\tilde{s} \in \mathcal{S}} \sum_{\tilde{r} \in \mathcal{R}, \tilde{o}: \psi(\tilde{h}\tilde{a}\tilde{o}\tilde{r}) = \tilde{s}} P(\tilde{o}\tilde{r}|\tilde{h}\tilde{a}) (\tilde{r} + \gamma V^\Pi(\tilde{h}\tilde{a}\tilde{o}\tilde{r})) \\ &\stackrel{(b)}{\leq} \sum_{\tilde{h} \in \mathcal{H}, \tilde{a} \in \mathcal{A}} B(\tilde{h}\tilde{a}|sb) \sum_{\tilde{s} \in \mathcal{S}, \tilde{r} \in \mathcal{R}} P_\psi(\tilde{s}\tilde{r}|\tilde{h}\tilde{a}) (\tilde{r} + \gamma v^\pi(\tilde{s}) \pm \gamma\varepsilon) \\ &\equiv \sum_{\tilde{s} \in \mathcal{S}, \tilde{r} \in \mathcal{R}} p_B(\tilde{s}\tilde{r}|sb) (\tilde{r} + \gamma v^\pi(\tilde{s})) \pm \gamma\varepsilon \equiv q^\pi(s,b) \pm \gamma\varepsilon \end{aligned}$$

(a) follows since  $\psi$  is surjective; (b) follows from the assumption. ■

**Lemma 15.** ( $\varepsilon_{Q\Pi}$ ) *Let  $|Q^\Pi(h,a) - Q^\Pi(h',a')| \leq \varepsilon$  for all  $\psi(h,a) = \psi(h',a')$ . Then  $\varepsilon_Q(s) \leq \varepsilon$  for all  $s \in \mathcal{S}$ .*

*Proof.* For all  $s \in \mathcal{S}$ ,

$$\varepsilon_Q(s) \equiv \sup_{\tilde{h}, \tilde{a}, \tilde{b}: \psi(\tilde{h}, \tilde{a}) = (s, \tilde{b})} |Q^\Pi(\psi^{-1}(s, \tilde{b})) - Q^\Pi(\tilde{h}, \tilde{a})| \stackrel{(a)}{\leq} \varepsilon$$

(a) follows from the assumption and the fact that  $Q^\Pi(\psi^{-1}(s, b))$  can be any member in the pre-image set of  $(s, b)$ . ■

Moreover, we need Lemma 8 from [6]. For completeness, we state the lemma in this work without repeating the proof. It gives value loss-bounds between the optimal policy and a policy that has bounded one-step action-value loss.

**Lemma 16.** ([6] Lemma 8.  $\Pi(h) \neq \Pi^*(h)$ ) *If  $Q^*(h, \Pi(h)) \geq V^*(h) - \varepsilon$  for all  $h$  and some policy  $\Pi$ , then for all  $h$  and  $a$  it follows:*

$$0 \leq Q^*(h, a) - Q^\Pi(h, a) \leq \frac{\gamma \varepsilon}{1 - \gamma} \quad \text{and} \quad 0 \leq V^*(h) - V^\Pi(h) \leq \frac{\varepsilon}{1 - \gamma}.$$

## C List of Notation

Symbol	Explanation
$\mathbb{R}$	set of real numbers
$\times$	Cartesian product
$\Delta(X)$	probability distribution over a set $X$
$X^T$	transpose of matrix $X$
$\tilde{x}$	local variable
$x'$	different member of a set
$P$	original process
$\mathcal{A}, \mathcal{O}, \mathcal{R}$	continuous action, observation and reward spaces of the original process
$a, o, r$	action, observation and reward tuple $\in \mathcal{A} \times \mathcal{O} \times \mathcal{R}$
$\mathcal{H}$	set of all histories
$h$	history $\in \mathcal{H}$
$Q, V, \Pi$	(action-)value function and policy of the original process
$\Pi^*$	optimal policy of the original process
$\check{\Pi}$	elevated policy on the original process from the abstract-process
$\gamma$	discount factor

$\psi$	homomorphism map
$P_\psi$	marginalized abstract-process
$p_B$	surrogate MDP process
$p_{\text{MDP}}$	abstract MDP process
$\mathcal{S}, \mathcal{B}$	finite state and action spaces of the abstract-process
$s, b$	state action pair $\in \mathcal{S} \times \mathcal{B}$
$q, v, \pi$	(action-)value function and policy of the surrogate MDP process
$\pi^*$	optimal policy of the surrogate MDP process
$\Pi_\psi$	history-dependent abstract policy
$\psi_b^{-1}(s)$	set of histories mapped to $(s, b)$ pair
$\psi_s(b)$	history-dependent set of (original) actions mapped to an $(s, b)$ pair
$Q^\Pi(\psi^{-1}(s, b))$	action value representative of an $(s, b)$ partition
$\varepsilon_\Pi$	maximum variation among abstracted policy members
$\varepsilon_Q$	maximum variation among abstracted action value members
$B$	stochastic inverse of the homomorphism map
$B^\pi$	$B$ and $\pi$ induced measure on the original action space
$\langle \cdot \rangle_B$	$B$ average
$\varepsilon, \varepsilon_B, \varepsilon_{\max}$	small positive error constants

## D Example Aggregations

In this section, we provide a toy example to illustrate the possibility of approximately joint state-action aggregation beyond MDPs. In the example, contrary to ESA, we also aggregate the approximately similar policy state-action pairs. These pairs are not aggregated in ESA due to the exact policy similarity condition, *cf.* Theorem 5 of [6].

For simplicity, we assume that actions and histories are mapped independently and the original environment  $P$  is an MDP. We define the original action and observation spaces as  $\mathcal{A} = \mathcal{O} = \{1, 2, 3, 4\}$ . Moreover,  $\mathcal{S} = \{X, Y\}$ ,  $\mathcal{B} = \{\alpha, \beta\}$ . The original state-action pairs are represented by dots and the shaded regions indicate the mapping function. The dynamics are (fictitious) region constant (see Figure 4)<sup>9</sup>.

We assume an arbitrary policy which may depend on the history rather than only on the last observation. This allows complex dynamics for the proof of concept. We

---

<sup>9</sup>Only the dots are elements of the joint space and the regions are fictitious. The aggregated pair  $(s, b)$  is indicated adjacent to each region.




Figure 4: **(Left)**: Non-MDP aggregation, **(Middle)**: Approximate aggregation, **(Right)**: Violating policy uniformity condition

express the environment as

$$\begin{aligned}
 P(o'|oa) &= \sum_{\tilde{a} \in \mathcal{A}} P(o'|oa)\Pi(\tilde{a}|oa o') \\
 &= \sum_{\tilde{a} \in \mathcal{A}} P^\Pi(o'\tilde{a}|oa).
 \end{aligned}$$

We express the dynamics with a joint measure  $P^\Pi$  and do not distinguish between the policy and the environment unless otherwise stated. Let the rewards be a function of the originating region and the problem has a finite set of real-valued rewards. We present three cases in this example: non-MDP dynamics, approximate action-values and policy disagreement.

**Non-MDP Example.** In the first case, this example (see Figure 4:Left) demonstrates a non-MDP aggregation. Let the problem have the following region uniform transition probability matrix in the observation-action space:

$$P^\Pi = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1/2 & 1/2 \\ 1 & 0 & 0 & 0 & 0 \\ 1/2 & 0 & 0 & 1/4 & 1/4 \end{bmatrix} \quad (36)$$

where  $P_{ij}^\Pi$  is the probability of reaching region  $R_j$  if the current  $(o,a) \in R_i$ <sup>10</sup>. Formally,  $P_{ij}^\Pi = \sum_{(\tilde{o}, \tilde{a}) \in R_j} P^\Pi(\tilde{o}\tilde{a}|oa \in R_j)$ . The marginalized process is expressed as

$$P_\psi(X|ha) = P_\psi(X|oa) = \sum_{\tilde{o}: \psi(oa\tilde{o})=X, \tilde{a} \in \mathcal{A}} P^\Pi(\tilde{o}\tilde{a}|oa).$$

<sup>10</sup>The indexes should be read in order — i.e. 1,2,3,4a,4b.

This is not even approximately an MDP. It is evident from the following probabilities of reaching state  $X$  from itself.

$$\begin{aligned} P_\psi(X|oa \in R_{4a}) &= P_{4a3}^\Pi + P_{4a4a}^\Pi + P_{4a4b}^\Pi = 0 \\ P_\psi(X|oa \in R_{4b}) &= P_{4b3}^\Pi + P_{4b4a}^\Pi + P_{4b4b}^\Pi = \frac{1}{2} \end{aligned}$$

The above equations show that reaching state  $X$  from the regions  $R_{4a}$  and  $R_{4b}$  are different. But, we can still aggregate the regions if they have similar action-values. Let the regional rewards are  $r = [0 \ 0 \ 0 \ \gamma \ 0]$ . The regional action values can be expressed as

$$Q^\Pi(o, a) = r(oa) + \gamma \sum_{\tilde{o} \in \mathcal{O}, \tilde{a} \in \mathcal{A}} Q^\Pi(\tilde{o}, \tilde{a}) P^\Pi(\tilde{o}\tilde{a}|oa). \quad (37)$$

It translates into a vectorization form with each region  $i$  has the action value expressed as

$$Q_i = r_i + \gamma \sum_j P_{ij}^\Pi Q_j. \quad (38)$$

By solving this system of equations we get the following region uniform action-value vector,

$$Q = [c - 2 \quad \gamma^2 c \quad \gamma c \quad c \quad c] \quad (39)$$

where  $c = \frac{2}{1-\gamma^3} \geq 2$ . Hence, the example shows that even the regions  $R_{4a}$  and  $R_{4b}$  have non-MDP dynamics, they can still be aggregated due to the same action-values.

**Approximate Q-value Example.** Now for the second case, we perform an approximate aggregation in region  $R_3$  (see Figure 4:Middle). Let the reward in  $R_{3b}$  is  $\varepsilon$ . The updated reward vector is  $r = [0 \ 0 \ 0 \ \varepsilon \ \gamma \ 0]$ . By keeping everything same as in the first example, the new transition matrix is given as

$$P^\Pi = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1/2 & 1/2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1/2 & 1/2 \\ 0 & 0 & 0 & 0 & 1/2 & 1/2 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 1/2 & 0 & 0 & 0 & 1/4 & 1/4 \end{bmatrix} \quad (40)$$

Similar to the first case, we get the aggregated action-values as

$$Q = [c - 2 \quad \frac{\gamma\varepsilon}{2} + \gamma^2 c \quad \gamma c \quad \gamma c + \varepsilon \quad c \quad c] \quad (41)$$

where  $c = \frac{\gamma^2\varepsilon + 4}{2(1-\gamma^3)} \geq 2$ . It shows that the regions  $R_{3a}$  and  $R_{3b}$  can be approximately aggregated together.

**Approximate Policy Example.** As the last case, let us divide the region  $R_2$  into two regions with different policies (see Figure 4:Right). Let the policy is approximately similar, i.e.  $|\Pi(a|o \in R_{2a}) - \Pi(a|o \in R_{2b})| = \varepsilon'$  for all  $\psi(o, a) = (Y, \beta)$ . It makes



region  $R_1$  approachable from itself because  $\varepsilon'$  weight of the aggregated action  $\beta$  in region  $R_{2b}$  is distributed to the aggregated action  $\alpha$  in region  $R_1$ . This effectively translates into the following region uniform transition matrix,

$$P^\Pi = \begin{bmatrix} \varepsilon' & 1/2 & 1/2 - \varepsilon' & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1/2 & 1/2 & 0 & 0 \\ 0 & 0 & 0 & 1/2 & 1/2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1/2 & 1/2 \\ 0 & 0 & 0 & 0 & 0 & 1/2 & 1/2 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1/2 & 0 & 0 & 0 & 0 & 1/4 & 1/4 \end{bmatrix}$$

The reward structure is the same as in the previous case,  $r = [0 \ 0 \ 0 \ 0 \ \varepsilon \ \gamma \ 0]$ . Finally, we get the action-value vector as

$$Q = \begin{bmatrix} \gamma^2 g(\varepsilon') (\varepsilon/2 + \gamma c) \\ \frac{\gamma \varepsilon}{2} + \gamma^2 c \\ \frac{\gamma \varepsilon}{2} + \gamma^2 c \\ \gamma c \\ \gamma c + \varepsilon \\ c \\ c \end{bmatrix}^T \quad (42)$$

where  $c = \frac{4 + \gamma^2 \varepsilon g(\varepsilon')}{2(1 - \gamma^3 g(\varepsilon'))}$  and  $g(\varepsilon') := \frac{1 - \varepsilon'}{1 - \gamma \varepsilon'}$ . This final case shows that although the regions  $R_{2a}$  and  $R_{2b}$  have different policies they still have same action-values. Hence, the regions can be aggregated together exactly. It allows us to have coarser maps than ESA permits.