
Reward-Free Policy Space Compression for Reinforcement Learning

Mirco Mutti*
Politecnico di Milano
Università di Bologna

Stefano Del Col
Politecnico di Milano

Marcello Restelli
Politecnico di Milano

Abstract

In reinforcement learning, we encode the potential behaviors of an agent interacting with an environment into an infinite set of policies, the *policy space*, typically represented by a family of parametric functions. Dealing with such a policy space is a hefty challenge, which often causes sample and computation inefficiencies. However, we argue that a limited number of policies are actually relevant when we also account for the structure of the environment and of the policy parameterization, as many of them would induce very similar interactions, i.e., state-action distributions. In this paper, we seek for a reward-free *compression* of the policy space into a finite set of representative policies, such that, given any policy π , the minimum Rényi divergence between the state-action distributions of the representative policies and the state-action distribution of π is bounded. We show that this compression of the policy space can be formulated as a set cover problem, and it is inherently NP-hard. Nonetheless, we propose a game-theoretic reformulation for which a locally optimal solution can be efficiently found by iteratively stretching the compressed space to cover an adversarial policy. Finally, we provide an empirical evaluation to illustrate the compression procedure in simple domains, and its ripple effects in reinforcement learning.

1 INTRODUCTION

In the Reinforcement Learning (RL) (Sutton and Barto, 2018) framework, an artificial agent interacts

with an environment, typically modeled through a Markov Decision Process (MDP) (Puterman, 2014), to maximize some form of long-term performance, which is usually the sum of the discounted rewards collected in the process. The agent’s behavior is encoded in a Markovian *policy*, i.e., a function that maps the current state of the environment with a probability distribution over the next action to be taken. In principle, if the underlying MDP is small enough, we can represent a Markovian policy with a table that includes an entry for each state-action pair, and we call it a *tabular* policy. However, most relevant scenarios have too many (possibly infinite) states and actions to allow for a tabular representation. In this case, we can turn to function approximation (Sutton and Barto, 2018) to encode the policy within a family of parametric functions, e.g., a linear basis combination or a deep neural network, and we call it a *parametric* policy. This set of parametric policies, which we call the *policy space*, is typically infinite. Therefore, learning a policy that maximizes the performance can be a hefty challenge, and the sheer size of the policy space often causes sample and computation inefficiencies.

A setting where these inefficiencies arise clearly and naturally is Policy Optimization (PO) (Deisenroth et al., 2013). In PO, we aim to find a policy that maximizes the performance within the policy space, i.e., an *optimal* policy, with the least amount of interactions (Sutton et al., 1999; Silver et al., 2014; Schulman et al., 2015; Metelli et al., 2018). If we also account for the performance of the policies that are actually deployed to collect these interactions, we come up with an online PO (Papini et al., 2019; Cai et al., 2020). In this setting, we try to minimize the *regret* that the agent suffers by taking interactions with a sub-optimal behavior before converging to an optimal policy. Recent results showed that the regret of online PO is directly related to the size of the policy space (Papini et al., 2019; Metelli et al., 2021a). In particular, online PO with a finite policy space can enjoy a constant regret, i.e., it does not scale with the number of interactions, under certain conditions (Metelli et al., 2021a). Instead, the regret of online PO with an infi-

Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151. Copyright 2022 by the author(s). (*) Correspondence to <mirco.mutti@polimi.it>

nite policy space does scale with the square root of the number of interactions in general (Papini et al., 2019), which means that we only have asymptotic guarantees of reaching an optimal policy. In view of these results, one could wonder whether the expressive power of an infinite policy space is worth the additional regret it causes: Are all of these infinitely many policies really necessary for PO? The expressive power of a policy space is related to the different distributions that its policies can induce over the states and actions of the environment, as the whole point of PO is to find a policy that maximizes the probability of reaching state-action pairs associated with high rewards. However, different parameterizations might actually induce equivalent policies due to the specific structure of the policy space. Similarly, even different policies can induce the same state-action distribution in a given environment. These two types of policies are arguably redundant for PO and we would like to find a policy space that does not include either. Especially, we aim to answer the following question:

Having an infinite parametric policy space Θ in a given environment \mathcal{M} , can we compress Θ into a finite subset that retains most of its expressive power?

In this paper, we formulate this question into the *Policy Space Compression* problem, where we exploit the inherent structure of \mathcal{M} and Θ to compute the compressed policy space. The general idea is to identify a finite set of representative policies, such that for any policy π of the original space, the minimum Rényi divergence between the state-action distributions of the representative policies and the state-action distribution of π is bounded by a given constant. This compression is agnostic to the reward function, and thus the resulting policy space can benefit the computational and sample complexity of any RL task one can later specify over \mathcal{M} , as it is typical in reward-free RL (Hazan et al., 2019; Jin et al., 2020a).

Specifically, the paper includes the following contributions. First, we provide a formal definition of the policy space compression problem (Section 3). We note that the problem can be formulated equivalently as a set cover, and that finding an optimal compression of the policy space is NP-hard in general (Feige, 1998). Despite this negative result, we propose a game-theoretic reformulation (Section 4) that casts the problem to the one of reaching a differential Stackelberg equilibrium (Fiez et al., 2020) of a two-player sequential game, in which the first player tries to cover the policy space with a finite set of policies and the second player tries to find a policy that falls outside this coverage. Then, we present a planning algorithm (Section 5) to efficiently compute a compression of the policy space in a given environment, by repeat-

edly solving, with a first-order method, the two-player game for an increasing number of covering policies, until the compression requirement is met globally. In Section 6, we provide a theoretical analysis of the performance guarantees attained by the compressed policy space in relevant RL tasks, namely policy evaluation and policy optimization. Finally, in Section 7 we provide a brief numerical validation of both the compression algorithm and RL with the compressed policy space. The proofs of the theorems can be found in Appendix A.

2 PRELIMINARIES

In this section, we introduce the essential background on controlled Markov processes, policy optimization, importance sampling estimation and Rényi divergence. Throughout the paper, we will denote a vector \mathbf{v} with a bold typeface, as opposed to a scalar v .

2.1 Controlled Markov Processes

A discrete-time Controlled Markov Process (CMP) is defined as a tuple $\mathcal{M} := (\mathcal{S}, \mathcal{A}, P, \mu, \gamma)$, in which \mathcal{S} is the state space, \mathcal{A} is the action space, $P : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ is a transition model such that the next state is drawn as $s' \sim P(\cdot|s, a)$ given the current state $s \in \mathcal{S}$ and action $a \in \mathcal{A}$, $\mu : \Delta(\mathcal{S})$ is an initial state distribution such that the initial state is drawn as $s \sim \mu(\cdot)$, and $\gamma \in [0, 1]$ is the discount factor. The behavior of an agent interacting with a CMP can be modeled through a Markovian parametric policy $\pi_{\theta} : \mathcal{S} \rightarrow \Delta(\mathcal{A})$ such that an action is drawn as $a \sim \pi_{\theta}(\cdot|s)$ given the current state $s \in \mathcal{S}$, where $\theta \in \Theta \subseteq \mathbb{R}^m$ are the policy parameters, and the set Π_{Θ} is called the *policy space*. A policy π_{θ} induces a γ -discounted state distribution $d_{\pi_{\theta}}^s : \Delta(\mathcal{S})$ over the state space of the CMP \mathcal{M} , which is given by $d_{\pi_{\theta}}^s(s) = (1 - \gamma) \sum_{t=1}^{\infty} \gamma^t Pr(s_t = s)$ or the equivalent recursive relation $d_{\pi_{\theta}}^s(s) = (1 - \gamma)\mu(s) - \gamma \int_{\mathcal{S}\mathcal{A}} d_{\pi_{\theta}}^s(s') \pi_{\theta}(a'|s') P(s|s', a') ds' da'$. Similarly, we define the γ -discounted state-action distribution $d_{\pi_{\theta}}^{sa} : \Delta(\mathcal{S} \times \mathcal{A})$ given by $d_{\pi_{\theta}}^{sa}(s, a) = \pi_{\theta}(a|s) d_{\pi_{\theta}}^s(s)$. With a slight overloading of notation, we will indifferently denote the parametric policy space Π_{Θ} by Θ , a parametric policy $\pi_{\theta} \in \Pi_{\Theta}$ by θ , and its induced distributions $d_{\pi_{\theta}}^s(s), d_{\pi_{\theta}}^{sa}(s, a)$ by $d_{\theta}^s(s), d_{\theta}^{sa}(s, a)$.

2.2 Policy Optimization

The process of looking for the policy that maximizes the agent’s performance on a given RL task with a direct search in the policy space is called Policy Optimization (PO) (Deisenroth et al., 2013). The task is generally modeled through a Markov Decision Process (MDP) (Puterman, 2014) $\mathcal{M}^{\mathcal{R}} := \mathcal{M} \cup \mathcal{R}$, i.e.,

the combination of a CMP \mathcal{M} and a reward function $\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow [-R_{\max}, R_{\max}]$ such that $R(s, a)$ is the bounded reward that the agent collects by selecting action $a \in \mathcal{A}$ in state $s \in \mathcal{S}$, and $R_{\max} < \infty$. The agent’s performance is defined by the expected sum of discounted rewards collected by its policy, i.e.,

$$\begin{aligned} J(\boldsymbol{\theta}) &:= \mathbb{E}_{\substack{s_0 \sim \mu(\cdot), a_t \sim \pi_{\boldsymbol{\theta}}(\cdot | s_t) \\ s_{t+1} \sim P(\cdot | s_t, a_t)}} \left[\sum_{t=1}^{\infty} \gamma^t \mathcal{R}(s_t, a_t) \right] \\ &= \frac{1}{(1-\gamma)} \mathbb{E}_{(s,a) \sim d_{\boldsymbol{\theta}}^{sa}} [\mathcal{R}(s, a)], \end{aligned}$$

A Monte-Carlo estimate of the performance can be computed from a batch of N samples $\{s_n, a_n\}_{n=1}^N$ taken with the policy $\pi_{\boldsymbol{\theta}}$ in the γ -discounted MDP $\mathcal{M}^{\mathcal{R}}$ as $\hat{J}(\boldsymbol{\theta}) = \frac{1}{(1-\gamma)N} \sum_{n=1}^N \mathcal{R}(s_n, a_n)$.

2.3 Importance Sampling and Rényi Divergence

Importance Sampling (IS) (Cochran, 2007; Owen, 2013) is a common technique to estimate the expectation of a function under a *target* distribution by taking samples from a different distribution. In PO, importance sampling allows for estimating the performance of a target policy $\pi_{\boldsymbol{\theta}'}$ through a batch of samples $\{s_n, a_n\}_{n=1}^N$ taken with a policy $\pi_{\boldsymbol{\theta}}$. Especially, we define the importance weight $w_{\boldsymbol{\theta}'/\boldsymbol{\theta}}(s, a) := d_{\boldsymbol{\theta}'}^{sa}(s, a)/d_{\boldsymbol{\theta}}^{sa}(s, a)$. A Monte-Carlo estimate of $J(\boldsymbol{\theta}')$ via importance sampling is given by

$$\hat{J}_{IS}(\boldsymbol{\theta}'/\boldsymbol{\theta}) = \frac{1}{(1-\gamma)N} \sum_{n=1}^N w_{\boldsymbol{\theta}'/\boldsymbol{\theta}}(s_n, a_n) \mathcal{R}(s_n, a_n).$$

The latter estimator is known to be unbiased, i.e., $\mathbb{E}_{\boldsymbol{\theta}}[\hat{J}_{IS}(\boldsymbol{\theta}'/\boldsymbol{\theta})] = J(\boldsymbol{\theta}')$ (Owen, 2013). However, $\hat{J}_{IS}(\boldsymbol{\theta}'/\boldsymbol{\theta})$ might suffer from a large variance whenever the importance weights $w_{\boldsymbol{\theta}'/\boldsymbol{\theta}}(s, a)$ have a large variance. The variance of the importance weights is related to the exponentiated 2-Rényi divergence $D_2(d_{\boldsymbol{\theta}'}^{sa} || d_{\boldsymbol{\theta}}^{sa})$ (Rényi et al., 1961) through $\mathbb{V}\text{ar}_{(s,a) \sim d_{\boldsymbol{\theta}}^{sa}}[w_{\boldsymbol{\theta}'/\boldsymbol{\theta}}(s, a)] = D_2(d_{\boldsymbol{\theta}'}^{sa} || d_{\boldsymbol{\theta}}^{sa}) - 1$ (Cortes et al., 2010), where

$$D_2(d_{\boldsymbol{\theta}'}^{sa} || d_{\boldsymbol{\theta}}^{sa}) := \int_{\mathcal{S}\mathcal{A}} d_{\boldsymbol{\theta}}^{sa}(s, a) \left(\frac{d_{\boldsymbol{\theta}'}^{sa}(s, a)}{d_{\boldsymbol{\theta}}^{sa}(s, a)} \right)^2 ds da.$$

The latter has been employed in (Metelli et al., 2018) to upper bound the variance of the importance sampling estimator as $\mathbb{V}\text{ar}_{(s,a) \sim d_{\boldsymbol{\theta}}^{sa}}[\hat{J}_{IS}(\boldsymbol{\theta}'/\boldsymbol{\theta})] \leq \left(\frac{R_{\max}}{1-\gamma}\right)^2 D_2(d_{\boldsymbol{\theta}'}^{sa} || d_{\boldsymbol{\theta}}^{sa})/N$. In the following, we will refer to the exponentiated 2-Rényi divergence as the Rényi divergence.

3 THE POLICY SPACE COMPRESSION PROBLEM

Let us suppose to have a CMP \mathcal{M} the agent can interact with, and a parametric policy space Θ from which the agent can select its strategy of interaction. For the common parameterization choices, ranging from linear policies to deep neural networks, the policy space Θ is typically infinite. Dealing with such a large policy space to address the usual RL tasks, e.g., finding a convenient task-agnostic sampling strategy (Hazan et al., 2019) or seeking for an optimal policy within the set (Deisenroth et al., 2013), is often a huge challenge. Furthermore, many policies in Θ are unnecessary for these purposes, as they induce very similar interactions, and thus they have very similar performance. On the one hand, different policy parameters $\boldsymbol{\theta} \in \Theta$ might induce nearly identical distributions over actions. On the other hand, even different distributions over actions can lead to comparable state-action distributions due to the structure of the environment. Since we do not have any reward encoded in \mathcal{M} , it would be unwise to deem any state-action distribution irrelevant without additional information on the task structure. In this work, we aim to identify a subset of the policy space $\Theta' \subseteq \Theta$ that retains most of the expressive power of Θ , i.e., the set of the state-action distributions it can induce, while dramatically reducing its size, to the advantage of the computational and sample efficiency of future RL tasks. Especially, we consider a σ -soft compression of Θ , where for any policy $\boldsymbol{\theta} \in \Theta$ we would like to have a policy $\boldsymbol{\theta}' \in \Theta'$ such that the Rényi divergence between their respective state-action distributions $d_{\boldsymbol{\theta}'}^{sa}, d_{\boldsymbol{\theta}}^{sa}$ is bounded by a positive constant σ . The Rényi divergence is particularly convenient in this setting due to its relationship with the variance of the importance sampling in the off-policy estimation (Cortes et al., 2010; Metelli et al., 2018). The following statement provides a more formal definition of this σ -soft compression.

Definition 3.1 (σ -compression). *Let \mathcal{M} be a CMP, let Θ be a parametric policy space for \mathcal{M} , and let $\sigma > 0$ be a constant. We call Θ_{σ} a σ -compression of Θ in \mathcal{M} if it holds that $|\Theta_{\sigma}| < \infty$ and*

$$\forall \boldsymbol{\theta} \in \Theta, \quad \min_{\boldsymbol{\theta}' \in \Theta_{\sigma}} D_2(d_{\boldsymbol{\theta}}^{sa} || d_{\boldsymbol{\theta}'}^{sa}) \leq \sigma.$$

We call the task of finding a σ -compression of Θ in \mathcal{M} the *policy space compression* problem. Notably, for some $\mathcal{M}, \Theta, \sigma$, a σ -compression of Θ in \mathcal{M} might not exist, as infinitely many policies $\boldsymbol{\theta} \in \Theta$ might induce relevant state-action distributions. However, we note that those scenarios are not interesting for our purposes, as the PO problem would be far-fetched as well, since one should try infinitely many policies

to find an optimal policy. Instead, we only consider scenarios in which the σ -compression is *feasible*. In these cases, given \mathcal{M} and Θ , we would like to extract the smallest set of policies Θ' that is a σ -compression of Θ in \mathcal{M} , and then keep this reduced policy space to address any RL task one can define over \mathcal{M} . Let $\Omega_\Theta := \{d_\theta^{sa} \mid \forall \theta \in \Theta\}$ be the set of state-action distributions induced by the policy space Θ , the compression problem can be formulated as a typical *set cover problem*, i.e.,

$$\begin{aligned} & \text{minimize} && \sum_{\omega \in \Omega_\Theta} x_\omega \\ & \text{subject to} && \sum_{\omega: D_2(v \parallel \omega) \leq \sigma} x_\omega \geq 1, \quad \forall v \in \Omega_\Theta \\ & && x_\omega \in \{0, 1\}, \quad \forall \omega \in \Omega_\Theta \end{aligned} \quad (1)$$

where the positive integers x_ω denote the state-action distributions that are active in the covering, and the corresponding σ -compression of Θ in \mathcal{M} can be retrieved as $\Theta_\sigma = \{\theta \in \Theta \mid d_\theta^{sa} = \omega \wedge x_\omega = 1\}$. Unfortunately, the problem (1) is known to be NP-hard (Feige, 1998), even when the model of \mathcal{M} is fully available. Two aspects arguably make this problem extremely hard: On the one hand, we are looking for an efficient solution in the number of active state-action distributions, secondly, we are covering the set Ω_Θ all at once rather than incrementally. Instead of considering common relaxations of (1) (Johnson, 1974; Lovász, 1975), which would not strictly meet the requirements of Definition 3.1 (Feige, 1998), in the next section we build on these insights to reformulate the policy space compression problem in a tractable way.

4 A GAME THEORETIC REFORMULATION

Due to its inherent hardness, we aim to find a tractable reformulation of the policy space compression problem (1) whose solution is a valid σ -compression of Θ in \mathcal{M} . Let us consider a game-theoretic perspective to the set cover problem. A first player distributes a set of K policies $(\theta_1, \dots, \theta_K) \in \Theta^K$ with the intention of covering the set of state-action distributions Ω_Θ . A second player tries to find a policy $\mu \in \Theta$ that is not well covered by $(\theta_1, \dots, \theta_K)$, i.e., a policy that maximizes the Rényi divergence between its state-action distribution and the one of the closest $\theta_k \in (\theta_1, \dots, \theta_K)$. The former player moves first, and we call it a *leader*. The latter player makes his move in response to the other player, and it is then called a *follower*. The two-player, zero-sum, sequential game that we have informally described can be represented

as the optimization problem

$$\begin{aligned} & \min_{\theta \in \Theta^K} \max_{\mu \in \Theta} f(\theta, \mu), \\ & f(\theta, \mu) := \min_{k \in [K]} D_2(d_\mu^{sa} \parallel d_{\theta_k}^{sa}), \end{aligned} \quad (2)$$

where $\theta = (\theta_1, \dots, \theta_K)$ and $[K] = \{1, \dots, K\}$. It is straightforward to see that if the σ -compression is feasible for Θ in \mathcal{M} and K is large enough, then any optimal leader's strategy for the game (2), i.e., $\theta^* \in \arg \min_{\theta \in \Theta^K} \max_{\mu \in \Theta} f(\theta, \mu)$, is a σ -compression of Θ in \mathcal{M} . Unfortunately, $f(\theta, \mu)$ is a non-convex non-concave function, and finding a globally optimal strategy for the game (2) is still a NP-hard problem. However, we do not actually need to find a globally optimal strategy for the leader, as any $\theta \in \Theta^K$ such that $\min_{\mu \in \Theta} f(\theta, \mu) \leq \sigma$ would be a valid σ -compression of Θ . Thus, we might instead target a locally optimal strategy for (2), which is a stationary point of f that is both a local maximum w.r.t. θ and a local minimum w.r.t. μ . We formalize this solution concept as a Differential Stackelberg Equilibrium (DSE) (Fiez et al., 2020).

Definition 4.1 (Differential Stackelberg (Fiez et al., 2020)). *The joint strategy $(\theta^*, \mu^*) \in \Theta^{K+1}$ in which $\theta_k^* \in \arg \min_{k \in [K]} (d_{\mu^*}^{sa} \parallel d_{\theta_k^*}^{sa})$ is a differential Stackelberg equilibrium of the game (2) if it holds $\nabla_{\theta_k^*} f(\theta^*, \mu^*) = 0$, $\nabla_{\mu^*} f(\theta^*, \mu^*) = 0$, $|\nabla_{\theta_k^*} \nabla_{\theta_k^*}^\top f(\theta^*, \mu^*)| > 0$, and $|\nabla_{\mu^*} \nabla_{\mu^*}^\top f(\theta^*, \mu^*)| < 0$.¹*

Luckily, several recent works have established a favorable complexity for the problem of finding a DSE (Jin et al., 2020b; Fiez et al., 2020; Fiez and Ratliff, 2020) in a sequential game. Especially, Jin et al. (2020b) showed that a basic first-order method, i.e., Gradient Descent Ascent (GDA), with an infinite time-scale separation between the leader's and follower's updates is guaranteed to converge to a DSE under mild conditions. This result might be surprising, as we started with a fundamentally hard problem (1) and ended up with a way easier formulation (2) that we can address with a common methodology, without making any strong assumption on the structure of the problem. However, we still have to deal with two crucial issues to solve the policy space compression problem through the game-theoretic formulation. On the one hand, it is not enough to look at the value $f(\theta^*, \mu^*)$ attained by a DSE (θ^*, μ^*) to guarantee that θ is a σ -compression of Θ , as we should check that $\max_{\mu \in \Theta} f(\theta^*, \mu) \leq \sigma$, where μ is a global maximizer. On the other hand, it is not clear how to set a convenient value of K beforehand. In the next section, we present a first-order

¹Let $f(x)$ be a function of $x \in \mathbb{R}^m$, we denote its gradient vector as $\nabla_x f(x)$, its Hessian matrix as $\nabla_x \nabla_x^\top f(x)$, and the determinant of its Hessian matrix as $|\nabla_x \nabla_x^\top f(x)|$.

method that addresses these two issues by finding a DSE of iteratively larger instances of the game (2) (which we will henceforth call the *cover game*) until a conservative approximation of the global condition $\max_{\mu \in \Theta} f(\theta^*, \mu) \leq \sigma$ is finally met.

5 A PLANNING ALGORITHM TO SOLVE THE PROBLEM

Optimization problems of the kind of (2) are typically addressed with a GDA procedure, in which the leader’s parameters (θ) and the follower’s parameters (μ) are updated iteratively according to

$$\theta \leftarrow \theta - \alpha \nabla_{\theta} f(\theta, \mu), \quad \mu \leftarrow \mu + \beta \nabla_{\mu} f(\theta, \mu),$$

where $\nabla_{\theta} f(\theta, \mu)$ and $\nabla_{\mu} f(\theta, \mu)$ are the respective gradients of the joint objective function, $\alpha > 0$ and $\beta > 0$ are learning rates. Especially, if we consider a sufficiently large time-scale separation $\tau := \beta/\alpha$, we are guaranteed to converge to a DSE of the game (2) (Jin et al., 2020b; Fiez and Ratliff, 2020). In this case, we can consider $\tau = \infty$, which means that we update the follower’s parameters until a stationary point is reached, i.e., $\nabla_{\mu} f(\theta, \mu) = 0$, before updating the leader’s parameters. However, to instantiate the cover game, we still need to specify the number K of leader-controlled policies $\theta = (\theta_1, \dots, \theta_K)$. A straightforward solution is to start with a small number of policies first, say $K = 1$, then retrieve a DSE (θ^*, μ^*) via GDA for a cover-game instance with K policies, and finally check if the resulting leader’s strategy θ^* meets the global requirement $\max_{\mu \in \Theta} f(\theta^*, \mu) \leq \sigma$. If the answer is positive, the policy space compression problem is solved, and θ^* is a σ -compression of Θ in \mathcal{M} . Otherwise, we increment K and we repeat the process to see if we can solve the problem with more policies in θ . If the policy space compression problem is feasible, with this simple procedure we are guaranteed to get a valid σ -compression eventually. We call this method the *Policy Space Compression Algorithm* (PSCA) and we report the pseudocode in Algorithm 1. In the following sections, we describe in details how the optimization of the follower’s parameters (Section 5.1) and the leader’s parameters (Section 5.2) are carried out in an adaptation of the GDA method to the specific setting of the cover game. In Section 5.3, we discuss how to verify the global requirement $\max_{\mu \in \Theta} f(\theta^*, \mu) \leq \sigma$ without actually having to find a globally optimal follower’s strategy, but instead optimizing a surrogate objective through a tractable linear program.

Algorithm 1 PSCA

Input: CMP \mathcal{M} , policy space Θ , constant σ
 initialize $K = 0$ and the cover guarantee $\bar{Z}_{\theta} = \infty$
while $(\bar{Z}_{\theta})^2 > \sigma$ **do**
 $K \leftarrow K + 1$
 initialize the leader $\theta = (\theta_1, \dots, \theta_K) \in \Theta^K$
 for epoch = 1, 2, \dots , until convergence **do**
 compute the best response μ_{br} to θ
 identify the active leader’s component θ_k
 update the leader $\theta_k \leftarrow \theta_k - \alpha \nabla_{\theta_k} f(\theta, \mu_{br})$
 end for
 compute the cover guarantee \bar{Z}_{θ} with (6)
end while
Output: return θ , a σ -compression of Θ in \mathcal{M}

5.1 Optimizing the Follower’s Parameters

In principle, we would like to compute the gradient $\nabla_{\mu} f(\theta, \mu)$ to perform the update $\mu \leftarrow \mu + \beta \nabla_{\mu} f(\theta, \mu)$ as in a common GDA procedure. Unfortunately, the objective function $f(\theta, \mu) = \min_{k \in [K]} D_2(d_{\mu}^{sa} || d_{\theta_k}^{sa})$ is not differentiable due to the minimum over the K components of θ . However, only the leader’s component θ_k that attains the minimum of f is actually relevant for the follower’s update, as the other $K - 1$ components do not affect the value of the objective. Thus, we call $\theta_k \in \arg \min_{\theta_i \in \theta} D_2(d_{\mu}^{sa} || d_{\theta_i}^{sa})$ the *active leader’s component*. Conveniently, we can update the follower’s parameters w.r.t. the gradient $\nabla_{\mu} f(\theta_k, \mu)$, which is differentiable w.r.t. μ . The following proposition provides the formula for this gradient.

Proposition 5.1 (Follower’s Gradient). *Let $(\theta, \mu) \in \Theta^K$, the gradient of $f(\theta, \mu)$ w.r.t. μ is given by*

$$\nabla_{\mu} f(\theta, \mu) = 2 \mathbb{E}_{(s,a) \sim d_{\theta_k}^{sa}} \left[\left(\frac{d_{\mu}^{sa}(s,a)}{d_{\theta_k}^{sa}(s,a)} \right)^2 \nabla_{\mu} \log d_{\mu}^{sa}(s,a) \right], \quad (3)$$

where θ_k is the active leader’s component such that $\theta_k \in \arg \min_{\theta_i \in \theta} D_2(d_{\mu}^{sa} || d_{\theta_i}^{sa})$.

To perform a full optimization of the follower’s parameters, we just need to repeatedly apply the gradient ascent update with the gradient $\nabla_{\mu} f(\theta, \mu)$ computed as in (3). Under mild conditions on the learning rate (Robbins and Monro, 1951), this process is guaranteed to converge to a stationary point such that $\nabla_{\mu} f(\theta, \mu) = 0$. We call the follower’s parameters μ at this stationary point the *best response* to the leader’s parameter θ , and we denote it as μ_{br} .

5.2 Optimizing the Leader’s Parameters

Whenever the follower converges at the best response μ_{br} to the current leader’s parameters, we would like

to make an update to θ in the direction of the gradient $\nabla_{\theta} f(\theta, \mu)$, i.e., $\theta \leftarrow \theta - \alpha \nabla_{\theta} f(\theta, \mu)$. Just as before, we can pre-compute the active leader’s component $\theta_k \in \arg \min_{\theta_i \in \Theta} D_2(d_{\mu}^{sa} \| d_{\theta_i}^{sa})$ to make an update to θ_k in the direction of the gradient $\nabla_{\theta_k} f(\theta_k, \mu)$, which is differentiable in θ_k . Indeed, an update to any other leader’s component would not have a meaningful impact on the value of the objective, whereas updating θ_k with a sufficiently small learning rate α is guaranteed to decrease $f(\theta, \mu)$, possibly forcing the follower to change its best response in the next epoch. The following proposition provides the formula for the gradient.

Proposition 5.2 (Leader’s Gradient). *Let $(\theta, \mu) \in \Theta^K$, the gradient of $f(\theta, \mu)$ w.r.t. θ_k is given by*

$$\nabla_{\theta_k} f(\theta, \mu) = - \mathbb{E}_{(s,a) \sim d_{\theta_k}^{sa}} \left[\left(\frac{d_{\mu}^{sa}(s,a)}{d_{\theta_k}^{sa}(s,a)} \right)^2 \nabla_{\theta_k} \log d_{\theta_k}^{sa}(s,a) \right]. \quad (4)$$

5.3 Assessing the Global Value of the Leader’s Parameters

The last missing piece of the PSCA algorithm requires verifying that the leader’s strategy in the DSE (θ^*, μ^*) obtained from the GDA procedure is actually a σ -compression of Θ in \mathcal{M} . In principle, we should verify that $\min_{k \in [K]} D_2(\theta_k^*, \mu) \leq \sigma$ for any $\mu \in \Theta$, which is equivalent to controlling if $\max_{\mu \in \Theta} f(\theta^*, \mu) \leq \sigma$. Unfortunately, the follower’s strategy μ^* is only locally optimal. Thus, checking $f(\theta^*, \mu^*) \leq \sigma$ is not sufficient, as the globally optimal follower’s strategy might attain a greater value of f than μ^* . Instead, we should check $\mathcal{Z}_{\theta^*} \leq \sigma$, where \mathcal{Z}_{θ^*} is given by

$$\mathcal{Z}_{\theta^*} = \max_{\omega \in \Omega_{\Theta}} \min_{k \in [K]} \int_{SA} \frac{(\omega(s,a))^2}{d_{\theta_k^*}^{sa}(s,a)} ds da, \quad (5)$$

which can be written as a quadratically constrained quadratic program (see Appendix B.1). It might come as no surprise that solving this problem is NP-hard. Indeed, this is equivalent to the problem (2) with a fixed leader’s strategy θ^* , but the objective $f(\theta^*, \mu)$ is still non-concave w.r.t. μ . Luckily, we can reformulate this NP-hard problem in the surrogate linear program (see Appendix B.2):

$$(\bar{\mathcal{Z}}_{\theta^*})^{-\frac{1}{2}} = \max_{\omega \in \Omega_{\Theta}} \min_{k \in [K]} \int_{SA} \frac{\omega(s,a)}{(d_{\theta_k^*}^{sa}(s,a))^{-\frac{1}{2}}} ds da, \quad (6)$$

where the value $\bar{\mathcal{Z}}_{\theta^*}$ is a conservative approximation of \mathcal{Z}_{θ^*} , as stated in the following theorem.

Theorem 5.3. *The value $\bar{\mathcal{Z}}_{\theta^*}$ is an upper bound to the value \mathcal{Z}_{θ^*} , i.e., $\bar{\mathcal{Z}}_{\theta^*} \geq \mathcal{Z}_{\theta^*}, \forall \theta^* \in \Theta^K$.*

6 GUARANTEES OF RL WITH A COMPRESSED POLICY SPACE

In the previous sections, we have motivated the pursuit of a compression Θ_{σ} of the original policy space Θ in the CMP \mathcal{M} as a way to improve the computation and sample efficiency of solving RL tasks defined upon \mathcal{M} . Since this compression procedure induces a loss, albeit bounded, in the expressive power of the policy space, it is worth investigating the performance guarantees that we have when addressing RL tasks with Θ_{σ} . We first analyze *policy evaluation* (Section 6.1) and then *policy optimization* (Section 6.2). The reported theoretical results mostly combine techniques from (Metelli et al., 2018; Papini et al., 2019).

6.1 Policy Evaluation

In policy evaluation (Sutton and Barto, 2018), we aim to estimate the performance $J(\theta)$ of a target policy $\theta \in \Theta$ through sampled interactions with an MDP $\mathcal{M}^{\mathcal{R}}$. In our case, we can only draw samples with the policies in Θ_{σ} , and we have to provide an off-policy estimate of $J(\theta)$ via importance sampling. Since for any target policy θ we are guaranteed to have a sampling policy $\theta' \in \Theta_{\sigma}$ such that $D_2(d_{\theta}^{sa} \| d_{\theta'}^{sa}) \leq \sigma$, by choosing a convenient sampling policy in Θ_{σ} , we can enjoy the following guarantee on the error we make when evaluating any target policy $\theta \in \Theta$ in any MDP $\mathcal{M}^{\mathcal{R}}$ one can build upon \mathcal{M} .

Theorem 6.1 (Policy Evaluation Error). *Let Θ_{σ} be a σ -compression of Θ in \mathcal{M} , let \mathcal{R} be a reward function for \mathcal{M} uniformly bounded by R_{\max} , let $\theta \in \Theta$ be a target policy, and let $\delta \in (0, 1)$ be a confidence. There exists $\theta' \in \Theta_{\sigma}$ such that, given N i.i.d. samples from $d_{\theta'}^{sa}$,² the error of the importance sampling evaluation of $J(\theta)$ in $\mathcal{M}^{\mathcal{R}}$, i.e.,*

$$\hat{J}_{IS}(\theta/\theta') = \frac{1}{(1-\gamma)N} \sum_{n=1}^N w_{\theta/\theta'}(s_n, a_n) \mathcal{R}(s_n, a_n),$$

is upper bounded with probability at least $1 - \delta$ as

$$|J(\theta) - \hat{J}_{IS}(\theta/\theta')| \leq \frac{R_{\max}}{1-\gamma} \sqrt{\frac{\sigma}{\delta N}}.$$

Notably, given a budget of samples N , a confidence δ , and a requirement on the evaluation error beforehand, we could select a proper σ to build a σ -compression that meets the requirement in any policy evaluation

²One can generate a sample from $d_{\theta'}^{sa}$ by drawing $s_0 \sim \mu$ and then following the policy θ' . At each step t , the state s_t and action a_t are accepted with probability γ , whereas the simulation ends with probability $1 - \gamma$ (Metelli et al., 2021b).

task. However, choosing a sampling policy $\theta' \in \Theta_\sigma$ that is best suited for a given task might be non-trivial. Thus, one can instead take a batch of N_k samples with each policy in Θ_σ , and then perform the policy evaluation via Multiple Importance Sampling (MIS) (Owen, 2013; Papini et al., 2019).

Corollary 6.2. *Let Θ_σ be a σ -compression of Θ in \mathcal{M} such that $|\Theta_\sigma| = K$, let \mathcal{R} be a reward function for \mathcal{M} uniformly bounded by R_{\max} , let $\theta \in \Theta$ be a target policy, and let $\delta \in (0, 1)$ be a confidence. Given N_k i.i.d. samples from each $d_{\theta_k}^{sa}$, $\theta_k \in \Theta_\sigma$, the error of the multiple importance sampling evaluation of $J(\theta)$ in $\mathcal{M}^{\mathcal{R}}$, i.e.,*

$$\widehat{J}_{MIS}(\theta/\theta_1, \dots, \theta_K) = \frac{1}{(1-\gamma)} \sum_{k=1}^K \sum_{n=1}^{N_k} \frac{d_{\theta}^{sa}(s_{n,k}, a_{n,k})}{\sum_{j=1}^K N_j d_{\theta_j}^{sa}(s_{n,k}, a_{n,k})} \mathcal{R}(s_{n,k}, a_{n,k}),$$

is upper bounded with probability at least $1 - \delta$ as

$$|J(\theta) - \widehat{J}_{MIS}(\theta/\theta_1, \dots, \theta_K)| \leq \frac{R_{\max}}{1-\gamma} \sqrt{\frac{D_2(d_{\theta}^{sa} || \Phi)}{\delta N}}$$

where $N = \sum_{k=1}^K N_k$ is the total number of samples and $\Phi = \sum_{k=1}^K \frac{N_k}{N} d_{\theta_k}^{sa}$ is a finite mixture.

Thanks to the result in (Metelli et al., 2020, Theorem 1), in tabular MDPs the evaluation error of the MIS estimator is guaranteed to be lower than the one of the IS estimator of Theorem 6.1 (as long as $N_k \geq N$, where N is the number of samples considered by the IS estimator).

6.2 Policy Optimization

In policy optimization (see Section 2.2), we seek for the policy θ that maximizes $J(\theta)$ within a parametric policy space. In principle, we could look for the policy that maximizes the performance within the σ -compression Θ_σ , which can be found efficiently with the OPTIMIST algorithm (Papini et al., 2019). Especially, in this setting OPTIMIST yields constant regret for tabular MDPs (Metelli et al., 2021a), as the set Θ_σ is finite and it is composed of stochastic policies such that $\forall \theta, \theta' \in \Theta_\sigma, D_2(d_{\theta}^{sa} || d_{\theta'}^{sa}) < \infty$. However, this optimal policy within Θ_σ might be sub-optimal w.r.t. the optimal policy within the original policy space Θ . We can still upper bound this sub-optimality, as reported in the following theorem.

Theorem 6.3 (Policy Optimization in Θ_σ). *Let Θ_σ be a σ -compression of Θ in \mathcal{M} , and let \mathcal{R} be a reward function for \mathcal{M} uniformly bounded by R_{\max} . The policy $\theta_\sigma^* \in \arg \max_{\theta \in \Theta_\sigma} J(\theta)$ is ϵ -optimal for the MDP $\mathcal{M}^{\mathcal{R}}$, where*

$$\epsilon := \left| \max_{\theta \in \Theta} J(\theta) - J(\theta_\sigma^*) \right| \leq \frac{R_{\max}}{1-\gamma} \sqrt{\log \sigma}.$$

Notably, the latter guarantee does not involve any estimation, and the policy θ^* can be obtained in a finite number of interactions. Nonetheless, one can shrink the sub-optimality ϵ , and without deteriorating the sample complexity, by coupling the OPTIMIST algorithm with an additional offline optimization procedure. The idea is to return the policy $\theta \in \Theta$ that maximizes the importance sampling evaluation obtained with the samples from the policies in Θ_σ .

Theorem 6.4 (Off-Policy Optimization in Θ). *Let Θ_σ be a σ -compression of Θ in \mathcal{M} such that $|\Theta_\sigma| = K$, let \mathcal{R} be a reward function for \mathcal{M} uniformly bounded by R_{\max} , and let $\delta \in (0, 1)$ be a confidence. Given N_k samples from each $d_{\theta_k}^{sa}$, $\theta_k \in \Theta_\sigma$, we can recover an ϵ -optimal policy for $\mathcal{M}^{\mathcal{R}}$ as*

$$\left(\cdot, \theta_{IS}^* \right) \in \arg \max_{\theta_k \in \Theta_\sigma, \theta \in \Theta: D_2(d_{\theta}^{sa} || d_{\theta_k}^{sa})} \frac{1}{(1-\gamma)N_k} \sum_{n=1}^{N_k} w_{\theta/\theta_k}(s_n, a_n) \mathcal{R}(s_n, a_n), \quad (7)$$

such that with probability at least $1 - \delta$

$$\epsilon := \left| \max_{\theta \in \Theta} J(\theta) - J(\theta_{IS}^*) \right| \leq \frac{R_{\max}}{1-\gamma} \sqrt{2\sigma/N_k \delta}.$$

Although, contrary to the guarantee in Theorem 6.3, ϵ vanishes with the number of samples in the latter result, solving the offline problem (7) is non-trivial in general, as the policy space Θ is often infinite.

7 NUMERICAL VALIDATION

In this section, we provide a brief numerical validation of the policy space compression problem (Section 7.1) and how it benefits RL (Section 7.2, 7.3). To the purpose of the analysis, we consider the *River Swim* domain (Strehl and Littman, 2008), in which an agent navigates a chain of six states by taking one of two actions: either *swim up*, to move upstream towards the upper states, or *swim down*, to go downstream back to the lower states. Swimming upstream is harder than swimming downstream, thus the action *swim up* fails with a positive probability, such that only a sequence of *swim up* is likely to lead to the final state (an illustration of the corresponding CMP is reported in Figure 1a). In Appendix C, we report further details on the experimental settings, along with some additional results in a *Grid World* environment. We leave as future work a more extensive experimental evaluation of the policy space compression problem beyond toy domains.

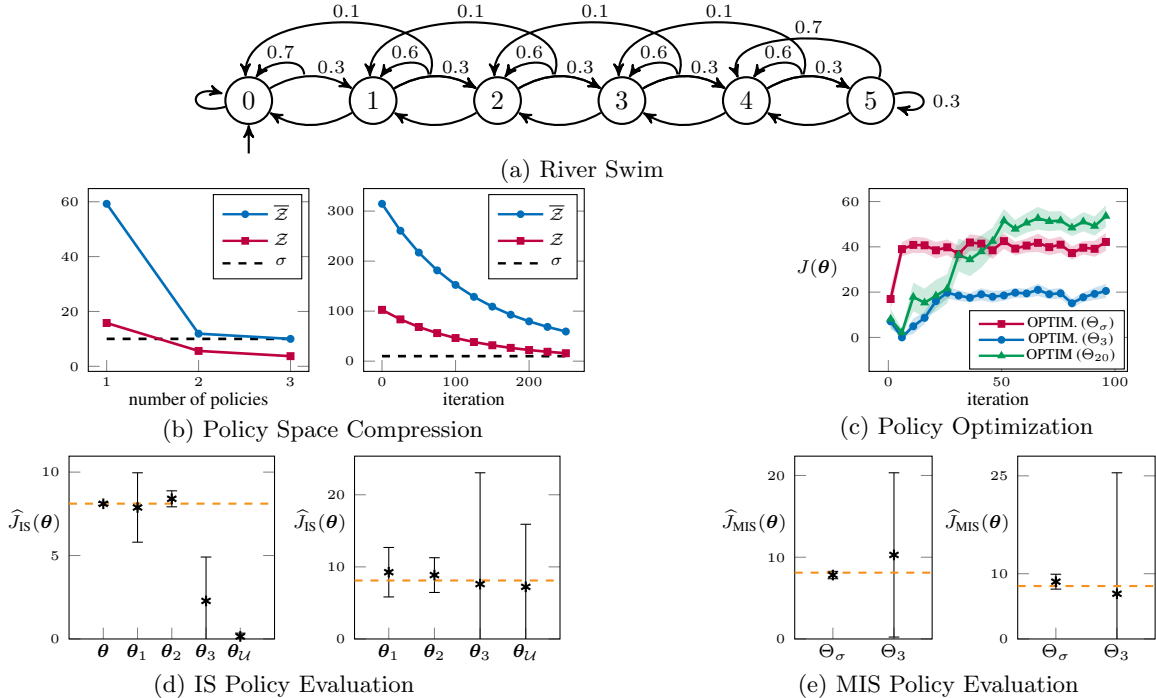


Figure 1: Set of experiments in the *River Swim* domain, which is illustrated in (a). (b) The value of the compression guarantee \mathcal{Z} , its upper bound $\bar{\mathcal{Z}}$, and the requirement σ as a function of the number of policies K (left) and as a function of the iterations with $K = 1$ (right) obtained with PSCA. (c) The average return $J(\theta)$ obtained by OPTIMIST with the σ -compression Θ_σ (3 policies), a 3-policies discretization Θ_3 , and a 20-policies discretization Θ_{20} (95% c.i. over 50 runs). (d,e) IS and MIS evaluation of $J(\theta)$ by taking samples with θ , $\theta_k \in \Theta_\sigma$, a uniform policy θ_U , the mixture Θ_σ , or a mixture of 3 random policies Θ_3 . We provide both the empirical (left, 95% c.i. over 50 runs) and the hindsight (right) values.

7.1 Policy Space Compression

In the River Swim, we consider the policy space $\Theta \subseteq \mathbb{R}^{|\mathcal{S}| \times (|\mathcal{A}|-1)}$ of the *softmax* policies $\pi_\theta(a|s) = \exp(\theta_{sa}) / \sum_{j \in \mathcal{A}} \exp(\theta_{sj})$, and we seek for a compression Θ_σ with the requirement $\sigma = 10$, such that Θ_σ is a valid σ -compression if $\min_{\theta \in \Theta_\sigma} \max_{\mu \in \Theta} f(\theta, \mu) \leq 10$. In Figure 1b, we report the values of $\mathcal{Z} = \max_{\mu \in \Theta} f(\theta, \mu)$ (5) and its upper bound $\bar{\mathcal{Z}} \geq \mathcal{Z}$ (6). Especially, we can see that PSCA effectively found a valid σ -compression Θ_σ of just $K = 3$ policies (Figure 1b, left), and that the values of \mathcal{Z} and $\bar{\mathcal{Z}}$ smoothly decreases during the GDA procedure for a fixed number of policies (Figure 1b, right). Notably, $K = 2$ policies are actually sufficient to meet the σ requirement in this setting. However, PSCA cannot access \mathcal{Z} but its conservative approximation $\bar{\mathcal{Z}}$, and thus stops whenever $\bar{\mathcal{Z}} \leq \sigma$. In Appendix C, we report an illustration of the obtained policies $\theta_k \in \Theta_\sigma$. This set coarsely includes two policies that swims up most of the time, either mixing the actions when the rightmost state is reached (θ_1) or swimming up there as well (θ_2), and a policy that swims down in the leftmost state and swims up in the others (θ_3).

7.2 Policy Evaluation with a Compressed Policy Space

We now show that the obtained σ -compression Θ_σ can be employed with benefit in the most challenging policy evaluation task one can define in the River Swim, which is the off-policy evaluation of an ϵ -greedy policy θ for the reward function that assigns $R_{\max} = 100$ for taking the action *swim up* in the rightmost state. In Figure 1d, we show that sampling with the policies $\theta_1, \theta_2 \in \Theta_\sigma$ lead to an IS off-policy evaluation that is comparable to the exact $J(\theta)$ (dashed line) and its on-policy estimate (θ). Instead, the policy θ_3 and a uniform policy θ_U lead to significantly worse evaluations, as they collect too many samples in the leftmost state. Even by sampling from a uniform mixture of the policies in Θ_σ , the performance of the MIS evaluation is significantly better than the one obtained by a uniform mixture of three random policies (Θ_3), as reported in Figure 1e. For both the IS and the MIS regime, we provide the empirical evaluations (on the left) and the hindsight evaluations (right) obtained with the exact values of the importance weights $w_{\theta/\theta'}$ and the confidence bounds of the Theorem 6.1, 6.2 respectively.

7.3 Policy Optimization with a Compressed Policy Space

Finally, we show that the compression Θ_σ allows for efficient policy optimization. We consider the same reward function of the previous section, and the OPTIMIST (Papini et al., 2019) algorithm equipped with Θ_σ , or a uniform discretization of the original policy space Θ with either three policies (Θ_3) or twenty policies (Θ_{20}). In Figure 1c, we show that OPTIMIST with Θ_σ swiftly converges (less than five iterations) to the optimal policy within the space. Instead, the policy space Θ_3 leads to a huge sub-optimality in the final performance, and OPTIMIST with Θ_{20} is way slower to converge to the optimal policy within the space. These results are a testament of the ability of PSCA to incorporate the peculiar structure of the domain in a small set of representative policies Θ_σ , and to allow for a remarkable balance between sample efficiency and sub-optimality in subsequent policy optimization.

8 DISCUSSION AND CONCLUSION

In this paper, we considered the problem of compressing an infinite parametric policy space into a finite set of representative policies for a given environment. First, we provided a formal definition of the problem, and we highlighted its inherent hardness. Then, we proposed a tractable game-theoretic reformulation, for which a locally optimal solution can be efficiently found through an iterative GDA procedure. Finally, we provided a theoretical characterization of the guarantees that the compression brings to subsequent RL tasks, and a numerical validation of the approach.

8.1 Related Works

Previous works (Gregor et al., 2016; Eysenbach et al., 2018; Achiam et al., 2018; Hansen et al., 2019) have considered heuristic methods to extract a convenient set of policies from the policy space, but they lack the formalization and the theoretical guarantees that we provided. Especially, Eysenbach et al. (2021) argue that the set of policies learned by those methods cannot be used to solve all the relevant policy optimization tasks. Those policies should be generally intended as effective initializations for subsequent adaptation procedures, operating in the original policy space once the task is revealed, rather than a minimal set of sufficient policies. To the best of our knowledge, the only other work considering a formal criterion to operate a selection of the policies is (Zahavy et al., 2021). Having some similarities, our work and (Zahavy et al., 2021) still differ for some crucial aspects. Whereas they look

for a set of policies that maximizes the performance under the worst-case reward, we look for a set of policies that guarantees ϵ -optimality for any task. They do not consider the parameterization of the policy space as an additional source of structure, and thus they do not fully exploit the interplay between the policy space and the environment as we do. Their problem formulation is multi-task, as they restrict the class of rewards to linear combinations of a feature vector, our formulation is instead fully reward-free. Overall, our policy space compression problem is more general, as it is solving the problem in (Zahavy et al., 2021) as a by-product. However, their problem might be easier in nature,³ and thus preferable if one only cares about the worst-case performance. Finally, Eysenbach et al. (2021) provide interesting insights on the information geometry of the space of the state distributions induced by a policy in a CMP, which can lead to compelling geometric interpretations of our policy space compression problem.

8.2 Limitations and Future Directions

The main limitation of our work is that the proposed algorithm is assuming full knowledge of the environment, which is uncommon in RL literature. However, we believe that PSCA is providing a clear blueprint for future works that might target the compression problem from interactions with an *unknown* environment, to pave the way for scalable policy space compression. Especially, such an extension would require sample-based estimates of the gradients (3), (4), and the global guarantee (6). Whereas estimating the gradients of state-action distributions is not an easy feat, previous works provide useful inspiration (Morimura et al., 2010; Schroecker and Isbell, 2017; Schroecker et al., 2018). Similarly, sample-based estimates of (6) can take inspiration from approximate linear programming methods for MDPs (De Farias and Van Roy, 2003; Pazis and Parr, 2011). Another potential limitation of the proposed approach is the memory complexity required to store the compression, in contrast to the compact representations of common policy spaces, such as a small set of basis functions or a neural network architecture. A future work might focus on compact representations for a given compression. Other interesting future directions include an extension of the policy space compression problem to the parameter-based perspective (Sehnke et al., 2008; Metelli et al., 2018; Papini et al., 2019), and the development of policy optimization algorithms that are tailored to exploit a compression of the policy space.

³This is purely speculative as (Zahavy et al., 2021) does not provide a formal study of the computational complexity of the problem.

References

- Achiam, J., Edwards, H., Amodei, D., and Abbeel, P. (2018). Variational option discovery algorithms. *arXiv preprint arXiv:1807.10299*.
- Cai, Q., Yang, Z., Jin, C., and Wang, Z. (2020). Provably efficient exploration in policy optimization. In *Proceedings of the International Conference on Machine Learning*.
- Cochran, W. G. (2007). *Sampling techniques*. John Wiley & Sons.
- Cortes, C., Mansour, Y., and Mohri, M. (2010). Learning bounds for importance weighting. In *Advances in Neural Information Processing Systems*.
- De Farias, D. P. and Van Roy, B. (2003). The linear programming approach to approximate dynamic programming. *Operations research*, 51(6):850–865.
- Deisenroth, M. P., Neumann, G., Peters, J., et al. (2013). A survey on policy search for robotics. *Foundations and trends in Robotics*, 2(1-2):388–403.
- Eysenbach, B., Gupta, A., Ibarz, J., and Levine, S. (2018). Diversity is all you need: Learning skills without a reward function. In *International Conference on Learning Representations*.
- Eysenbach, B., Salakhutdinov, R., and Levine, S. (2021). The information geometry of unsupervised reinforcement learning. *arXiv preprint arXiv:2110.02719*.
- Feige, U. (1998). A threshold of $\ln n$ for approximating set cover. *Journal of the ACM (JACM)*.
- Fiez, T., Chasnov, B., and Ratliff, L. (2020). Implicit learning dynamics in stackelberg games: Equilibria characterization, convergence analysis, and empirical study. In *Proceedings of the International Conference on Machine Learning*.
- Fiez, T. and Ratliff, L. J. (2020). Local convergence analysis of gradient descent ascent with finite timescale separation. In *International Conference on Learning Representations*.
- Gregor, K., Rezende, D. J., and Wierstra, D. (2016). Variational intrinsic control. *arXiv preprint arXiv:1611.07507*.
- Hansen, S., Dabney, W., Barreto, A., Warde-Farley, D., Van de Wiele, T., and Mnih, V. (2019). Fast task inference with variational intrinsic successor features. In *International Conference on Learning Representations*.
- Hazan, E., Kakade, S., Singh, K., and Van Soest, A. (2019). Provably efficient maximum entropy exploration. In *Proceedings of the International Conference on Machine Learning*.
- Jin, C., Krishnamurthy, A., Simchowitz, M., and Yu, T. (2020a). Reward-free exploration for reinforcement learning. In *Proceedings of the International Conference on Machine Learning*.
- Jin, C., Netrapalli, P., and Jordan, M. (2020b). What is local optimality in nonconvex-nonconcave minimax optimization? In *Proceedings of the International Conference on Machine Learning*.
- Johnson, D. S. (1974). Approximation algorithms for combinatorial problems. *Journal of computer and system sciences*.
- Lovász, L. (1975). On the ratio of optimal integral and fractional covers. *Discrete mathematics*.
- Metelli, A. M., Papini, M., D’Oro, P., and Restelli, M. (2021a). Policy optimization as online learning with mediator feedback. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Metelli, A. M., Papini, M., Faccio, F., and Restelli, M. (2018). Policy optimization via importance sampling. *Advances in Neural Information Processing Systems*.
- Metelli, A. M., Papini, M., Montali, N., and Restelli, M. (2020). Importance sampling techniques for policy optimization. *Journal of Machine Learning Research*, 21(141):1–75.
- Metelli, A. M., Pirotta, M., Calandriello, D., and Restelli, M. (2021b). Safe policy iteration: A monotonically improving approximate policy iteration approach. *Journal of Machine Learning Research*, 22(97):1–83.
- Morimura, T., Uchibe, E., Yoshimoto, J., Peters, J., and Doya, K. (2010). Derivatives of logarithmic stationary distributions for policy gradient reinforcement learning. *Neural computation*.
- Owen, A. B. (2013). *Monte Carlo theory, methods and examples*.
- Papini, M., Metelli, A. M., Lupo, L., and Restelli, M. (2019). Optimistic policy optimization via multiple importance sampling. In *Proceedings of the International Conference on Machine Learning*.
- Paulin, D. (2015). Concentration inequalities for markov chains by marton couplings and spectral methods. *Electronic Journal of Probability*, 20:1–32.
- Pazis, J. and Parr, R. (2011). Non-parametric approximate linear programming for mdps. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Puterman, M. L. (2014). *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons.
- Rényi, A. et al. (1961). On measures of entropy and information. In *Proceedings of the Fourth Berkeley*

Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics. The Regents of the University of California.

- Robbins, H. and Monro, S. (1951). A stochastic approximation method. *The annals of mathematical statistics.*
- Schroecker, Y. and Isbell, C. L. (2017). State aware imitation learning. In *Advances in Neural Information Processing Systems.*
- Schroecker, Y., Vecerik, M., and Scholz, J. (2018). Generative predecessor models for sample-efficient imitation learning. In *International Conference on Learning Representations.*
- Schulman, J., Levine, S., Abbeel, P., Jordan, M., and Moritz, P. (2015). Trust region policy optimization. In *Proceedings of the International Conference on Machine Learning.*
- Sehnke, F., Osendorfer, C., Rückstieß, T., Graves, A., Peters, J., and Schmidhuber, J. (2008). Policy gradients with parameter-based exploration for control. In *International Conference on Artificial Neural Networks.*
- Silver, D., Lever, G., Heess, N., Degris, T., Wierstra, D., and Riedmiller, M. (2014). Deterministic policy gradient algorithms. In *Proceedings of the International Conference on Machine Learning.*
- Sobel, M. J. (1982). The variance of discounted markov decision processes. *Journal of Applied Probability.*
- Strehl, A. L. and Littman, M. L. (2008). An analysis of model-based interval estimation for markov decision processes. *Journal of Computer and System Sciences.*
- Sutton, R. S. and Barto, A. G. (2018). *Reinforcement learning: An introduction.* MIT press.
- Sutton, R. S., McAllester, D. A., Singh, S. P., Mansour, Y., et al. (1999). Policy gradient methods for reinforcement learning with function approximation. In *Advances in Neural Information Processing Systems.*
- Xie, T., Ma, Y., and Wang, Y.-X. (2019). Towards optimal off-policy evaluation for reinforcement learning with marginalized importance sampling. In *Advances in Neural Information Processing Systems.*
- Zahavy, T., Barreto, A., Mankowitz, D. J., Hou, S., O’Donoghue, B., Kemaev, I., and Singh, S. (2021). Discovering a set of policies for the worst case reward. In *International Conference on Learning Representations.*

A Proofs

A.1 Proofs of Section 5

Proposition 5.1 (Follower’s Gradient). *Let $(\boldsymbol{\theta}, \boldsymbol{\mu}) \in \Theta^K$, the gradient of $f(\boldsymbol{\theta}, \boldsymbol{\mu})$ w.r.t. $\boldsymbol{\mu}$ is given by*

$$\begin{aligned} \nabla_{\boldsymbol{\mu}} f(\boldsymbol{\theta}, \boldsymbol{\mu}) &= \\ & 2 \mathbb{E}_{(s,a) \sim d_{\boldsymbol{\theta}_k}^{sa}} \left[\left(\frac{d_{\boldsymbol{\mu}}^{sa}(s,a)}{d_{\boldsymbol{\theta}_k}^{sa}(s,a)} \right)^2 \nabla_{\boldsymbol{\mu}} \log d_{\boldsymbol{\mu}}^{sa}(s,a) \right], \end{aligned} \quad (3)$$

where $\boldsymbol{\theta}_k$ is the active leader’s component such that $\boldsymbol{\theta}_k \in \arg \min_{\boldsymbol{\theta}_i \in \boldsymbol{\theta}} D_2(d_{\boldsymbol{\mu}}^{sa} || d_{\boldsymbol{\theta}_i}^{sa})$.

Proof. Let $\boldsymbol{\theta}_k$ be the active leader’s component, i.e., $\boldsymbol{\theta}_k \in \arg \min_{\boldsymbol{\theta}_i \in \boldsymbol{\theta}} D_2(d_{\boldsymbol{\mu}}^{sa} || d_{\boldsymbol{\theta}_i}^{sa})$. We can compute the gradient of the objective $f(\boldsymbol{\theta}, \boldsymbol{\mu})$ w.r.t. $\boldsymbol{\mu}$ as

$$\begin{aligned} \nabla_{\boldsymbol{\mu}} f(\boldsymbol{\theta}, \boldsymbol{\mu}) &= \nabla_{\boldsymbol{\mu}} D_2(d_{\boldsymbol{\mu}}^{sa} || d_{\boldsymbol{\theta}_k}^{sa}) \\ &= \nabla_{\boldsymbol{\mu}} \int_{\mathcal{S}\mathcal{A}} d_{\boldsymbol{\theta}_k}^{sa}(s,a) \left(\frac{d_{\boldsymbol{\mu}}^{sa}(s,a)}{d_{\boldsymbol{\theta}_k}^{sa}(s,a)} \right)^2 ds da \\ &= 2 \int_{\mathcal{S}\mathcal{A}} d_{\boldsymbol{\theta}_k}^{sa}(s,a) \left(\frac{d_{\boldsymbol{\mu}}^{sa}(s,a)}{d_{\boldsymbol{\theta}_k}^{sa}(s,a)} \right)^2 \nabla_{\boldsymbol{\mu}} \log d_{\boldsymbol{\mu}}^{sa}(s,a) ds da. \end{aligned}$$

□

Proposition 5.2 (Leader’s Gradient). *Let $(\boldsymbol{\theta}, \boldsymbol{\mu}) \in \Theta^K$, the gradient of $f(\boldsymbol{\theta}, \boldsymbol{\mu})$ w.r.t. $\boldsymbol{\theta}_k$ is given by*

$$\begin{aligned} \nabla_{\boldsymbol{\theta}_k} f(\boldsymbol{\theta}, \boldsymbol{\mu}) &= \\ & - \mathbb{E}_{(s,a) \sim d_{\boldsymbol{\theta}_k}^{sa}} \left[\left(\frac{d_{\boldsymbol{\mu}}^{sa}(s,a)}{d_{\boldsymbol{\theta}_k}^{sa}(s,a)} \right)^2 \nabla_{\boldsymbol{\theta}_k} \log d_{\boldsymbol{\theta}_k}^{sa}(s,a) \right]. \end{aligned} \quad (4)$$

Proof. We can compute the gradient of the objective $f(\boldsymbol{\theta}, \boldsymbol{\mu})$ w.r.t. $\boldsymbol{\theta}_k \in \boldsymbol{\theta}$ as

$$\begin{aligned} \nabla_{\boldsymbol{\theta}_k} f(\boldsymbol{\theta}, \boldsymbol{\mu}) &= \nabla_{\boldsymbol{\theta}_k} D_2(d_{\boldsymbol{\mu}}^{sa} || d_{\boldsymbol{\theta}_k}^{sa}) \\ &= \nabla_{\boldsymbol{\theta}_k} \int_{\mathcal{S}\mathcal{A}} d_{\boldsymbol{\theta}_k}^{sa}(s,a) \left(\frac{d_{\boldsymbol{\mu}}^{sa}(s,a)}{d_{\boldsymbol{\theta}_k}^{sa}(s,a)} \right)^2 ds da \\ &= - \int_{\mathcal{S}\mathcal{A}} d_{\boldsymbol{\theta}_k}^{sa}(s,a) \left(\frac{d_{\boldsymbol{\mu}}^{sa}(s,a)}{d_{\boldsymbol{\theta}_k}^{sa}(s,a)} \right)^2 \nabla_{\boldsymbol{\theta}_k} \log d_{\boldsymbol{\theta}_k}^{sa}(s,a) ds da. \end{aligned}$$

□

Theorem 5.3. *The value $\bar{\mathcal{Z}}_{\boldsymbol{\theta}^*}$ is an upper bound to the value $\mathcal{Z}_{\boldsymbol{\theta}^*}$, i.e., $\bar{\mathcal{Z}}_{\boldsymbol{\theta}^*} \geq \mathcal{Z}_{\boldsymbol{\theta}^*}, \forall \boldsymbol{\theta}^* \in \Theta^K$.*

Proof. The result is straightforward from

$$\begin{aligned} \bar{\mathcal{Z}}_{\boldsymbol{\theta}^*} &= ((\bar{\mathcal{Z}}_{\boldsymbol{\theta}^*})^{-\frac{1}{2}})^2 = \max_{\omega \in \Omega_{\boldsymbol{\theta}}} \min_{k \in [K]} \left(\int_{\mathcal{S}\mathcal{A}} \omega(s,a) (d_{\boldsymbol{\theta}_k}^{sa}(s,a))^{-\frac{1}{2}} ds da \right)^2 \\ &\geq \max_{\omega \in \Omega_{\boldsymbol{\theta}}} \min_{k \in [K]} \int_{\mathcal{S}\mathcal{A}} \left(\omega(s,a) (d_{\boldsymbol{\theta}_k}^{sa}(s,a))^{-\frac{1}{2}} \right)^2 ds da = \mathcal{Z}_{\boldsymbol{\theta}^*}. \end{aligned}$$

□

A.2 Proofs of Section 6

Lemma A.1 (Variance of the IS Estimator). *Let \mathcal{M} be a CMP, and let $\theta \in \Theta$ be a target policy. Let $\{s_n, a_n\}_{n=1}^N$ be a sample of state-action pairs taken with the policy θ' in \mathcal{M} . Then, the variance of the importance sampling evaluation of $J(\theta)$ in \mathcal{M} , i.e., $\widehat{J}_{IS}(\theta/\theta') = \frac{1}{(1-\gamma)N} \sum_{n=1}^N w_{\theta/\theta'}(s_n, a_n) \mathcal{R}(s_n, a_n)$, can be upper bounded as*

$$\mathbb{V}\text{ar}_{(s,a) \sim d_{\theta'}^{sa}} [\widehat{J}_{IS}(\theta/\theta')] \leq \frac{(\mathbf{R}_{\max})^2 D_2(d_{\theta}^{sa} \| d_{\theta'}^{sa})}{(1-\gamma)^2 N}.$$

Proof. The proof follows the derivation in (Metelli et al., 2018, Lemma 4.1). When considering state-action pairs (as opposed to trajectories in (Metelli et al., 2018)) one should account for the dependency between state-actions in the same trajectory. Here we consider a batch of N i.i.d. samples taken with the discounted state distribution $d_{\theta'}^{sa}$, in which the dependency vanishes. Especially, we write

$$\begin{aligned} \mathbb{V}\text{ar}_{(s,a) \sim d_{\theta'}^{sa}} [\widehat{J}_{IS}(\theta/\theta')] &\leq \frac{1}{(1-\gamma)^2 N} \mathbb{V}\text{ar}_{(s,a) \sim d_{\theta'}^{sa}} [w_{\theta/\theta'}(s, a) \mathcal{R}(s, a)] \\ &\leq \frac{1}{(1-\gamma)^2 N} \mathbb{E}_{(s,a) \sim d_{\theta'}^{sa}} \left[\left(\frac{d_{\theta}^{sa}(s, a)}{d_{\theta'}^{sa}(s, a)} \mathcal{R}(s, a) \right)^2 \right] \\ &\leq \frac{(\mathbf{R}_{\max})^2}{(1-\gamma)^2 N} \mathbb{E}_{(s,a) \sim d_{\theta'}^{sa}} \left[\left(\frac{d_{\theta}^{sa}(s, a)}{d_{\theta'}^{sa}(s, a)} \right)^2 \right] = \frac{(\mathbf{R}_{\max})^2 D_2(d_{\theta}^{sa} \| d_{\theta'}^{sa})}{(1-\gamma)^2 N}. \end{aligned}$$

Note that the sampling procedure from a discounted state distribution is wasteful, as one should draw exactly N trajectories from the discounted CMP \mathcal{M} to collect just N i.i.d. samples, while the other samples in the trajectories are discarded (Metelli et al., 2021b). Nonetheless, one could refine this result to account for dependent data, by either exploiting the Bellman equation of the variance (see Sobel, 1982; Xie et al., 2019) or concentration inequalities for Markov chains (Paulin, 2015), which allows to upper bound the variance of the estimate computed over N dependent samples from $d_{\theta'}^{sa}$. \square

Theorem 6.1 (Policy Evaluation Error). *Let Θ_{σ} be a σ -compression of Θ in \mathcal{M} , let \mathcal{R} be a reward function for \mathcal{M} uniformly bounded by \mathbf{R}_{\max} , let $\theta \in \Theta$ be a target policy, and let $\delta \in (0, 1)$ be a confidence. There exists $\theta' \in \Theta_{\sigma}$ such that, given N i.i.d. samples from $d_{\theta'}^{sa}$,⁴ the error of the importance sampling evaluation of $J(\theta)$ in $\mathcal{M}^{\mathcal{R}}$, i.e.,*

$$\widehat{J}_{IS}(\theta/\theta') = \frac{1}{(1-\gamma)N} \sum_{n=1}^N w_{\theta/\theta'}(s_n, a_n) \mathcal{R}(s_n, a_n),$$

is upper bounded with probability at least $1 - \delta$ as

$$|J(\theta) - \widehat{J}_{IS}(\theta/\theta')| \leq \frac{\mathbf{R}_{\max}}{1-\gamma} \sqrt{\frac{\sigma}{\delta N}}.$$

Proof. We would like to bound the difference $|J(\theta) - \widehat{J}_{IS}(\theta/\theta')|$ for a policy $\theta' \in \Theta_{\sigma}$. By the definition of σ -compression, there exists at least a policy $\theta' \in \Theta_{\sigma}$ such that $D_2(d_{\theta}^{sa} \| d_{\theta'}^{sa}) \leq \sigma$. Since the IS estimator $\widehat{J}_{IS}(\theta/\theta')$ is unbiased, and $\mathbb{V}\text{ar}_{(s,a) \sim d_{\theta'}^{sa}} [\widehat{J}_{IS}(\theta/\theta')] < \infty$ through Lemma A.1, we can use the Chebichev's inequality to write, $\forall \epsilon > 0$,

$$\Pr(|J(\theta) - \widehat{J}_{IS}(\theta/\theta')| \geq \epsilon) \leq \frac{\mathbb{V}\text{ar}_{(s,a) \sim d_{\theta'}^{sa}} [\widehat{J}_{IS}(\theta/\theta')]}{\epsilon^2}.$$

Then, by calling $\delta = \frac{\mathbb{V}\text{ar}_{(s,a) \sim d_{\theta'}^{sa}} [\widehat{J}_{IS}(\theta/\theta')]}{\epsilon^2}$ and considering the complimentary event, we get

$$\Pr\left(|J(\theta) - \widehat{J}_{IS}(\theta/\theta')| \leq \frac{\mathbf{R}_{\max}}{1-\gamma} \sqrt{\sigma/\delta N}\right) \geq 1 - \delta$$

where we upper bounded the variance of $\widehat{J}_{IS}(\theta/\theta')$ as in Lemma A.1 and the Rényi $D_2(d_{\theta}^{sa} \| d_{\theta'}^{sa})$ with σ . \square

⁴One can generate a sample from $d_{\theta'}^{sa}$ by drawing $s_0 \sim \mu$ and then following the policy θ' . At each step t , the state s_t and action a_t are accepted with probability γ , whereas the simulation ends with probability $1 - \gamma$ (Metelli et al., 2021b).

Corollary 6.2. Let Θ_σ be a σ -compression of Θ in \mathcal{M} such that $|\Theta_\sigma| = K$, let \mathcal{R} be a reward function for \mathcal{M} uniformly bounded by R_{\max} , let $\boldsymbol{\theta} \in \Theta$ be a target policy, and let $\delta \in (0, 1)$ be a confidence. Given N_k i.i.d. samples from each $d_{\boldsymbol{\theta}_k}^{sa}$, $\boldsymbol{\theta}_k \in \Theta_\sigma$, the error of the multiple importance sampling evaluation of $J(\boldsymbol{\theta})$ in $\mathcal{M}^{\mathcal{R}}$, i.e.,

$$\begin{aligned} \widehat{J}_{MIS}(\boldsymbol{\theta}/\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_K) &= \\ \frac{1}{(1-\gamma)} \sum_{k=1}^K \sum_{n=1}^{N_k} \frac{d_{\boldsymbol{\theta}}^{sa}(s_{n,k}, a_{n,k})}{\sum_{j=1}^K N_j d_{\boldsymbol{\theta}_j}^{sa}(s_{n,k}, a_{n,k})} \mathcal{R}(s_{n,k}, a_{n,k}), \end{aligned}$$

is upper bounded with probability at least $1 - \delta$ as

$$|J(\boldsymbol{\theta}) - \widehat{J}_{MIS}(\boldsymbol{\theta}/\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_K)| \leq \frac{R_{\max}}{1-\gamma} \sqrt{\frac{D_2(d_{\boldsymbol{\theta}}^{sa} || \Phi)}{\delta N}}$$

where $N = \sum_{k=1}^K N_k$ is the total number of samples and $\Phi = \sum_{k=1}^K \frac{N_k}{N} d_{\boldsymbol{\theta}_k}^{sa}$ is a finite mixture.

Proof. Through the combination of (Papini et al., 2019, Lemma 1) and Lemma A.1, it is straightforward to derive

$$\text{Var}_{(s,a) \sim d_{\boldsymbol{\theta}_k}^{sa}} [\widehat{J}_{MIS}(\boldsymbol{\theta}/\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_K)] \leq \frac{(R_{\max})^2 D_2(d_{\boldsymbol{\theta}}^{sa} || \Phi)}{(1-\gamma)^2 N}. \quad (8)$$

Then, similarly as in Theorem 6.1, we can use the Chebichev's inequality to write, $\forall \epsilon > 0$,

$$\Pr(|J(\boldsymbol{\theta}) - \widehat{J}_{MIS}(\boldsymbol{\theta}/\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_K)| \geq \epsilon) \leq \frac{\text{Var}_{(s,a) \sim d_{\boldsymbol{\theta}_k}^{sa}} [\widehat{J}_{MIS}(\boldsymbol{\theta}/\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_K)]}{\epsilon^2}.$$

By calling $\delta = \frac{\text{Var}_{(s,a) \sim d_{\boldsymbol{\theta}_k}^{sa}} [\widehat{J}_{MIS}(\boldsymbol{\theta}/\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_K)]}{\epsilon^2}$ and considering the complimentary event, we get

$$\Pr\left(|J(\boldsymbol{\theta}) - \widehat{J}_{MIS}(\boldsymbol{\theta}/\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_K)| \leq \frac{R_{\max}}{1-\gamma} \sqrt{\frac{D_2(d_{\boldsymbol{\theta}}^{sa} || \Phi)}{\delta N}}\right) \geq 1 - \delta$$

where we upper bounded the variance of $\widehat{J}_{MIS}(\boldsymbol{\theta}/\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_K)$ as in (8). \square

Theorem 6.3 (Policy Optimization in Θ_σ). Let Θ_σ be a σ -compression of Θ in \mathcal{M} , and let \mathcal{R} be a reward function for \mathcal{M} uniformly bounded by R_{\max} . The policy $\boldsymbol{\theta}_\sigma^* \in \arg \max_{\boldsymbol{\theta} \in \Theta_\sigma} J(\boldsymbol{\theta})$ is ϵ -optimal for the MDP $\mathcal{M}^{\mathcal{R}}$, where

$$\epsilon := \left| \max_{\boldsymbol{\theta} \in \Theta} J(\boldsymbol{\theta}) - J(\boldsymbol{\theta}_\sigma^*) \right| \leq \frac{R_{\max}}{1-\gamma} \sqrt{\log \sigma}.$$

Proof. Let be $\boldsymbol{\theta}^* \in \arg \max_{\boldsymbol{\theta} \in \Theta} J(\boldsymbol{\theta})$. From the definition of σ -compression we have that there exists at least a policy $\boldsymbol{\theta}' \in \Theta_\sigma$ such that $D_2(d_{\boldsymbol{\theta}^*}^{sa} || d_{\boldsymbol{\theta}'}^{sa}) \leq \sigma$. Then, we can write

$$(1-\gamma) |J(\boldsymbol{\theta}^*) - J(\boldsymbol{\theta}')| = \left| \int_{\mathcal{S}\mathcal{A}} \mathcal{R}(s, a) (d_{\boldsymbol{\theta}^*}^{sa} - d_{\boldsymbol{\theta}'}^{sa}) ds da \right| \quad (9)$$

$$\leq R_{\max} \int_{\mathcal{S}\mathcal{A}} |d_{\boldsymbol{\theta}^*}^{sa} - d_{\boldsymbol{\theta}'}^{sa}| ds da \quad (10)$$

$$\leq R_{\max} \sqrt{d_{KL}(d_{\boldsymbol{\theta}^*}^{sa} || d_{\boldsymbol{\theta}'}^{sa})} \quad (11)$$

$$\leq R_{\max} \sqrt{\log (D_2(d_{\boldsymbol{\theta}^*}^{sa} || d_{\boldsymbol{\theta}'}^{sa}))} = R_{\max} \sqrt{\log \sigma} \quad (12)$$

where (9) is from the definition of J given in Section 2.2, (11) is obtained from (10) through the Pinsker's inequality, and (12) derives from $d_{KL}(p||q) = d_1(p||q) \leq d_2(p||q) = D_2(p||q)$, which is straightforward from the definition of Rényi divergence. Finally, it is trivial to see that $J(\boldsymbol{\theta}_\sigma^*) \geq J(\boldsymbol{\theta}')$ for $\boldsymbol{\theta}_\sigma^* \in \arg \max_{\boldsymbol{\theta} \in \Theta_\sigma} J(\boldsymbol{\theta})$. \square

Theorem 6.4 (Off-Policy Optimization in Θ). *Let Θ_σ be a σ -compression of Θ in \mathcal{M} such that $|\Theta_\sigma| = K$, let \mathcal{R} be a reward function for \mathcal{M} uniformly bounded by R_{\max} , and let $\delta \in (0, 1)$ be a confidence. Given N_k samples from each $d_{\theta_k}^{sa}$, $\theta_k \in \Theta_\sigma$, we can recover an ϵ -optimal policy for $\mathcal{M}^{\mathcal{R}}$ as*

$$\begin{aligned} (\cdot, \theta_{IS}^*) \in & \arg \max_{\theta_k \in \Theta_\sigma, \theta \in \Theta: D_2(d_{\theta}^{sa} \| d_{\theta_k}^{sa})} \\ & \frac{1}{(1-\gamma)N_k} \sum_{n=1}^{N_k} w_{\theta/\theta_k}(s_n, a_n) \mathcal{R}(s_n, a_n), \end{aligned} \quad (7)$$

such that with probability at least $1 - \delta$

$$\epsilon := \left| \max_{\theta \in \Theta} J(\theta) - J(\theta_{IS}^*) \right| \leq \frac{R_{\max}}{1-\gamma} \sqrt{2\sigma/N_k\delta}.$$

Proof. Thanks to the definition of σ -compression and the guarantee provided by Theorem 6.1, from the collected samples we have that there exists $\theta_k \in \Theta_\sigma$ such that

$$\widehat{J}_{IS}(\theta/\theta_k) - \frac{R_{\max}}{1-\gamma} \sqrt{\frac{2\sigma}{N_k\delta}} \leq J(\theta) \leq \widehat{J}_{IS}(\theta/\theta_k) + \frac{R_{\max}}{1-\gamma} \sqrt{\frac{2\sigma}{N_k\delta}}$$

holds $\forall \theta \in \Theta$ with probability at least $1 - \delta/2$. Then, let θ_{IS}^* be a policy obtained as in (7), and let $\theta^* \in \arg \max_{\theta \in \Theta} J(\theta)$. We consider the event in which $J(\theta_{IS}^*)$ falls below its lower confidence bound and $J(\theta^*)$ exceeds its upper confidence bound. It is easy to see that this event happens with probability at most δ , whereas the complimentary event guarantees that

$$|J(\theta^*) - J(\theta_{IS}^*)| \leq \frac{R_{\max}}{1-\gamma} \sqrt{\frac{2\sigma}{N_k\delta}}.$$

□

B Optimization Problems

B.1 Quadratic Program Formulation of (5)

The optimization problem in (5) can be formulated into a quadratically constrained quadratic program as

$$\begin{aligned} & \underset{z \in \mathbb{R}, \omega \in \mathbb{R}^{\mathcal{S}\mathcal{A}}}{\text{maximize}} && z \\ & \text{subject to} && z - \int_{\mathcal{S}\mathcal{A}} \frac{(\omega(s, a))^2}{d_{\theta_k}^{sa}(s, a)} ds da \leq 0, \quad \forall k \in [K] \\ & && \int_{\mathcal{A}} \omega(s, a) da = (1-\gamma)\mu(s) + \gamma \int_{\mathcal{S}\mathcal{A}} \omega(s', a') P(s|s', a') ds' da', \quad \forall s \in \mathcal{S} \\ & && \omega(s, a) \geq 0, \quad \forall s \in \mathcal{S}, \forall a \in \mathcal{A}. \end{aligned}$$

B.2 Linear Program Formulation of (6)

The optimization problem in (6) can be formulated into a linear program as

$$\begin{aligned} & \underset{z \in \mathbb{R}, \omega \in \mathbb{R}^{\mathcal{S}\mathcal{A}}}{\text{maximize}} && z \\ & \text{subject to} && z - \int_{\mathcal{S}\mathcal{A}} \frac{\omega(s, a)}{d_{\theta_k}^{sa}(s, a)} ds da \leq 0, \quad \forall k \in [K] \\ & && \int_{\mathcal{A}} \omega(s, a) da = (1-\gamma)\mu(s) + \gamma \int_{\mathcal{S}\mathcal{A}} \omega(s', a') P(s|s', a') ds' da', \quad \forall s \in \mathcal{S} \\ & && \omega(s, a) \geq 0, \quad \forall s \in \mathcal{S}, \forall a \in \mathcal{A}. \end{aligned}$$

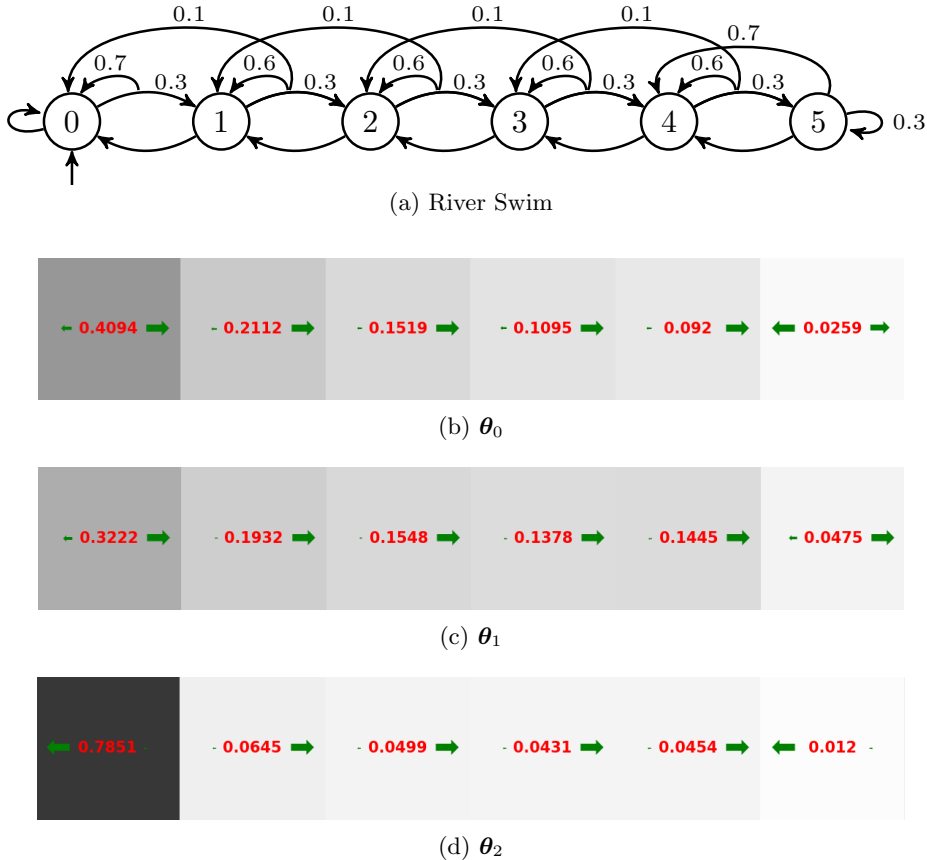


Figure 2: **(a)** Illustration of the River Swim CMP. **(b, c, d)** Heatmap visualization of the policies in the σ -compression $\theta_k \in \Theta_\sigma$ for the River Swim domain. The background color and the label denote the state probability, the green arrows represent the policy in the state.

C Further Details on the Numerical Validation

In Section 7.1, we commented the results of PSCA in the River Swim domain. For the sake of clarity, here we report an illustration of the River Swim CMP (Figure 2a), heatmap visualizations of the policies in the σ -compression obtained by PSCA (Figure 2b-2d), and the set of parameters we employed ($\sigma = 10, \alpha = 0.005, \beta = 0.1$). We further report the results of an additional policy space compression experiment in a Gridworld domain ($|\mathcal{S}| = 9, |\mathcal{A}| = 4$). In this setting, we considered $\sigma = 40, \alpha = 0.005, \beta = 0.1$, and the resulting σ -compression is composed of $K = 4$ policies (a visualization is provided in Figure 3a-3d).

In Section 7.2, we reported a set of policy evaluation experiments in the River Swim domain. Especially, we considered an IS off-policy evaluation setting, in which we take a batch of samples with each policy $\theta_k \in \Theta_\sigma$, or with a uniform policy $\theta_{\mathcal{U}}$, or with the target policy itself θ . For every policy, the batch is composed of $N = 100000$ samples, and it is obtained by drawing 5000 trajectories of 20 steps. Similarly, we considered a MIS off-policy evaluation setting, in which we take a batch of samples with the σ -compression Θ_σ , or a set of three random policies Θ_3 . In both the cases, the batch is composed of $N = 300000$ samples ($N_k = 100000$ for each policy in the space), obtained by drawing 15000 trajectories of 20 steps.

In Section 7.3, we reported a policy optimization experiment in the River Swim domain. To run this experiment, we implemented the action-based formulation of the OPTIMIST algorithm (Papini et al., 2019, Algorithm 1). For each seed, we run the algorithm for 100 iterations, in each iteration we collect $N = 1000$ samples, which are obtained from 50 trajectories of 20 steps. The value of the importance weights truncation M and the confidence schedule δ_t are taken from the theoretical analysis in (Papini et al., 2019).




Figure 3: **(a, b, c, d)** Heatmap visualization of the policies in the σ -compression $\theta_k \in \Theta_\sigma$ for the Gridworld domain. The background color and the label denote the state probability, the green arrows represent the policy in the state.