Control Frequency Adaptation via Action Persistence in Batch Reinforcement Learning

Alberto Maria Metelli [1] **Flavio Mazzolini** [1] **Lorenzo Bisi** [1 2] **Luca Sabbioni** [1 2] **Marcello Restelli** [1 2]

Abstract

The choice of the control frequency of a system has a relevant impact on the ability of *reinforcement learning algorithms* to learn a highly performing policy. In this paper, we introduce the notion of *action persistence* that consists in the repetition of an action for a fixed number of decision steps, having the effect of modifying the control frequency. We start analyzing how action persistence affects the performance of the optimal policy, and then we present a novel algorithm, *Persistent Fitted Q-Iteration* (PFQI), that extends FQI, with the goal of learning the optimal value function at a given persistence. After having provided a theoretical study of PFQI and a heuristic approach to identify the optimal persistence, we present an experimental campaign on benchmark domains to show the advantages of action persistence and proving the effectiveness of our persistence selection method.

1. Introduction

In recent years, Reinforcement Learning (RL, Sutton & Barto, 2018) has proven to be a successful approach to address complex control tasks: from robotic locomotion (e.g., Peters & Schaal, 2008; Kober & Peters, 2014; Haarnoja et al., 2019; Kilinc et al., 2019) to continuous system control (e.g., Schulman et al., 2015; Lillicrap et al., 2016; Schulman et al., 2017). These classes of problems are usually formalized in the framework of the *discrete–time* Markov Decision Processes (MDP, Puterman, 2014), assuming that the control signal is issued at discrete time instants. However, many relevant real–world problems are more naturally defined in the continuous–time domain (Luenberger, 1979). Even though a branch of literature has studied RL in *continuous–time* MDPs (Bradtke & Duff, 1994; Munos & Bourgine, 1997; Doya, 2000), the majority of the research has focused on the discrete–time formulation, which appears to be a necessary, but effective, approximation.

Intuitively, increasing the *control frequency* of the system offers the agent more control opportunities, possibly leading to improved performance as the agent has access to a larger *policy space*. This might wrongly suggest that we should control the system with the highest frequency possible, within its physical limits. However, in the RL framework, the environment dynamics is unknown, thus, a too fine discretization could result in the opposite effect, making the problem harder to solve. Indeed, any RL algorithm needs samples to figure out (implicitly or explicitly) how the environment evolves as an effect of the agent's actions. When increasing the control frequency, the *advantage* of individual actions becomes infinitesimal, making them almost indistinguishable for standard *value-based* RL approaches (Tallec et al., 2019). As a consequence, the *sample complexity* increases. Instead, low frequencies allow the environment to evolve longer, making the effect of individual actions more easily detectable. Furthermore, in the presence of a system characterized by a "slowly evolving" dynamics, the gain obtained by increasing the control frequency might become negligible. Finally, in robotics, lower frequencies help to overcome some partial observability issues, like action execution delays (Kober & Peters, 2014).

Therefore, we experience a fundamental *trade–off* in the control frequency choice that involves the policy space (larger at high frequency) and the sample complexity (smaller at low frequency). Thus, it seems natural to wonder: "*what is the optimal control frequency?*" An answer to this question can disregard neither the task we are facing nor the learning algorithm we intend to employ. Indeed, the performance loss we experience by reducing the control frequency depends strictly on the properties of the system and, thus, of the task. Similarly, the dependence of the sample complexity on the control frequency is related to how the learning algorithm will employ the collected samples.

In this paper, we analyze and exploit this trade–off in the context of batch RL (Lange et al., 2012), with the goal of enhancing the learning process and achieving higher perfor-

[1]Politecnico di Milano, Milan, Italy. [2]Institute for Scientific Interchange Foundation, Turin, Italy. Correspondence to: Alberto Maria Metelli <albertomaria.metelli@polimi.it>.

Proceedings of the 37^{th} International Conference on Machine Learning, Online, PMLR 119, 2020.

mance. We assume to have access to a discrete–time MDP $\mathcal{M}_{\Delta t_0}$, called base MDP, which is obtained from the time discretization of a continuous–time MDP with fixed base control time step Δt_0, or equivalently, a control frequency equal to $f_0 = \frac{1}{\Delta t_0}$. In this setting, we want to select a suitable *control time step* Δt that is an integer multiple of the base time step Δt_0, i.e., $\Delta t = k\Delta t_0$ with $k \in \mathbb{N}_{\geqslant 1}$.[1] Any choice of k generates an MDP $\mathcal{M}_{k\Delta t_0}$ obtained from the base one $\mathcal{M}_{\Delta t_0}$ by altering the transition model so that each action is repeated for k times. For this reason, we refer to k as the action *persistence*, i.e., the number of decision epochs in which an action is kept fixed. It is possible to appreciate the same effect in the base MDP $\mathcal{M}_{\Delta t_0}$ by executing a (non-Markovian and non-stationary) policy that persists every action for k time steps. The idea of repeating actions has been previously employed, although heuristically, with deep RL architectures (Lakshminarayanan et al., 2017).

The contributions of this paper are theoretical, algorithmic, and experimental. We first prove that action persistence (with a fixed k) can be represented by a suitable modification of the Bellman operators, which preserves the contraction property and, consequently, allows deriving the corresponding value functions (Section 3). Since increasing the duration of the control time step $k\Delta t_0$ has the effect of degrading the performance of the optimal policy, we derive an algorithm–independent bound for the difference between the optimal value functions of MDPs $\mathcal{M}_{\Delta t_0}$ and $\mathcal{M}_{k\Delta t_0}$, which holds under Lipschitz conditions. The result confirms the intuition that the performance loss is strictly related to how fast the environment evolves as an effect of the actions (Section 4). Then, we apply the notion of action persistence in the batch RL scenario, proposing and analyzing an extension of Fitted Q-Iteration (FQI, Ernst et al., 2005). The resulting algorithm, *Persistent Fitted Q-Iteration* (PFQI) takes as input a target persistence k and estimates the corresponding optimal value function, assuming to have access to a dataset of samples collected in the base MDP $\mathcal{M}_{\Delta t_0}$ (Section 5). Once we estimate the value function for a set of candidate persistences $\mathcal{K} \subset \mathbb{N}_{\geqslant 1}$, we aim at selecting the one that yields the best performing greedy policy. Thus, we introduce a persistence selection heuristic able to approximate the optimal persistence, without requiring further interactions with the environment (Section 6). After having revised the literature (Section 7), we present an experimental evaluation on benchmark domains, to confirm our theoretical findings and evaluate our persistence selection method (Section 8). We conclude by discussing some open questions related to action persistence (Section 9). The proofs of all the results can be found in Appendix A. The code is available at `github.com/albertometelli/pfqi`.

[1] We are considering the *near–continuous* setting. This is almost w.l.o.g. compared to the continuous time since the discretization time step Δt_0 can be chosen to be arbitrarily small. Typically, a lower bound on Δt_0 is imposed by the physical limitations of the system. Thus, we restrict the search of Δt from the continuous set $\mathbb{R}_{>0}$ to the discrete set $\{k\Delta t_0, k \in \mathbb{N}_{\geqslant 1}\}$. Moreover, considering an already discretized MDP simplifies the mathematical treatment.

2. Preliminaries

In this section, we introduce the notation and the basic notions that we will employ in the remainder of the paper.

Mathematical Background Let $\mathcal{X}$ be a set with a σ-algebra $\sigma_{\mathcal{X}}$, we denote with $\mathscr{P}(\mathcal{X})$ the set of all probability measures and with $\mathscr{B}(\mathcal{X})$ the set of all bounded measurable functions over $(\mathcal{X}, \sigma_{\mathcal{X}})$. If $x \in \mathcal{X}$, we denote with δ_x the Dirac measure defined on x. Given a probability measure $\rho \in \mathscr{P}(\mathcal{X})$ and a measurable function $f \in \mathscr{B}(\mathcal{X})$, we abbreviate $\rho f = \int_{\mathcal{X}} f(x)\rho(\mathrm{d}x)$ (i.e., we use ρ as an operator). Moreover, we define the $L_p(\rho)$-norm of f as $\|f\|_{p,\rho}^p = \int_{\mathcal{X}} |f(x)|^p \rho(\mathrm{d}x)$ for $p \geqslant 1$, whereas the L_∞-norm is defined as $\|f\|_\infty = \sup_{x \in \mathcal{X}} f(x)$. Let $\mathcal{D} = \{x_i\}_{i=1}^n \subseteq \mathcal{X}$ we define the $L_p(\rho)$ empirical norm as $\|f\|_{p,\mathcal{D}}^p = \frac{1}{n}\sum_{i=1}^n |f(x_i)|^p$.

Markov Decision Processes A discrete-time Markov Decision Process (MDP, Puterman, 2014) is a 5-tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, R, \gamma)$, where $\mathcal{S}$ is a measurable set of states, $\mathcal{A}$ is a measurable set of actions, $P : \mathcal{S} \times \mathcal{A} \rightarrow \mathscr{P}(\mathcal{S})$ is the transition kernel that for each state-action pair $(s,a) \in \mathcal{S} \times \mathcal{A}$ provides the probability distribution $P(\cdot|s,a)$ of the next state, $R : \mathcal{S} \times \mathcal{A} \rightarrow \mathscr{P}(\mathbb{R})$ is the reward distribution $R(\cdot|s,a)$ for performing action $a \in \mathcal{A}$ in state $s \in \mathcal{S}$, whose expected value is denoted by $r(s,a) = \int_{\mathbb{R}} x R(\mathrm{d}x|s,a)$ and uniformly bounded by $R_{\max} < +\infty$, and $\gamma \in [0,1)$ is the discount factor.

A policy $\pi = (\pi_t)_{t \in \mathbb{N}}$ is a sequence of functions $\pi_t : \mathcal{H}_t \rightarrow \mathscr{P}(\mathcal{A})$ mapping a history $H_t = (S_0, A_0, \dots, S_{t-1}, A_{t-1}, S_t)$ of length $t \in \mathbb{N}$ to a probability distribution over $\mathcal{A}$, where $\mathcal{H}_t = (\mathcal{S} \times \mathcal{A})^t \times \mathcal{S}$. If π_t depends only on the last visited state S_t then it is called Markovian, i.e., $\pi_t : \mathcal{S} \rightarrow \mathscr{P}(\mathcal{A})$. Moreover, if π_t does not depend on explicitly t it is stationary, in this case we remove the subscript t. We denote with Π the set of Markovian stationary policies. A policy $\pi \in \Pi$ induces a (state-action) transition kernel $P^\pi : \mathcal{S} \times \mathcal{A} \rightarrow \mathscr{P}(\mathcal{S} \times \mathcal{A})$, defined for any measurable set $\mathcal{B} \subseteq \mathcal{S} \times \mathcal{A}$ as (Farahmand, 2011):

$$(P^\pi)(\mathcal{B}|s,a) = \int_{\mathcal{S}} P(\mathrm{d}s'|s,a) \int_{\mathcal{A}} \pi(\mathrm{d}a'|s') \delta_{(s',a')}(\mathcal{B}). \quad (1)$$

The *action-value function*, or Q-function, of a policy $\pi \in \Pi$ is the expected discounted sum of the rewards obtained by performing action a in state s and following policy π thereafter $Q^\pi(s,a) = \mathbb{E}\left[\sum_{t=0}^{+\infty} \gamma^t R_t | S_0 = s, A_0 = a\right]$, where $R_t \sim R(\cdot|S_t, A_t)$, $S_{t+1} \sim P(\cdot|S_t, A_t)$, and $A_{t+1} \sim \pi(\cdot|S_{t+1})$ for all $t \in \mathbb{N}$. The *value function* is the expectation of the Q-function over the actions: $V^\pi(s) = \int_{\mathcal{A}} \pi(\mathrm{d}a|s) Q^\pi(s,a)$. Given a distribution $\rho \in \mathscr{P}(\mathcal{S})$, we define the *expected return* as $J^{\rho,\pi}(s) = \int_{\mathcal{S}} \rho(\mathrm{d}s) V^\pi(s)$. The optimal Q-function is given by: $Q^*(s,a) = \sup_{\pi \in \Pi} Q^\pi(s,a)$ for all $(s,a) \in \mathcal{S} \times \mathcal{A}$. A policy π is *greedy* w.r.t. a function $f \in \mathscr{B}(\mathcal{S} \times \mathcal{A})$ if it plays only greedy actions, i.e., $\pi(\cdot|s) \in \mathscr{P}(\operatorname{argmax}_{a \in \mathcal{A}} f(s,a))$. An *optimal policy* $\pi^* \in \Pi$ is any policy greedy w.r.t. Q^*.

Given a policy $\pi \in \Pi$, the *Bellman Expectation Operator* $T^\pi : \mathscr{B}(\mathcal{S} \times \mathcal{A}) \to \mathscr{B}(\mathcal{S} \times \mathcal{A})$ and the *Bellman Optimal Operator* $T^* : \mathscr{B}(\mathcal{S} \times \mathcal{A}) \to \mathscr{B}(\mathcal{S} \times \mathcal{A})$ are defined for a bounded measurable function $f \in \mathscr{B}(\mathcal{S} \times \mathcal{A})$ and $(s,a) \in \mathcal{S} \times \mathcal{A}$ as (Bertsekas & Shreve, 2004):

$$(T^\pi f)(s,a) = r(s,a) + (P^\pi f)(s,a),$$
$$(T^* f)(s,a) = r(s,a) + \gamma \int_{\mathcal{S}} P(\mathrm{d}s'|s,a) \max_{a' \in \mathcal{A}} f(s',a').$$

Both T^π and T^* are γ-contractions in L_∞-norm and, consequently, they have a unique fixed point, that are the Q-function of policy π ($T^\pi Q^\pi = Q^\pi$) and the optimal Q-function ($T^* Q^* = Q^*$) respectively.

Lipschitz MDPs Let $(\mathcal{X}, d_\mathcal{X})$ and $(\mathcal{Y}, d_\mathcal{Y})$ be two metric spaces, a function $f : \mathcal{X} \to \mathcal{Y}$ is called L_f-Lipschitz continuous (L_f-LC), where $L_f \geqslant 0$, if for all $x, x' \in \mathcal{X}$ we have:

$$d_\mathcal{Y}(f(x), f(x')) \leqslant L_f d_\mathcal{X}(x, x'). \tag{2}$$

Moreover, we define the Lipschitz semi-norm as $\|f\|_L = \sup_{x,x' \in \mathcal{X} : x \neq x'} \frac{d_\mathcal{Y}(f(x), f(x'))}{d_\mathcal{X}(x,x')}$. For real functions we employ Euclidean distance $d_\mathcal{Y}(y,y') = \|y - y'\|_2$, while for probability distributions we use the Kantorovich (L_1-Wasserstein) metric defined for $\mu, \nu \in \mathscr{P}(\mathcal{Z})$ as (Villani, 2008):

$$d_\mathcal{Y}(\mu,\nu) = \mathcal{W}_1(\mu,\nu) = \sup_{f : \|f\|_L \leqslant 1} \left| \int_\mathcal{Z} f(z)(\mu - \nu)(\mathrm{d}z) \right|. \tag{3}$$

We now introduce the notions of Lipschitz MDP and Lipschitz policy that we will employ in the following (Rachelson & Lagoudakis, 2010; Pirotta et al., 2015).

Assumption 2.1 (Lipschitz MDP)**.** *Let $\mathcal{M}$ be an MDP. $\mathcal{M}$ is called (L_P, L_r)-LC if for all $(s,a), (\overline{s}, \overline{a}) \in \mathcal{S} \times \mathcal{A}$:*

$$\mathcal{W}_1(P(\cdot|s,a), P(\cdot|\overline{s},\overline{a})) \leqslant L_P d_{\mathcal{S} \times \mathcal{A}}((s,a),(\overline{s},\overline{a})),$$
$$|r(s,a) - r(\overline{s},\overline{a})| \leqslant L_r d_{\mathcal{S} \times \mathcal{A}}((s,a),(\overline{s},\overline{a})).$$

Assumption 2.2 (Lipschitz Policy)**.** *Let $\pi \in \Pi$ be a Markovian stationary policy. π is called L_π-LC if for all $s, \overline{s} \in \mathcal{S}$:*

$$\mathcal{W}_1(\pi(\cdot|s), \pi(\cdot|\overline{s})) \leqslant L_\pi d_\mathcal{S}(s, \overline{s}).$$

3. Persisting Actions in MDPs

By the phrase "executing a policy π at persistence k", with $k \in \mathbb{N}_{\geqslant 1}$, we mean the following type of agent-environment interaction. At decision step $t = 0$, the agent selects an action according to its policy $A_0 \sim \pi(\cdot|S_0)$. Action A_0 is kept fixed, or *persisted*, for the subsequent $k-1$ decision steps, i.e., actions $A_1, \ldots, A_{k-1}$ are all equal to A_0. Then, at decision step $t = k$, the agent queries again the policy $A_k \sim \pi(\cdot|S_k)$ and persists action A_k for the subsequent $k-1$ decision steps and so on. In other words, the agent employs its policy only at decision steps t that are integer multiples of the persistence k ($t \bmod k = 0$). Clearly, the usual execution of π corresponds to persistence 1.

3.1. Duality of Action Persistence

Unsurprisingly, the execution of a Markovian stationary policy π at persistence $k > 1$ produces a behavior that, in general, cannot be represented by executing any Markovian stationary policy at persistence 1. Indeed, at any decision step t, such a policy needs to remember which action was taken at the previous decision step $t-1$ (thus it is non-Markovian with memory 1) and has to understand whether to select a new action based on t (so it is non-stationary).

Definition 3.1 (k-persistent policy)**.** *Let $\pi \in \Pi$ be a Markovian stationary policy. For any $k \in \mathbb{N}_{\geqslant 1}$, the* k-persistent policy *induced by π is a non–Markovian non–stationary policy, defined for any measurable set $\mathcal{B} \subseteq \mathcal{A}$ and $t \in \mathbb{N}$ as:*

$$\pi_{t,k}(\mathcal{B}|H_t) = \begin{cases} \pi(\mathcal{B}|S_t) & \text{if } t \bmod k = 0 \\ \delta_{A_{t-1}}(\mathcal{B}) & \text{otherwise} \end{cases}. \tag{4}$$

Moreover, we denote with $\Pi_k = \{(\pi_{t,k})_{t \in \mathbb{N}} : \pi \in \Pi\}$ the set of the k-persistent policies.

Clearly, for $k = 1$ we recover policy π as we always satisfy the condition $t \bmod k = 0$ i.e., $\pi = \pi_{t,1}$ for all $t \in \mathbb{N}$. We refer to this interpretation of action persistence as *policy view*.

A different perspective towards action persistence consists in looking at the effect of the original policy π in a suitably modified MDP. To this purpose, we introduce the (state-action) persistent transition probability kernel $P^\delta : \mathcal{S} \times \mathcal{A} \to \mathscr{P}(\mathcal{S} \times \mathcal{A})$ defined for any measurable set $\mathcal{B} \subseteq \mathcal{S} \times \mathcal{A}$ as:

$$(P^\delta)(\mathcal{B}|s,a) = \int_\mathcal{S} P(\mathrm{d}s'|s,a) \delta_{(s',a)}(\mathcal{B}). \tag{5}$$

The crucial difference between P^π and P^δ is that the former samples the action a' to be executed in the next state s' according to π, whereas the latter replicates in state s' action a. We are now ready to define the k-persistent MDP.

Definition 3.2 (k-persistent MDP)**.** *Let $\mathcal{M}$ be an MDP. For any $k \in \mathbb{N}_{\geqslant 1}$, the* k-persistent MDP *is the following MDP $\mathcal{M}_k = (\mathcal{S}, \mathcal{A}, P_k, R_k, \gamma^k)$, where P_k and R_k are the k-persistent transition model and reward distribution respectively, defined for any measurable sets $\mathcal{B} \subseteq \mathcal{S}$, $\mathcal{C} \subseteq \mathbb{R}$ and state-action pair $(s,a) \in \mathcal{S} \times \mathcal{A}$ as:*

$$P_k(\mathcal{B}|s,a) = \left((P^\delta)^{k-1} P\right)(\mathcal{B}|s,a), \tag{6}$$

$$R_k(\mathcal{C}|s,a) = \sum_{i=0}^{k-1} \gamma^i \left((P^\delta)^i R\right)(\mathcal{C}|s,a), \tag{7}$$

and $r_k(s,a) = \int_\mathbb{R} x R_k(\mathrm{d}x|s,a) = \sum_{i=0}^{k-1} \gamma^i \left((P^\delta)^i r\right)(s,a)$ is the expected reward, uniformly bounded by $R_{\max} \frac{1-\gamma^k}{1-\gamma}$.

The k-persistent transition model P_k keeps action a fixed for $k-1$ steps while making the state evolve according to P.

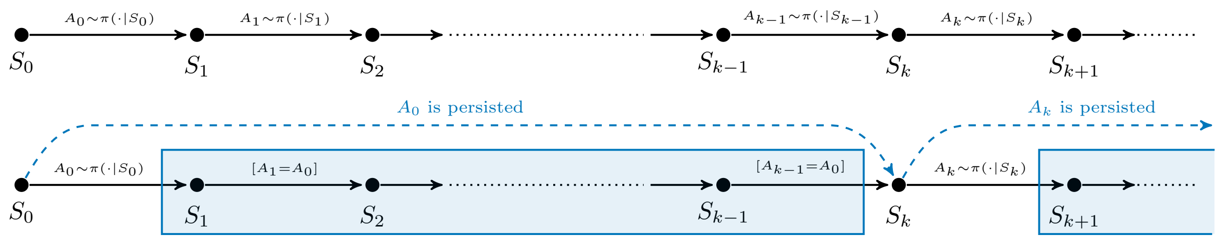

Figure 1. Agent-environment interaction without (top) and with (bottom) action persistence, highlighting duality. The transition generated by the k-persistent MDP $\mathcal{M}_k$ is the cyan dashed arrow, while the actions played by the k-persistent policy are inside the cyan rectangle.

Similarly, the k-persistent reward R_k provides the cumulative discounted reward over k steps in which a is persisted. We define the transition kernel P_k^π, analogously to P^π, as in Equation (1). Clearly, for $k=1$ we recover the base MDP, i.e., $\mathcal{M}=\mathcal{M}_1$.[2] Therefore, executing policy π in $\mathcal{M}_k$ at persistence 1 is equivalent to executing policy π at persistence k in the original MDP $\mathcal{M}$. We refer to this interpretation of persistence as *environment view* (Figure 1). Thus, solving the base MDP $\mathcal{M}$ in the space of k-persistent policies Π_k (Definition 3.1), thanks to this *duality*, is equivalent to solving the k-persistent MDP $\mathcal{M}_k$ (Definition 3.2) in the space of Markovian stationary policies Π.

It is worth noting that the persistence $k\in\mathbb{N}_{\geqslant 1}$ can be seen as an *environmental parameter* (affecting P, R, and γ), which can be externally configured with the goal to improve the learning process for the agent. In this sense, the MDP $\mathcal{M}_k$ can be seen as a Configurable Markov Decision Process with parameter $k\in\mathbb{N}_{\geqslant 1}$ (Metelli et al., 2018; 2019).

Furthermore, a persistence of k induces a k-persistent MDP $\mathcal{M}_k$ with smaller discount factor γ^k. Therefore, the effective horizon in $\mathcal{M}_k$ is $\frac{1}{1-\gamma^k}<\frac{1}{1-\gamma}$. Interestingly, the end effect of persisting actions is similar to reducing the planning horizon, by explicitly reducing the discount factor of the task (Petrik & Scherrer, 2008; Jiang et al., 2016) or setting a maximum trajectory length (Farahmand et al., 2016).

3.2. Persistent Bellman Operators

When executing policy π at persistence k in the base MDP $\mathcal{M}$, we can evaluate its performance starting from any state-action pair $(s,a)\in\mathcal{S}\times\mathcal{A}$, inducing a Q-function that we denote with Q_k^π and call *k-persistent action-value function* of π. Thanks to duality, Q_k^π is also the action-value function of policy π when executed in the k-persistent MDP $\mathcal{M}_k$. Therefore, Q_k^π is the fixed point of the Bellman Expectation Operator of $\mathcal{M}_k$, i.e., the operator defined for any $f\in\mathscr{B}(\mathcal{S}\times\mathcal{A})$ as $(T_k^\pi f)(s,a)=r_k(s,a)+\gamma^k(P_k^\pi f)(s,a)$, that we call *$k$-persistent Bellman Expectation Operator*. Similarly, again thanks to duality, the optimal Q-function in the space of k-persistent policies Π_k, denoted by Q_k^* and called *k-persistent optimal action-value function*, corresponds to the optimal Q-function of the k-persistent MDP, i.e., $Q_k^*(s,a)=\sup_{\pi\in\Pi}Q_k^\pi(s,a)$ for all $(s,a)\in\mathcal{S}\times\mathcal{A}$. As a consequence, Q_k^* is the fixed point of the Bellman Optimal Operator of $\mathcal{M}_k$, defined for $f\in\mathscr{B}(\mathcal{S}\times\mathcal{A})$ as $(T_k^* f)(s,a)=r_k(s,a)+\gamma^k\int_{\mathcal{S}}P_k(\mathrm{d}s'|s,a)\max_{a'\in\mathcal{A}}f(s',a')$, that we call *$k$-persistent Bellman Optimal Operator*. Clearly, both T_k^π and T_k^* are γ^k-contractions in L_∞-norm.

We now prove that the k-persistent Bellman operators are obtained as composition of the base operators T^π and T^*.

Theorem 3.1. *Let $\mathcal{M}$ be an MDP, $k\in\mathbb{N}_{\geqslant 1}$ and $\mathcal{M}_k$ be the k-persistent MDP. Let $\pi\in\Pi$ be a Markovian stationary policy. Then, T_k^π and T_k^* can be expressed as:*

$$T_k^\pi=\left(T^\delta\right)^{k-1}T^\pi \quad \text{and} \quad T_k^*=\left(T^\delta\right)^{k-1}T^*, \tag{8}$$

where $T^\delta:\mathscr{B}(\mathcal{S}\times\mathcal{A})\to\mathscr{B}(\mathcal{S}\times\mathcal{A})$ is the Bellman Persistent Operator, *defined for $f\in\mathscr{B}(\mathcal{S}\times\mathcal{A})$ and $(s,a)\in\mathcal{S}\times\mathcal{A}$:*

$$\left(T^\delta f\right)(s,a)=r(s,a)+\gamma\left(P^\delta f\right)(s,a). \tag{9}$$

The fixed point equations for the k-persistent Q-functions become: $Q_k^\pi=\left(T^\delta\right)^{k-1}T^\pi Q_k^\pi$ and $Q_k^*=\left(T^\delta\right)^{k-1}T^* Q_k^*$.

4. Bounding the Performance Loss

Learning in the space of k-persistent policies Π_k can only lower the performance of the optimal policy, i.e., $Q^*(s,a)\geqslant Q_k^*(s,a)$ for $k\in\mathbb{N}_{\geqslant 1}$. The goal of this section is to bound $\|Q^*-Q_k^*\|_{p,\rho}$ as a function of the persistence $k\in\mathbb{N}_{\geqslant 1}$. To this purpose, we focus on $\|Q^\pi-Q_k^\pi\|_{p,\rho}$ for a fixed policy $\pi\in\Pi$, since denoting with π^* an optimal policy of $\mathcal{M}$ and with π_k^* an optimal policy of $\mathcal{M}_k$, we have that:

$$Q^*-Q_k^*=Q^{\pi^*}-Q_k^{\pi_k^*}\leqslant Q^{\pi^*}-Q_k^{\pi^*},$$

since $Q_k^{\pi_k^*}(s,a)\geqslant Q_k^{\pi^*}(s,a)$. We start with the following result which makes no assumption about the structure of the MDP and then we particularize it for the Lipschitz MDPs.

Theorem 4.1. *Let $\mathcal{M}$ be an MDP and $\pi\in\Pi$ be a Markovian stationary policy. Let $\mathcal{Q}_k=\{\left(T^\delta\right)^{k-2-l}T^\pi Q_k^\pi : l\in\{0,\dots,k-2\}\}$ and for all $(s,a)\in\mathcal{S}\times\mathcal{A}$ let us define:*

$$d_{\mathcal{Q}_k}^\pi(s,a)=\sup_{f\in\mathcal{Q}_k}\left|\int_{\mathcal{S}}\int_{\mathcal{A}}\left(P^\pi(\mathrm{d}s',\mathrm{d}a'|s,a)-P^\delta(\mathrm{d}s',\mathrm{d}a'|s,a)\right)f(s',a')\right|.$$

[2] If $\mathcal{M}$ is the base MDP $\mathcal{M}_{\Delta t_0}$, the k–persistent MDP $\mathcal{M}_k$ corresponds to $\mathcal{M}_{k\Delta t_0}$. We remove the subscript Δt_0 for brevity.

Then, for any $\rho \in \mathscr{P}(\mathcal{S}\times\mathcal{A})$, $p \geqslant 1$, and $k \in \mathbb{N}_{\geqslant 1}$, it holds that:

$$\|Q^\pi - Q_k^\pi\|_{p,\rho} \leqslant \frac{\gamma(1-\gamma^{k-1})}{(1-\gamma)(1-\gamma^k)} \|d_{\mathcal{Q}_k}^\pi\|_{p,\eta_k^{\rho,\pi}},$$

where $\eta_k^{\rho,\pi} \in \mathscr{P}(\mathcal{S}\times\mathcal{A})$ is a probability measure defined for any measurable set $\mathcal{B} \subseteq \mathcal{S}\times\mathcal{A}$ as:

$$\eta_k^{\rho,\pi}(\mathcal{B}) = \frac{(1-\gamma)(1-\gamma^k)}{\gamma(1-\gamma^{k-1})} \sum_{\substack{i\in\mathbb{N}\\ i \bmod k \neq 0}} \gamma^i \left(\rho (P^\pi)^{i-1}\right)(\mathcal{B}).$$

The bound shows that the Q-function difference depends on the discrepancy $d_{\mathcal{Q}_k}^\pi$ between the transition-kernel P^π and the corresponding persistent version P^δ, which is a form of *integral probability metric* (Müller, 1997), defined in terms of the set $\mathcal{Q}_k$. This term is averaged with the distribution $\eta_k^{\rho,\pi}$, which encodes the (discounted) probability of visiting a state-action pair, ignoring the visitations made at decision steps i that are multiple of the persistence k. Indeed, in those steps, we play policy π regardless of whether persistence is used.[3] The dependence on k is represented in the term $\frac{1-\gamma^{k-1}}{1-\gamma^k}$. When $k \rightarrow 1$ this term displays a linear growth in k, being asymptotic to $(k-1)\log\frac{1}{\gamma}$, and, clearly, vanishes for $k=1$. Instead, when $k \rightarrow \infty$ this term tends to 1.

If no structure on the MDP/policy is enforced, the dissimilarity term $d_{\mathcal{Q}_k}^\pi$ may become large enough to make the bound vacuous, i.e., larger than $\frac{\gamma R_{\max}}{1-\gamma}$, even for $k=2$ (see Appendix B.1). Intuitively, since the persistence will execute old actions in new states, we need to guarantee that the environment state changes slowly w.r.t. to time and the policy must play similar actions in similar states. This means that if an action is good in a state, it will also be almost good for states encountered in the near future. Although the condition on π is directly enforced by Assumption 2.2, we need a new notion of regularity over time for the MDP.

Assumption 4.1. *Let $\mathcal{M}$ be an MDP. $\mathcal{M}$ is L_T–*Time-Lipschitz Continuous *(L_T–TLC) if for all $(s,a) \in \mathcal{S}\times\mathcal{A}$:*

$$\mathcal{W}_1(P(\cdot|s,a),\delta_s) \leqslant L_T. \tag{10}$$

This assumption requires that the Kantorovich distance between the distribution of the next state s' and the deterministic distribution centered in the current state s is bounded by L_T, i.e., the system does not evolve "too fast" (see Appendix B.3). We can now state the following result.

Theorem 4.2. *Let $\mathcal{M}$ be an MDP and $\pi \in \Pi$ be a Markovian stationary policy. Under Assumptions 2.1, 2.2, and 4.1, if $\gamma\max\{L_P+1, L_P(1+L_\pi)\} < 1$ and if $\rho(s,a) = \rho_{\mathcal{S}}(s)\pi(a|s)$ with $\rho_{\mathcal{S}} \in \mathscr{P}(\mathcal{S})$, then for any $k \in \mathbb{N}_{\geqslant 1}$:*

$$\|d_{\mathcal{Q}_k}^\pi\|_{p,\eta_k^{\rho,\pi}} \leqslant L_{\mathcal{Q}_k}\left[(L_\pi+1)L_T + \sigma_p\right].$$

where $\sigma_p^p = \sup_{s\in\mathcal{S}} \int_{\mathcal{A}}\int_{\mathcal{A}} d_{\mathcal{A}}(a,a')^p \pi(\mathrm{d}a|s)\pi(\mathrm{d}a'|s)$, and $L_{\mathcal{Q}_k} = \frac{L_r}{1-\gamma\max\{L_P+1, L_P(1+L_\pi)\}}$.

Thus, the dissimilarity $d_{\mathcal{Q}_k}^\pi$ between P^π and P^δ can be bounded with four terms. i) $L_{\mathcal{Q}_k}$ is (an upper-bound of) the Lipschitz constant of the functions in the set $\mathcal{Q}_k$. Indeed, under Assumptions 2.1 and 2.2 we can reduce the dissimilatity term to the Kantorivich distance (Lemma A.5):

$$d_{\mathcal{Q}_k}^\pi(s,a) \leqslant L_{\mathcal{Q}_k}\mathcal{W}_1\left(P^\pi(\cdot|s,a), P^\delta(\cdot|s,a)\right).$$

ii) $(L_\pi+1)$ accounts for the Lipschitz continuity of the policy, i.e., policies that prescribe similar actions in similar states have a small value of this quantity. iii) L_T represents the speed at which the environment state evolves over time. iv) σ_p denotes the average distance (in L_p-norm) between two actions prescribed by the policy in the same state. This term is zero for deterministic policies and can be related to the maximum policy variance (Lemma A.6). A more detailed discussion on the conditions requested in Theorem 4.2 is reported in Appendix B.4.

5. Persistent Fitted Q-Iteration

In this section, we introduce an extension of Fitted Q-Iteration (FQI, Ernst et al., 2005) that employs the notion of persistence.[4] *Persisted Fitted Q-Iteration* (PFQI(k)) takes as input a *target persistence* $k \in \mathbb{N}_{\geqslant 1}$ and its goal is to approximate the k-persistent optimal action-value function Q_k^*. Starting from an initial estimate $Q^{(0)}$, at each iteration we compute the next estimate $Q^{(j+1)}$ by performing an approximate application of k-persistent Bellman optimal operator to the previous estimate $Q^{(j)}$, i.e., $Q^{(j+1)} \approx T_k^* Q^{(j)}$. In practice, we have two sources of approximation in this process: i) the representation of the Q-function; ii) the estimation of the k-persistent Bellman optimal operator. (i) comes from the necessity of using functional space $\mathcal{F} \subset \mathscr{B}(\mathcal{S}\times\mathcal{A})$ to represent $Q^{(j)}$ when dealing with continuous state spaces. (ii) derives from the approximate computation of T_k^* which needs to be estimated from samples.

Clearly, with samples collected in the k-persistent MDP $\mathcal{M}_k$, the process described above reduces to the standard FQI. However, our algorithm needs to be able to estimate Q_k^* for different values of k, using the same dataset of samples collected in the base MDP $\mathcal{M}$ (at persistence 1).[5] For this purpose, we can exploit the decomposition $T_k^* = (T^\delta)^{k-1} T^*$ of Theorem 3.1 to reduce a single application of T_k^* to a sequence of k applications of the 1-persistent operators. Specifically, at each iteration j with $j \bmod k = 0$, given the current estimate $Q^{(j)}$, we need to perform (in this order) a single application of T^* followed by $k-1$ applica-

[3] $\eta_k^{\rho,\pi}$ resambles the γ-discounted state-action distribution (Sutton et al., 1999a), but ignoring the decision steps multiple of k.

[4] From now on, we assume that $|\mathcal{A}| < +\infty$.

[5] In real–world cases, we might be unable to interact with the physical system to collect samples for any persistence k of interest.

Algorithm 1 Persistent Fitted Q-Iteration PFQI(k).

Input: k persistence, J number of iterations ($J \bmod k = 0$), $Q^{(0)}$ initial action-value function, $\mathcal{F}$ functional space, $\mathcal{D} = \{(S_i, A_i, S'_i, R_i)\}_{i=1}^n$ batch samples
Output: greedy policy $\pi^{(J)}$

for $j = 0, \dots, J-1$ **do**
  Phase 1
  if $j \bmod k = 0$ **then**
    $Y_i^{(j)} = \widehat{T}^* Q^{(j)}(S_i, A_i), \quad i = 1, \dots, n$
  else
    $Y_i^{(j)} = \widehat{T}^\delta Q^{(j)}(S_i, A_i), \quad i = 1, \dots, n$
  end if
  Phase 2
  $Q^{(j+1)} \in \operatorname{arginf}_{f \in \mathcal{F}} \|f - Y^{(j)}\|_{2,\mathcal{D}}^2$
end for
Phase 3
$\pi^{(J)}(s) \in \operatorname{argmax}_{a \in \mathcal{A}} Q^{(J)}(s,a), \quad \forall s \in \mathcal{S}$

tions of T^δ, leading to the sequence of approximations:

$$Q^{(j+1)} \approx \begin{cases} T^* Q^{(j)} & \text{if } j \bmod k = 0 \\ T^\delta Q^{(j)} & \text{otherwise} \end{cases}. \qquad (11)$$

In order to estimate the Bellman operators, we have access to a dataset $\mathcal{D} = \{(S_i, A_i, S'_i, R_i)\}_{i=1}^n$ collected in the base MDP $\mathcal{M}$, where $(S_i, A_i) \sim \nu$, $S'_i \sim P(\cdot|S_i, A_i)$, $R_i \sim R(\cdot|S_i, A_i)$, and $\nu \in \mathscr{P}(\mathcal{S} \times \mathcal{A})$ is a sampling distribution. We employ $\mathcal{D}$ to compute the *empirical Bellman operators* (Farahmand, 2011) defined for $f \in \mathscr{B}(\mathcal{S} \times \mathcal{A})$ as:

$$(\widehat{T}^* f)(S_i, A_i) = R_i + \gamma \max_{a \in \mathcal{A}} f(S'_i, a) \quad i = 1, \dots, n$$
$$(\widehat{T}^\delta f)(S_i, A_i) = R_i + \gamma f(S'_i, A_i) \quad i = 1, \dots, n.$$

These operators are unbiased conditioned to $\mathcal{D}$ (Farahmand, 2011): $\mathbb{E}[(\widehat{T}^* f)(S_i, A_i)|S_i, A_i] = (T^* f)(S_i, A_i)$ and $\mathbb{E}[(\widehat{T}^\delta f)(S_i, A_i)|S_i, A_i] = (T^\delta f)(S_i, A_i)$.

The pseudocode of PFQI(k) is summarized in Algorithm 1. At each iteration $j = 0, \dots J-1$, we first compute the target values $Y^{(j)}$ by applying the empirical Bellman operators, $\widehat{T}^*$ or $\widehat{T}^\delta$, on the current estimate $Q^{(j)}$ (Phase 1). Then, we project the target $Y^{(j)}$ onto the functional space $\mathcal{F}$ by solving the least squares problem (Phase 2):

$$Q^{(j+1)} \in \operatorname*{arginf}_{f \in \mathcal{F}} \left\| f - Y^{(j)} \right\|_{2,\mathcal{D}}^2 = \frac{1}{n} \sum_{i=1}^n \left| f(S_i, A_i) - Y_i^{(j)} \right|^2.$$

Finally, we compute the approximation of the optimal policy $\pi^{(J)}$, i.e., the greedy policy w.r.t. $Q^{(J)}$ (Phase 3).

5.1. Theoretical Analysis

In this section, we present the computational complexity analysis and the study of the error propagation in PFQI(k).

Computational Complexity The computational complexity of PFQI(k) decreases monotonically with the persistence k. Whenever applying $\widehat{T}^\delta$, we need a single evaluation of $Q^{(j)}$, while $|\mathcal{A}|$ evaluations are needed for $\widehat{T}^*$ due to the max over $\mathcal{A}$. Thus, the overall complexity of J iterations of PFQI(k) with n samples, disregarding the cost of regression and assuming that a single evaluation of $Q^{(j)}$ takes constant time, is given by $\mathcal{O}(Jn(1 + (|\mathcal{A}| - 1)/k))$ (Proposition A.1).

Error Propagation We now consider the error propagation in PFQI(k). Given the sequence of Q-functions estimates $(Q^{(j)})_{j=0}^J \subset \mathcal{F}$ produced by PFQI(k), we define the approximation error at each iteration $j = 0, \dots, J-1$ as:

$$\epsilon^{(j)} = \begin{cases} T^* Q^{(j)} - Q^{(j+1)} & \text{if } j \bmod k = 0 \\ T^\delta Q^{(j)} - Q^{(j+1)} & \text{otherwise} \end{cases}. \qquad (12)$$

The goal of this analysis is to bound the distance between the k–persistent optimal Q-function Q_k^* and the Q-function $Q_k^{\pi^{(J)}}$ of the greedy policy $\pi^{(J)}$ w.r.t. $Q^{(J)}$, after J iterations of PFQI(k). The following result extends Theorem 3.4 of Farahmand (2011) to account for action persistence.

Theorem 5.1 (Error Propagation for PFQI(k))**.** *Let $p \geqslant 1$, $k \in \mathbb{N}_{\geqslant 1}$, $J \in \mathbb{N}_{\geqslant 1}$ with $J \bmod k = 0$ and $\rho \in \mathscr{P}(\mathcal{S} \times \mathcal{A})$. Then for any sequence $(Q^{(j)})_{j=0}^J \subset \mathcal{F}$ uniformly bounded by $Q_{\max} \leqslant \frac{R_{\max}}{1-\gamma}$, the corresponding $(\epsilon^{(j)})_{j=0}^{J-1}$ defined in Equation (12) and for any $r \in [0,1]$ and $q \in [1, +\infty]$ it holds that:*

$$\left\| Q_k^* - Q_k^{\pi^{(J)}} \right\|_{p,\rho} \leqslant \frac{2\gamma^k}{(1-\gamma)(1-\gamma^k)} \left[\frac{2}{1-\gamma} \gamma^{\frac{J}{p}} R_{\max} + C_{\mathrm{VI},\rho,\nu}^{\frac{1}{2p}}(J; r, q) \mathcal{E}^{\frac{1}{2p}}(\epsilon^{(0)}, \dots, \epsilon^{(J-1)}; r, q) \right].$$

The expression of $C_{\mathrm{VI},\rho,\nu}(J; r, q)$ and $\mathcal{E}(\cdot; r, q)$ can be found in Appendix A.3.

We immediately observe that for $k = 1$ we recover Theorem 3.4 of Farahmand (2011). The term $C_{\mathrm{VI},\rho,\nu}(J; r, q)$ is defined in terms of suitable *concentrability coefficients* (Definition A.1) and encodes the distribution shift between the sampling distribution ν and the one induced by the greedy policy sequence $(\pi^{(j)})_{j=0}^J$ encountered along the execution of PFQI(k). $\mathcal{E}(\cdot; r, q)$ incorporates the approximation errors $(\epsilon^{(j)})_{j=0}^{J-1}$. In principle, it is hard to compare the values of these terms for different persistences k since both the greedy policies and the regression problems are different. Nevertheless, it is worth noting that the multiplicative term $\frac{\gamma^k}{1-\gamma^k}$ decreases in $k \in \mathbb{N}_{\geqslant 1}$. Thus, other things being equal, the bound value decreases when increasing the persistence.

Thus, the trade-off in the choice of control frequency, which motivates action persistence, can now be stated more formally. We aim at finding the persistence $k \in \mathbb{N}_{\geqslant 1}$ that, for a fixed J, allows learning a policy $\pi^{(J)}$ whose Q-function $Q_k^{\pi^{(J)}}$ is the closest to Q^*. Consider the decomposition:

$$\left\| Q^* - Q_k^{\pi^{(J)}} \right\|_{p,\rho} \leqslant \| Q^* - Q_k^* \|_{p,\rho} + \left\| Q_k^* - Q_k^{\pi^{(J)}} \right\|_{p,\rho}.$$

The term $\|Q^* - Q_k^*\|_{p,\rho}$ accounts for the performance degradation due to action persistence: it is algorithm–independent, and it increases in k (Theorem 4.1). Instead, the second term $\|Q_k^* - Q_k^{\pi^{(J)}}\|_{p,\rho}$ decreases with k and depends on the algo-

Algorithm 2 Heuristic Persistence Selection.

Input: batch samples $\mathcal{D}=\{(S_0^i,A_0^i,\dots,S_{H_i-1}^i,A_{H_i-1}^i,S_{H_i}^i)\}_{i=1}^m$, set of persistences $\mathcal{K}$, set of Q-function $\{Q_k : k\in\mathcal{K}\}$, regressor Reg
Output: approximately optimal persistence $\widetilde{k}$
for $k\in\mathcal{K}$ **do**
  $\widehat{J}_k^\rho=\frac{1}{m}\sum_{i=1}^m V_k(S_0^i)$
  Use the Reg to get an estimate $\widetilde{Q}_k$ of $T_k^* Q_k$
  $\|\widetilde{Q}_k-Q_k\|_{1,\mathcal{D}}=\frac{1}{\sum_{i=1}^m H_i}\sum_{i=1}^m\sum_{t=0}^{H_i-1}|\widetilde{Q}_k(S_t^i,A_t^i)-Q_k(S_t^i,A_t^i)|$
end for
$\widetilde{k}\in\operatorname{argmax}_{k\in\mathcal{K}} B_k=\widehat{J}_k^\rho-\frac{1}{1-\gamma^k}\|\widetilde{Q}_k-Q_k\|_{1,\mathcal{D}}$.

rithm (Theorem 5.1). Unfortunately, optimizing their sum is hard since the individual bounds contain terms that are not known in general (e.g., Lipschitz constants, $\epsilon^{(j)}$). The next section proposes heuristics to overcome this problem.

6. Persistence Selection

In this section, we discuss how to select a persistence k in a set $\mathcal{K}\subset\mathbb{N}_{\geqslant 1}$ of candidate persistences, when we are given a set of estimated Q-functions: $\{Q_k : k\in\mathcal{K}\}$.[6] Each Q_k induces a greedy policy π_k. Our goal is to find the persistence $k\in\mathcal{K}$ such that π_k has the maximum expected return in the corresponding k–persistent MDP $\mathcal{M}_k$:

$$k^*\in\operatorname*{argmax}_{k\in\mathcal{K}} J_k^{\rho,\pi_k},\quad \rho\in\mathscr{P}(\mathcal{S}). \tag{13}$$

In principle, we could execute π_k in $\mathcal{M}_k$ to get an estimate of J_k^{ρ,π_k} and employ it to select the persistence k. However, in the batch setting, further interactions with the environment might be not allowed. On the other hand, directly using the estimated Q-function Q_k is inappropriate, since we need to take into account how well Q_k approximates $Q_k^{\pi_k}$. This trade–off is encoded in the following result, which makes use of the *expected Bellman residual*.

Lemma 6.1. *Let $Q\in\mathscr{B}(\mathcal{S}\times\mathcal{A})$ and π be a greedy policy w.r.t. Q. Let $J^\rho=\int\rho(\mathrm{d}s)V(s)$, with $V(s)=\max_{a\in\mathcal{A}}Q(s,a)$ for all $s\in\mathcal{S}$. Then, for any $k\in\mathbb{N}_{\geqslant 1}$, it holds that:*

$$J_k^{\rho,\pi}\geqslant J^\rho-\frac{1}{1-\gamma^k}\|T_k^*Q-Q\|_{1,\eta^{\rho,\pi}}, \tag{14}$$

where $\eta^{\rho,\pi}=(1-\gamma^k)\rho\pi\left(\mathrm{Id}-\gamma^k P_k^\pi\right)^{-1}$, is the γ-discounted stationary distribution induced by policy π and distribution ρ in MDP $\mathcal{M}_k$.

To get a usable bound, we need to make some simplifications. First, we assume that $\mathcal{D}\sim\nu$ is composed of m *trajectories*, i.e., $\mathcal{D}=\{(S_0^i,A_0^i,\dots,S_{H_i-1}^i,A_{H_i-1}^i,S_{H_i}^i)\}_{i=1}^m$, where H_i is the trajectory length and the initial states are sampled as $S_0^i\sim\rho$. In this way, J^ρ can be estimated from samples as $\widehat{J}^\rho=\frac{1}{m}\sum_{i=1}^m V(S_0^i)$. Second, since we are unable to compute expectations over $\eta^{\rho,\pi}$, we replace it with the sampling distribution ν.[7] Lastly, estimating the expected Bellman residual is problematic since its empirical version is biased (Antos et al., 2008). Thus, we resort to an approach similar to (Farahmand & Szepesvári, 2011), assuming to have a regressor Reg able to output an approximation $\widetilde{Q}_k$ of T_k^*Q. In this way, we replace $\|T_k^*Q-Q\|_{1,\nu}$ with $\|\widetilde{Q}_k-Q\|_{1,\mathcal{D}}$ (details in Appendix C). In practice, we set $Q=Q^{(J)}$ and we obtain $\widetilde{Q}_k$ running PFQI(k) for k additional iterations, setting $\widetilde{Q}_k=Q^{(J+k)}$. Thus, the procedure (Algorithm 2) reduces to optimizing the index:

$$\widetilde{k}\in\operatorname*{argmax}_{k\in\mathcal{K}} B_k=\widehat{J}_k^\rho-\frac{1}{1-\gamma^k}\left\|\widetilde{Q}_k-Q_k\right\|_{1,\mathcal{D}}. \tag{15}$$

7. Related Works

In this section, we revise the works connected to persistence, focusing on continuous–time RL and temporal abstractions.

Continuous–time RL Among the first attempts to extend value–based RL to the continuous–time domain there is *advantage updating* (Bradtke & Duff, 1994), in which Q-learning (Watkins, 1989) is modified to account for infinitesimal control timesteps. Instead of storing the Q-function, the *advantage function* $A(s,a)=Q(s,a)-V(s)$ is recorder. The continuous time is addressed in Baird (1994) by means of the semi-Markov decision processes (Howard, 1963) for finite–state problems. The optimal control literature has extensively studied the solution of the Hamilton-Jacobi-Bellman equation, i.e., the continuous–time counterpart of the Bellman equation, when assuming the knowledge of the environment (Bertsekas, 2005; Fleming & Soner, 2006). The model–free case has been tackled by resorting to time (and space) discretizations (Peterson, 1993), with also convergence guarantees (Munos, 1997; Munos & Bourgine, 1997), and coped with function approximation (Dayan & Singh, 1995; Doya, 2000). More recently, the sensitivity of deep RL algorithm to the time discretization has been analyzed in Tallec et al. (2019), proposing an adaptation of advantage updating to deal with small time scales, that can be employed with deep architectures.

Temporal Abstractions The notion of action persistence can be seen as a form of *temporal abstraction* (Sutton et al., 1999b; Precup, 2001). Temporally extended actions have been extensively used in the hierarchical RL literature to model different time resolutions (Singh, 1992a;b), subgoals (Dietterich, 1998), and combined with the actor–critic architectures (Bacon et al., 2017). Persisting an action is a particular instance of a semi-Markov *option*, always lasting k steps. According to the flat option representation (Precup, 2001), we have as initiation set $\mathcal{I}=\mathcal{S}$ the set of all states, as internal policy the policy that plays deter-

[6] For instance, the Q_k can be obtained by executing PFQI(k) with different persistences $k\in\mathcal{K}$.

[7] This introduces a bias that is negligible if $\|\eta^{\rho,\pi}/\nu\|_\infty\approx 1$ (details in Appendix C.1).

Table 1. Results of PFQI in different environments and persistences. For each persistence k, we report the sample mean and the standard deviation of the estimated return of the last policy $\widehat{J}_k^{\rho,\pi_k}$. For each environment, the persistence with highest average performance and the ones not statistically significantly different from that one (Welch's t-test with $p<0.05$) are in bold. The last column reports the mean and the standard deviation of the performance loss δ between the optimal persistence and the one selected by the index B_k (Equation (15)).

Environment	Expected return at persistence k ($\widehat{J}_k^{\rho,\pi_k}$, mean $\pm$ std)							Performance loss (δ mean $\pm$ std)
	$k=1$	$k=2$	$k=4$	$k=8$	$k=16$	$k=32$	$k=64$	
Cartpole	169.9 ± 5.8	176.5 ± 5.0	$\mathbf{239.5\pm4.4}$	10.0 ± 0.0	9.8 ± 0.0	9.8 ± 0.0	9.8 ± 0.0	0.0 ± 0.0
MountainCar	-111.1 ± 1.5	-103.6 ± 1.6	-97.2 ± 2.0	$\mathbf{-93.6\pm2.1}$	$\mathbf{-94.4\pm1.8}$	$\mathbf{-92.4\pm1.5}$	-136.7 ± 0.9	1.88 ± 0.85
LunarLander	-165.8 ± 50.4	-12.8 ± 4.7	$\mathbf{1.2\pm3.6}$	$\mathbf{2.0\pm3.4}$	-44.1 ± 6.9	-122.8 ± 10.5	-121.2 ± 8.6	2.12 ± 4.21
Pendulum	$\mathbf{-116.7\pm16.7}$	$\mathbf{-113.1\pm16.3}$	$\mathbf{-153.8\pm23.0}$	-283.1 ± 18.0	-338.9 ± 16.3	-364.3 ± 22.1	-377.2 ± 21.7	3.52 ± 0.0
Acrobot	-89.2 ± 1.1	$\mathbf{-82.5\pm1.7}$	$\mathbf{-83.4\pm1.3}$	-122.8 ± 1.3	-266.2 ± 1.9	-287.3 ± 0.3	-286.7 ± 0.6	0.80 ± 0.27
Swimmer	21.3 ± 1.1	$\mathbf{25.2\pm0.8}$	$\mathbf{25.0\pm0.5}$	$\mathbf{24.0\pm0.3}$	22.4 ± 0.3	12.8 ± 1.2	14.0 ± 0.2	2.69 ± 1.71
Hopper	58.6 ± 4.8	61.9 ± 4.2	62.2 ± 1.7	59.7 ± 3.1	60.8 ± 1.0	66.7 ± 2.7	$\mathbf{73.4\pm1.2}$	5.33 ± 2.32
Walker 2D	61.6 ± 5.5	37.6 ± 4.0	62.7 ± 18.2	$\mathbf{80.8\pm6.6}$	$\mathbf{102.1\pm19.3}$	$\mathbf{91.5\pm13.0}$	$\mathbf{97.2\pm17.6}$	5.10 ± 3.74

ministically the action taken when the option was initiated, i.e., the k–persistent policy, and as termination condition whether k timesteps have passed after the option started, i.e., $\beta(H_t)=\mathbb{1}_{\{t \bmod k=0\}}$. Interestingly, in Mann et al. (2015) an approximate value iteration procedure for options lasting at least a given number of steps is proposed and analyzed. This approach shares some similarities with action persistence. Nevertheless, we believe that the option framework is more general and usually the time abstractions are related to the semantic of the tasks, rather than based on the modification of the control frequency, like action persistence.

8. Experimental Evaluation

In this section, we provide the empirical evaluation of PFQI, with the threefold goal: i) proving that a persistence $k>1$ can boost learning, leading to more profitable policies, ii) assessing the quality of our persistence selection method, and iii) studying how the batch size influences the performance of PFQI policies for different persistences. Refer to Appendix D for detailed experimental settings.

We train PFQI, using extra-trees (Geurts et al., 2006) as a regression model, for J iterations and different values of k, starting with the same dataset $\mathcal{D}$ collected at persistence 1. To compare the performance of the learned policies π_k at the different persistences, we estimate their expected return J_k^{ρ,π_k} in the corresponding MDP $\mathcal{M}_k$. Table 1 shows the results for different continuous environments and different persistences averaged over 20 runs and highlighting in bold the persistence with the highest average performance and the ones that are not statistically significantly different from that one. Across the different environments we observe some common trends in line with our theory: i) persistence 1 rarely leads to the best performance; ii) excessively increasing persistence prevents the control at all. In Cartpole (Barto et al., 1983), we easily identify a persistence ($k=4$) that outperforms all the others. In the Lunar Lander (Brockman et al., 2016) persistences $k\in\{4,8\}$ are the only ones that lead to positive return (i.e., the lander does not crash) and in the Acrobot domain (Geramifard et al., 2015) we identify $k\in\{2,4\}$ as optimal persistences. A qualitatively different behavior is displayed in Mountain Car (Moore, 1991), Pendulum (Brockman et al., 2016), and Swimmer (Coulom, 2002), where we observe a plateau of three persistences with similar performance. An explanation for this phenomenon is that, in those domains, the optimal policy tends to persist actions on its own, making the difference less evident. Intriguingly, the more complex Mujoco domains, like Hopper and Walker 2D (Erickson et al., 2019), seem to benefit from the higher persistences.

To test the quality of our persistence selection method, we compare the performance of the estimated optimal persistence, i.e., the one with the highest estimated expected return $\widehat{k}\in\operatorname{argmax}\widehat{J}_k^{\rho,\pi_k}$, and the performance of the persistence $\widetilde{k}$ selected by maximizing the index B_k (Equation (15)). For each run $i=1,\dots,20$, we compute the *performance loss* $\delta_i=\widehat{J}_{\widehat{k}}^{\rho,\pi_{\widehat{k}}}-\widehat{J}_{\widetilde{k}_i}^{\rho,\pi_{\widetilde{k}_i}}$ and we report it in the last column of Table 1. In the Cartpole experiment, we observe a zero loss, which means that our heuristic always selects the optimal persistence ($k=4$). Differently, non–zero loss occurs in the other domains, which means that sometimes the index B_k mispredicts the optimal persistence. Nevertheless, in almost all cases the average performance loss is significantly smaller than the magnitude of the return, proving the effectiveness of our heuristics.

In Figure 2, we show the learning curves for the Cartpole experiment, highlighting the components that contribute to the index B_k. The first plot reports the estimated *expected return* $\widehat{J}_k^{\rho,\pi_k}$, obtained by averaging 10 trajectories executing π_k in the environment $\mathcal{M}_k$, which confirms that $k=4$ is the optimal persistence. The second plot shows the *estimated return* $\widehat{J}_k^{\rho}$ obtained by averaging the Q-function Q_k learned with PFQI(k), over the initial states sampled from ρ. We can see that for $k\in\{1,2\}$, PFQI(k) tends to overestimate the return, while for $k=4$ we notice a slight underestimation. The overestimation phenomenon can be explained by the fact that with small persistences we perform a large number

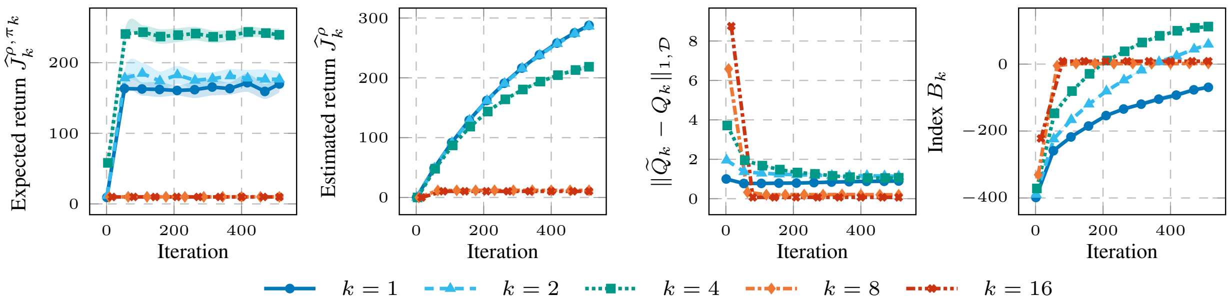

Figure 2. Expected return $\widehat{J}_k^{\rho,\pi_k}$, estimated return $\widehat{J}_k^{\rho}$, estimated expected Bellman residual $\|\widetilde{Q}_k - Q_k\|_{1,\mathcal{D}}$, and persistence selection index B_k in the Cartpole experiment as a function of the number of iterations for different persistences. 20 runs, 95 % c.i.

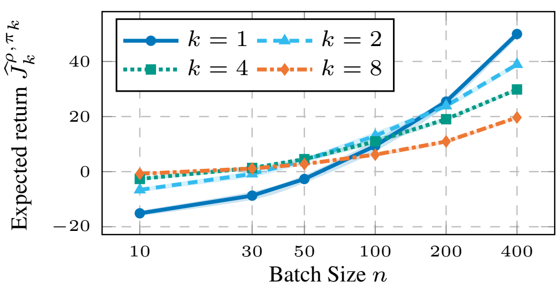

Figure 3. Expected return $\widehat{J}_k^{\rho,\pi_k}$ in the Trading experiment as a function of the batch size. 10 runs, 95 % c.i.

of applications of the operator $\widehat{T}^*$, which involves a maximization over the action space, injecting an overestimation bias. By combining this curve with the expected Bellman residual (third plot), we get the value of our persistence selection index B_k (fourth plot). Finally, we observe that B_k correctly ranks persistences 4 and 8, but overestimates persistences 8 and 16, compared to persistence 1.

To analyze the effect of the batch size, we run PFQI on the Trading environment (see Appendix D.4) varying the number of sampled trajectories. In Figure 3, we notice that the performance improves as the batch size increases, for all persistences. Moreover, we observe that if the batch size is small ($n \in \{10, 30, 50\}$), higher persistences ($k \in \{2, 4, 8\}$) result in better performances, while for larger batch sizes, $k = 1$ becomes the best choice. Since data is taken from real market prices, this environment is very noisy, thus, when the amount of samples is limited, PFQI can exploit higher persistences to mitigate the poor estimation.

9. Open Questions

Improving Exploration with Persistence We analyzed the effect of action persistence on FQI with a fixed dataset, collected in the base MDP $\mathcal{M}$. In principle, samples can be collected at arbitrary persistence. We may wonder how well the same sampling policy (e.g., the uniform policy over $\mathcal{A}$), executed at different persistences, explores the environment. For instance, in Mountain Car, high persistences increase the probability of reaching the goal, generating more informative datasets (preliminary results in Appendix E.1).

Learn in $\mathcal{M}_k$ and execute in $\mathcal{M}_{k'}$ Deploying each policy π_k in the corresponding MDP $\mathcal{M}_k$ allows for some guarantees (Lemma 6.1). However, we empirically discovered that using π_k in an MDP $\mathcal{M}_{k'}$ with smaller persistence k' sometimes improves its performance. (preliminary results in Appendix E.2). We wonder what regularity conditions on the environment are needed to explain this phenomenon.

Persistence in On–line RL Our approach focuses on batch off–line RL. However, the on–line framework could open up new opportunities for action persistence. Specifically, we could *dynamically* adapt the persistence (and so the control frequency) to speed up learning. Intuition suggests that we should start with a low frequency, reaching a fairly good policy with few samples, and then increase it to refine the learned policy.

10. Discussion and Conclusions

In this paper, we formalized the notion of action persistence, i.e., the repetition of a single action for a fixed number k of decision epochs, having the effect of altering the control frequency of the system. We have shown that persistence leads to the definition of new Bellman operators and that we are able to bound the induced performance loss, under some regularity conditions on the MDP. Based on these considerations, we presented and analyzed a novel batch RL algorithm, PFQI, able to approximate the value function at a given persistence. The experimental evaluation justifies the introduction of persistence, since reducing the control frequency can lead to an improvement when dealing with a limited number of samples. Furthermore, we introduced a persistence selection heuristic, which is able to identify good persistence in most cases. We believe that our work makes a step towards understanding why repeating actions may be useful for solving complex control tasks. Numerous questions remain unanswered, leading to several appealing future research directions.

Acknowledgements

The research was conducted under a cooperative agreement between ISI Foundation, Banca IMI and Intesa Sanpaolo Innovation Center.

References

Antos, A., Szepesvári, C., and Munos, R. Learning near-optimal policies with bellman-residual minimization based fitted policy iteration and a single sample path. *Machine Learning*, 71(1):89–129, 2008. doi: 10.1007/s10994-007-5038-2.

Bacon, P., Harb, J., and Precup, D. The option-critic architecture. In Singh, S. P. and Markovitch, S. (eds.), *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pp. 1726–1734. AAAI Press, 2017.

Baird, L. C. Reinforcement learning in continuous time: Advantage updating. In *Proceedings of 1994 IEEE International Conference on Neural Networks (ICNN'94)*, volume 4, pp. 2448–2453. IEEE, 1994.

Barto, A. G., Sutton, R. S., and Anderson, C. W. Neuronlike adaptive elements that can solve difficult learning control problems. *IEEE Trans. Systems, Man, and Cybernetics*, 13(5):834–846, 1983. doi: 10.1109/TSMC.1983.6313077.

Bertsekas, D. P. *Dynamic programming and optimal control, 3rd Edition*. Athena Scientific, 2005. ISBN 1886529264.

Bertsekas, D. P. and Shreve, S. *Stochastic optimal control: the discrete-time case*. 2004.

Bradtke, S. J. and Duff, M. O. Reinforcement learning methods for continuous-time markov decision problems. In Tesauro, G., Touretzky, D. S., and Leen, T. K. (eds.), *Advances in Neural Information Processing Systems 7, [NIPS Conference, Denver, Colorado, USA, 1994]*, pp. 393–400. MIT Press, 1994.

Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. Openai gym, 2016.

Coulom, R. *Reinforcement Learning Using Neural Networks, with Applications to Motor Control. (Apprentissage par renforcement utilisant des réseaux de neurones, avec des applications au contrôle moteur)*. PhD thesis, Grenoble Institute of Technology, France, 2002.

Dayan, P. and Singh, S. P. Improving policies without measuring merits. In Touretzky, D. S., Mozer, M., and Hasselmo, M. E. (eds.), *Advances in Neural Information Processing Systems 8, NIPS, Denver, CO, USA, November 27-30, 1995*, pp. 1059–1065. MIT Press, 1995.

Dietterich, T. G. The MAXQ method for hierarchical reinforcement learning. In Shavlik, J. W. (ed.), *Proceedings of the Fifteenth International Conference on Machine Learning (ICML 1998), Madison, Wisconsin, USA, July 24-27, 1998*, pp. 118–126. Morgan Kaufmann, 1998.

Doya, K. Reinforcement learning in continuous time and space. *Neural Computation*, 12(1):219–245, 2000. doi: 10.1162/089976600300015961.

Erickson, Z. M., Gangaram, V., Kapusta, A., Liu, C. K., and Kemp, C. C. Assistive gym: A physics simulation framework for assistive robotics. *CoRR*, abs/1910.04700, 2019.

Ernst, D., Geurts, P., and Wehenkel, L. Tree-based batch mode reinforcement learning. *J. Mach. Learn. Res.*, 6: 503–556, 2005.

Farahmand, A., Nikovski, D. N., Igarashi, Y., and Konaka, H. Truncated approximate dynamic programming with task-dependent terminal value. In Schuurmans, D. and Wellman, M. P. (eds.), *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*, pp. 3123–3129. AAAI Press, 2016.

Farahmand, A. M. *Regularization in Reinforcement Learning*. PhD thesis, University of Alberta, 2011.

Farahmand, A. M. and Szepesvári, C. Model selection in reinforcement learning. *Machine Learning*, 85(3):299–332, 2011. doi: 10.1007/s10994-011-5254-7.

Fleming, W. H. and Soner, H. M. *Controlled Markov processes and viscosity solutions*, volume 25. Springer Science & Business Media, 2006.

Geramifard, A., Dann, C., Klein, R. H., Dabney, W., and How, J. P. Rlpy: a value-function-based reinforcement learning framework for education and research. *J. Mach. Learn. Res.*, 16:1573–1578, 2015.

Geurts, P., Ernst, D., and Wehenkel, L. Extremely randomized trees. *Machine Learning*, 63(1):3–42, 2006. doi: 10.1007/s10994-006-6226-1.

Györfi, L., Kohler, M., Krzyzak, A., and Walk, H. *A Distribution-Free Theory of Nonparametric Regression*. Springer series in statistics. Springer, 2002. ISBN 978-0-387-95441-7. doi: 10.1007/b97848.

Haarnoja, T., Ha, S., Zhou, A., Tan, J., Tucker, G., and Levine, S. Learning to walk via deep reinforcement learning. In Bicchi, A., Kress-Gazit, H., and Hutchinson, S. (eds.), *Robotics: Science and Systems XV, University of Freiburg, Freiburg im Breisgau, Germany, June 22-26, 2019*, 2019. doi: 10.15607/RSS.2019.XV.011.

Howard, R. A. Semi-markov decision-processes. *Bulletin of the International Statistical Institute*, 40(2):625–652, 1963.

Jiang, N., Kulesza, A., Singh, S. P., and Lewis, R. L. The dependence of effective planning horizon on model accuracy. In Kambhampati, S. (ed.), *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, pp. 4180–4189. IJCAI/AAAI Press, 2016.

Kilinc, O., Hu, Y., and Montana, G. Reinforcement learning for robotic manipulation using simulated locomotion demonstrations. *CoRR*, abs/1910.07294, 2019.

Kober, J. and Peters, J. *Learning Motor Skills - From Algorithms to Robot Experiments*, volume 97 of *Springer Tracts in Advanced Robotics*. Springer, 2014. ISBN 978-3-319-03193-4. doi: 10.1007/978-3-319-03194-1.

Lakshminarayanan, A. S., Sharma, S., and Ravindran, B. Dynamic action repetition for deep reinforcement learning. In Singh, S. P. and Markovitch, S. (eds.), *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pp. 2133–2139. AAAI Press, 2017.

Lange, S., Gabel, T., and Riedmiller, M. A. Batch reinforcement learning. In Wiering, M. and van Otterlo, M. (eds.), *Reinforcement Learning*, volume 12 of *Adaptation, Learning, and Optimization*, pp. 45–73. Springer, 2012. doi: 10.1007/978-3-642-27645-3_2.

Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. Continuous control with deep reinforcement learning. In Bengio, Y. and LeCun, Y. (eds.), *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016.

Luenberger, D. G. Introduction to dynamic systems; theory, models, and applications. Technical report, New York: John Wiley & Sons, 1979.

Mann, T. A., Mannor, S., and Precup, D. Approximate value iteration with temporally extended actions. *J. Artif. Intell. Res.*, 53:375–438, 2015. doi: 10.1613/jair.4676.

Metelli, A. M., Mutti, M., and Restelli, M. Configurable markov decision processes. In Dy, J. G. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pp. 3488–3497. PMLR, 2018.

Metelli, A. M., Ghelfi, E., and Restelli, M. Reinforcement learning in configurable continuous environments. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pp. 4546–4555. PMLR, 2019.

Moore, A. W. Efficient memory based learning for robot control. *PhD Thesis, Computer Laboratory, University of Cambridge*, 1991.

Müller, A. Integral probability metrics and their generating classes of functions. *Advances in Applied Probability*, 29 (2):429–443, 1997.

Munos, R. A convergent reinforcement learning algorithm in the continuous case based on a finite difference method. In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence, IJCAI 97, Nagoya, Japan, August 23-29, 1997, 2 Volumes*, pp. 826–831. Morgan Kaufmann, 1997.

Munos, R. Performance bounds in l_p-norm for approximate value iteration. *SIAM journal on control and optimization*, 46(2):541–561, 2007.

Munos, R. and Bourgine, P. Reinforcement learning for continuous stochastic control problems. In Jordan, M. I., Kearns, M. J., and Solla, S. A. (eds.), *Advances in Neural Information Processing Systems 10, [NIPS Conference, Denver, Colorado, USA, 1997]*, pp. 1029–1035. The MIT Press, 1997.

Munos, R. and Szepesvári, C. Finite-time bounds for fitted value iteration. *J. Mach. Learn. Res.*, 9:815–857, 2008.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. Scikit-learn: Machine learning in Python. *J. Mach. Learn. Res.*, 12:2825–2830, 2011.

Peters, J. and Schaal, S. Reinforcement learning of motor skills with policy gradients. *Neural Networks*, 21(4): 682–697, 2008. doi: 10.1016/j.neunet.2008.02.003.

Peterson, J. K. On-line estimation of the optimal value function: Hjb-estimators. In *Advances in Neural Information Processing Systems*, pp. 319–326, 1993.

Petrik, M. and Scherrer, B. Biasing approximate dynamic programming with a lower discount factor. In Koller, D., Schuurmans, D., Bengio, Y., and Bottou, L. (eds.), *Advances in Neural Information Processing Systems 21, Proceedings of the Twenty-Second Annual Conference on Neural Information Processing Systems, Vancouver,* *British Columbia, Canada, December 8-11, 2008*, pp. 1265–1272. Curran Associates, Inc., 2008.

Pirotta, M., Restelli, M., and Bascetta, L. Policy gradient in lipschitz markov decision processes. *Machine Learning*, 100(2-3):255–283, 2015. doi: 10.1007/s10994-015-5484-1.

Precup, D. *Temporal abstraction in reinforcement learning*. PhD thesis, University of Massachusetts Amherst, 2001.

Puterman, M. L. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, 2014.

Rachelson, E. and Lagoudakis, M. G. On the locality of action domination in sequential decision making. In *International Symposium on Artificial Intelligence and Mathematics, ISAIM 2010, Fort Lauderdale, Florida, USA, January 6-8, 2010*, 2010.

Schulman, J., Levine, S., Abbeel, P., Jordan, M. I., and Moritz, P. Trust region policy optimization. In Bach, F. R. and Blei, D. M. (eds.), *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, volume 37 of *JMLR Workshop and Conference Proceedings*, pp. 1889–1897. JMLR.org, 2015.

Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *CoRR*, abs/1707.06347, 2017.

Singh, S. P. Reinforcement learning with a hierarchy of abstract models. In *Proceedings of the National Conference on Artificial Intelligence*, number 10, pp. 202. JOHN WILEY & SONS LTD, 1992a.

Singh, S. P. Scaling reinforcement learning algorithms by learning variable temporal resolution models. In *Machine Learning Proceedings 1992*, pp. 406–415. Elsevier, 1992b.

Sutton, R. S. and Barto, A. G. *Reinforcement learning: An introduction*. MIT press, 2018.

Sutton, R. S., McAllester, D. A., Singh, S. P., and Mansour, Y. Policy gradient methods for reinforcement learning with function approximation. In Solla, S. A., Leen, T. K., and Müller, K. (eds.), *Advances in Neural Information Processing Systems 12, [NIPS Conference, Denver, Colorado, USA, November 29 - December 4, 1999]*, pp. 1057–1063. The MIT Press, 1999a.

Sutton, R. S., Precup, D., and Singh, S. P. Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. *Artif. Intell.*, 112(1-2):181–211, 1999b. doi: 10.1016/S0004-3702(99)00052-1.

Tallec, C., Blier, L., and Ollivier, Y. Making deep q-learning methods robust to time discretization. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pp. 6096–6104. PMLR, 2019.

Villani, C. *Optimal transport: old and new*, volume 338. Springer Science & Business Media, 2008.

Watkins, C. J. C. H. *Learning from delayed rewards*. PhD thesis, King's College, University of Cambridge, 1989.

Index of the Appendix

In the following, we briefly recap the contents of the Appendix.

– Appendix A reports all proofs and derivations.

– Appendix B provides additional considerations and discussion concerning the regularity conditions for bounding the performance loss due to action persistence.

– Appendix C illustrates the motivations behind the choice we made for defining our persistence selection index.

– Appendix D presents the experimental setting, together with additional experimental results (including some experiments with neural networks as regressor).

– Appendix E reports some preliminary experiments to motivate the open questions stated in the main paper.

A. Proofs and Derivations

In this appendix, we report the proofs of all the results presented in the main paper.

A.1. Proofs of Section 3

Theorem 3.1. *Let $\mathcal{M}$ be an MDP, $k \in \mathbb{N}_{\geqslant 1}$ and $\mathcal{M}_k$ be the k-persistent MDP. Let $\pi \in \Pi$ be a Markovian stationary policy. Then, T_k^π and T_k^* can be expressed as:*

$$T_k^\pi = \left(T^\delta\right)^{k-1} T^\pi \quad \text{and} \quad T_k^* = \left(T^\delta\right)^{k-1} T^*, \tag{8}$$

where $T^\delta : \mathscr{B}(\mathcal{S}\times\mathcal{A}) \rightarrow \mathscr{B}(\mathcal{S}\times\mathcal{A})$ is the Bellman Persistent Operator*, defined for $f \in \mathscr{B}(\mathcal{S}\times\mathcal{A})$ and $(s,a) \in \mathcal{S}\times\mathcal{A}$:*

$$\left(T^\delta f\right)(s,a) = r(s,a) + \gamma \left(P^\delta f\right)(s,a). \tag{9}$$

Proof. We derive the result by explicitly writing the definitions of the k-persistent transition model P_k and k-persistent reward distribution R_k in terms of P, R and γ in the definition of the k-persistent Bellman expectation operator T_k^π. Let $f \in \mathscr{B}(\mathcal{S}\times\mathcal{A})$ and $(s,a) \in \mathcal{S}\times\mathcal{A}$:

$$\begin{aligned}
(T_k^\pi f)(s,a) &= r_k(s,a) + \gamma^k (P_k^\pi f)(s,a) \\
&= \sum_{i=0}^{k-1} \gamma^i \left((P^\delta)^i r \right)(s,a) + \gamma^k ((P^\delta)^{k-1} P^\pi f)(s,a) && \text{(P.1)} \\
&= \left(\sum_{i=0}^{k-1} \gamma^i (P^\delta)^i r + \gamma^k (P^\delta)^{k-1} P^\pi f \right)(s,a) \\
&= \left(\sum_{i=0}^{k-2} \gamma^i (P^\delta)^i r + \gamma^{k-1} (P^\delta)^{k-1} (r + \gamma P^\pi f) \right)(s,a) && \text{(P.2)} \\
&= \left(\sum_{i=0}^{k-2} \gamma^i (P^\delta)^i r + \gamma^{k-1} (P^\delta)^{k-1} T^\pi f \right)(s,a), && \text{(P.3)}
\end{aligned}$$

where line (P.1) follows from Definition 3.2, line (P.2) is obtained by isolating the last term in the summation $\gamma^{k-1}(P^\delta)^{k-1} r$ and collecting $\gamma^{k-1}(P^\delta)^{k-1}$ thanks to the linearity of $(P^\delta)^{k-1}$, and line (P.3) derives from the definition of the Bellman expectation operator T^π. It remains to prove that for $g \in \mathscr{B}(\mathcal{S}\times\mathcal{A})$ and $(s,a) \in \mathcal{S}\times\mathcal{A}$, we have the following identity:

$$(T^\delta)^{k-1} g = \sum_{i=0}^{k-2} \gamma^i (P^\delta)^i r + \gamma^{k-1} (P^\delta)^{k-1} g. \tag{P.4}$$

We prove it by induction on $k \in \mathbb{N}_{\geqslant 1}$. For $k=1$ we have only $g = (T^\delta)^0 g$. Let us assume that the identity hold for all integers $h<k$, we prove the statement for k:

$$\left((T^\delta)^{k-1} g \right)(s,a) = \left(\sum_{i=0}^{k-2} \gamma^i (P^\delta)^i r + \gamma^{k-1} (P^\delta)^{k-1} g \right)(s,a)$$

$$= \left(\sum_{i=0}^{k-3} \gamma^i (P^\delta)^i r + \gamma^{k-2} (P^\delta)^{k-2} (r + \gamma P^\delta g) \right)(s,a) \tag{P.5}$$

$$= \left(\sum_{i=0}^{k-3} \gamma^i (P^\delta)^i r + \gamma^{k-2} (P^\delta)^{k-2} T^\delta g \right)(s,a) \tag{P.6}$$

$$= (T^\delta)^{k-2} T^\delta g = (T^\delta)^{k-1} g. \tag{P.7}$$

where line (P.5) derives from isolating the last term in the summation and collecting $\gamma^{k-2}(P^\delta)^{k-2}$ thanks to the linearity of $(P^\delta)^{k-2}$, line (P.6) comes from the definition of the Bellman persisted operator T^δ, and finally line (P.7) follows from the inductive hypothesis. We get the result by taking $g = T^\pi f$.

Concerning the k-persistent Bellman optimal operator the derivation is analogous. For simplicity, we define the $\max$-operator $M : \mathscr{B}(\mathcal{S} \times \mathcal{A}) \to \mathscr{B}(\mathcal{S})$ defined for a bounded measurable function $f \in \mathscr{B}(\mathcal{S} \times \mathcal{A})$ and a state $s \in \mathcal{S}$ as $(Mf)(s) = \max_{a \in \mathcal{A}} f(s,a)$. As a consequence the Bellman optimal operator becomes: $T^* f = r + \gamma P M f$. Therefore, we have:

$$(T_k^* f)(s,a) = r_k(s,a) + \gamma^k \int_{\mathcal{S}} P_k(\mathrm{d}s'|s,a) \max_{a' \in \mathcal{A}} f(s',a')$$

$$= r_k(s,a) + \gamma^k \int_{\mathcal{S}} P_k(\mathrm{d}s'|s,a) M f(s') \tag{P.8}$$

$$= \left(r_k + \gamma^k P_k M f \right)(s,a) \tag{P.9}$$

$$= \left(\sum_{i=0}^{k-1} \gamma^i (P^\delta)^i r + \gamma^k (P^\delta)^{k-1} P M f \right)(s,a)$$

$$= \left(\sum_{i=0}^{k-2} \gamma^i (P^\delta)^i r + \gamma^{k-1} (P^\delta)^{k-1} (r + \gamma P M f) \right)(s,a) \tag{P.10}$$

$$= \left(\sum_{i=0}^{k-2} \gamma^i (P^\delta)^i r + \gamma^{k-1} (P^\delta)^{k-1} T^* f \right)(s,a), \tag{P.11}$$

where line (P.8) derives from the definition of the $\max$-operator M and line (P.8) from the definition of the operator P_k. By applying Equation (P.4) we get the result. □

A.2. Proofs of Section 4

Lemma A.1. *Let $\mathcal{M}$ be an MDP and $\pi \in \Pi$ be a Markovian stationary policy, then for any $k \in \mathbb{N}_{\geqslant 1}$ the following two identities hold:*

$$Q^\pi - Q_k^\pi = \left(\mathrm{Id} - \gamma^k (P^\pi)^k \right)^{-1} \left((T^\pi)^k Q_k^\pi - \left(T^\delta\right)^{k-1} T^\pi Q_k^\pi \right)$$

$$= \left(\mathrm{Id} - \gamma^k \left(P^\delta\right)^{k-1} P^\pi \right)^{-1} \left((T^\pi)^k Q^\pi - \left(T^\delta\right)^{k-1} T^\pi Q^\pi \right),$$

where $\mathrm{Id} : \mathscr{B}(\mathcal{S} \times \mathcal{A}) \to \mathscr{B}(\mathcal{S} \times \mathcal{A})$ is the identity operator over $\mathcal{S} \times \mathcal{A}$.

Proof. We prove the equalities by exploiting the facts that Q^π and Q_k^π are the fixed points of T^π and T_k^π:

$$Q^\pi - Q_k^\pi = T^\pi Q^\pi - T_k^\pi Q_k^\pi$$

$$= (T^\pi)^k Q^\pi - \left(T^\delta\right)^{k-1} T^\pi Q_k^\pi \tag{P.12}$$

$$= (T^\pi)^k Q^\pi - \left(T^\delta\right)^{k-1} T^\pi Q_k^\pi \pm (T^\pi)^k Q_k^\pi \tag{P.13}$$

$$= \gamma^k (P^\pi)^k (Q^\pi - Q_k^\pi) + \left((T^\pi)^k Q_k^\pi - \left(T^\delta\right)^{k-1} T^\pi Q_k^\pi \right), \tag{P.14}$$

where line (P.12) derives from recalling that $Q^\pi = T^\pi Q^\pi$ and exploiting Theorem 3.1, line (P.14) is obtained by exploiting the identity that holds for two generic bounded measurable functions $f, g \in \mathscr{B}(\mathcal{S} \times \mathcal{A})$:

$$(T^\pi)^k f - (T^\pi)^k g = \gamma^k (P^\pi)^k (f - g). \tag{P.15}$$

We prove this identity by induction. For $k=1$ the identity clearly holds. Suppose Equation (P.15) holds for all integers $h < k$, we prove that it holds for k too:

$$(T^\pi)^k f - (T^\pi)^k g = T^\pi (T^\pi)^{k-1} f - T^\pi (T^\pi)^{k-1} g$$
$$= r + \gamma P^\pi (T^\pi)^{k-1} f - r - P^\pi \gamma (T^\pi)^{k-1} g$$
$$= \gamma P^\pi \left((T^\pi)^{k-1} f - (T^\pi)^{k-1} g \right) \tag{P.16}$$
$$= \gamma P^\pi \gamma^{k-1} (P^\pi)^{k-1} (f-g) \tag{P.17}$$
$$= \gamma^k (P^\pi)^k (f-g),$$

where line (P.16) derives from the linearity of operator P^π and line (P.17) follows from the inductive hypothesis. From line (P.14) the result follows immediately, recalling that since $\gamma < 1$ the inversion of the operator is well-defined:

$$Q^\pi - Q^\pi_k = \gamma^k (P^\pi)^k (Q^\pi - Q^\pi_k) + \left((T^\pi)^k Q^\pi_k - \left(T^\delta\right)^{k-1} T^\pi Q^\pi_k \right) \implies$$
$$\left(\mathrm{Id} - \gamma^k (P^\pi)^k \right) (Q^\pi - Q^\pi_k) = \left((T^\pi)^k Q^\pi_k - \left(T^\delta\right)^{k-1} T^\pi Q^\pi_k \right) \implies$$
$$Q^\pi - Q^\pi_k = \left(\mathrm{Id} - \gamma^k (P^\pi)^k \right)^{-1} \left((T^\pi)^k Q^\pi_k - \left(T^\delta\right)^{k-1} T^\pi Q^\pi_k \right).$$

The second identity of the statement is obtained with an analogous derivation, in which at line (P.13) we sum and subtract $\left(T^\delta\right)^{k-1} T^\pi Q^\pi$ and we exploit the identity for two bounded measurable functions $f, g \in \mathscr{B}(\mathcal{S} \times \mathcal{A})$:

$$\left(T^\delta\right)^{k-1} T^\pi Q f - \left(T^\delta\right)^{k-1} T^\pi Q g = \gamma^k \left(P^\delta\right)^{k-1} P^\pi (f-g). \tag{P.18}$$

□

Lemma A.2. *Let $\mathcal{M}$ be an MDP and $\pi \in \Pi$ be a Markovian stationary policy, then for any $k \in \mathbb{N}_{\geqslant 1}$ and any bounded measurable function $f \in \mathscr{B}(\mathcal{S} \times \mathcal{A})$ the following two identities hold:*

$$(T^\pi)^{k-1} f - \left(T^\delta\right)^{k-1} f = \sum_{i=0}^{k-2} \gamma^{i+1} (P^\pi)^i \left(P^\pi - P^\delta\right) \left(T^\delta\right)^{k-2-i} f$$
$$= \sum_{i=0}^{k-2} \gamma^{i+1} \left(P^\delta\right)^i \left(P^\pi - P^\delta\right) (T^\pi)^{k-2-i} f.$$

Proof. We start with the first identity and we prove it by induction on k. For $k=1$, we have that the left hand side is zero and the summation on the right hand side has no terms. Suppose that the statement holds for every $h<k$, we prove the statement for k:

$$(T^\pi)^{k-1} f - \left(T^\delta\right)^{k-1} f = (T^\pi)^{k-1} f - \left(T^\delta\right)^{k-1} f \pm (T^\pi)^{k-2} T^\delta f \tag{P.19}$$
$$= \left((T^\pi)^{k-2} T^\pi f - (T^\pi)^{k-2} T^\delta f \right) + \left((T^\pi)^{k-2} T^\delta f - \left(T^\delta\right)^{k-2} T^\delta f \right)$$
$$= \gamma^{k-2} (P^\pi)^{k-2} \left(T^\pi f - T^\delta f \right) + \left((T^\pi)^{k-2} T^\delta f - \left(T^\delta\right)^{k-2} T^\delta f \right) \tag{P.20}$$
$$= \gamma^{k-1} (P^\pi)^{k-2} \left(P^\pi - P^\delta \right) f + \sum_{i=0}^{k-3} \gamma^{i+1} (P^\pi)^i \left(P^\pi - P^\delta \right) \left(T^\delta\right)^{k-3-i} T^\delta f \tag{P.21}$$
$$= \sum_{i=0}^{k-2} \gamma^{i+1} (P^\pi)^i \left(P^\pi - P^\delta \right) \left(T^\delta\right)^{k-2-i} f, \tag{P.22}$$

where in line (P.20) we exploited the identity at Equation (P.15), line (P.21) derives from observing that $T^\pi f - T^\delta f = \gamma \left(P^\pi - P^\delta\right) f$ and by inductive hypothesis applied on $T^\delta f$ which is a bounded measurable function as well. Finally, line (P.22) follows from observing that the first term completes the summation up to $k-2$. The second identity in the statement can be obtained by an analogous derivation in which at line (P.19) we sum and subtract $\left(T^\delta\right)^{k-2} T^\pi f$ and, later, exploit the identity at Equation (P.18). □

Lemma A.3 (Persistence Lemma)**.** *Let $\mathcal{M}$ be an MDP and $\pi \in \Pi$ be a Markovian stationary policy, then for any $k \in \mathbb{N}_{\geqslant 1}$ the following two identities hold:*

$$
\begin{aligned}
Q^\pi - Q^\pi_k &= \sum_{\substack{i \in \mathbb{N} \\ i \bmod k \neq 0}} \gamma^i (P^\pi)^{i-1} \left(P^\pi - P^\delta\right) \left(T^\delta\right)^{k-2-(i-1) \bmod k} T^\pi Q^\pi_k \\
&= \sum_{\substack{i \in \mathbb{N} \\ i \bmod k \neq 0}} \gamma^i \left(\left(P^\delta\right)^{k-1} P^\pi \right)^{i \operatorname{div} k} \left(P^\delta\right)^{i \bmod k - 1} \left(P^\pi - P^\delta\right) (T^\pi)^{k - i \bmod k} Q^\pi,
\end{aligned}
$$

where for two non-negative integers $a,b \in \mathbb{N}$, we denote with $a \bmod b$ and $a \operatorname{div} b$ the remainder and the quotient of the integer division between a and b respectively.

Proof. We start proving the first identity. Let us consider the first identity of Lemma A.1:

$$
\begin{aligned}
Q^\pi - Q^\pi_k &= \left(\mathrm{Id} - \gamma^k (P^\pi)^k\right)^{-1} \left((T^\pi)^k Q^\pi_k - \left(T^\delta\right)^{k-1} T^\pi Q^\pi_k \right) \\
&= \left(\sum_{j=0}^{+\infty} \gamma^{kj} (P^\pi)^{kj} \right) \left((T^\pi)^k Q^\pi_k - \left(T^\delta\right)^{k-1} T^\pi Q^\pi_k \right) && \text{(P.23)} \\
&= \left(\sum_{j=0}^{+\infty} \gamma^{kj} (P^\pi)^{kj} \right) \sum_{l=0}^{k-2} \gamma^{l+1} (P^\pi)^l \left(P^\pi - P^\delta\right) \left(T^\delta\right)^{k-2-l} T^\pi Q^\pi_k && \text{(P.24)} \\
&= \sum_{j=0}^{+\infty} \gamma^{kj} (P^\pi)^{kj} \sum_{l=0}^{k-2} \gamma^{l+1} (P^\pi)^l \left(P^\pi - P^\delta\right) \left(T^\delta\right)^{k-2-l} T^\pi Q^\pi_k \\
&= \sum_{j=0}^{+\infty} \sum_{l=0}^{k-2} \gamma^{kj+l+1} (P^\pi)^{kj+l} \left(P^\pi - P^\delta\right) \left(T^\delta\right)^{k-2-l} T^\pi Q^\pi_k,
\end{aligned}
$$

where line (P.23) follows from applying the Neumann series at the first factor, line (P.24) is obtained by applying the first identity of Lemma A.2 to the bounded measurable function $T^\pi Q^\pi_k$. The subsequent lines are obtained by straightforward algebraic manipulations. Now we rename the indexes by setting $i = kj + l + 1$. Since $l \in \{0, \dots, k-2\}$ we have that $j = (i-1) \operatorname{div} k$ and $l = (i-1) \bmod k$. Moreover, we observe that i ranges over all non-negative integers values except for the multiples of the persistence k, i.e., $i \in \{n \in \mathbb{N} : n \bmod k \neq 0\}$. Now, recalling that $i \bmod k \neq 0$, we observe that for the distributive property of the modulo operator we have $(i-1) \bmod k = (i \bmod k - 1 \bmod k) \bmod k = (i \bmod k - 1) \bmod k = i \bmod k - 1$. The second identity is obtained by an analogous derivation in which we exploit the second identities at Lemmas A.1 and A.2. ☐

Theorem 4.1. *Let $\mathcal{M}$ be an MDP and $\pi \in \Pi$ be a Markovian stationary policy. Let $\mathcal{Q}_k = \{ \left(T^\delta\right)^{k-2-l} T^\pi Q^\pi_k : l \in \{0, \dots, k-2\}\}$ and for all $(s,a) \in \mathcal{S} \times \mathcal{A}$ let us define:*

$$
d^\pi_{\mathcal{Q}_k}(s,a) = \sup_{f \in \mathcal{Q}_k} \left| \int_{\mathcal{S}} \int_{\mathcal{A}} \left(P^\pi(\mathrm{d}s', \mathrm{d}a' | s,a) - P^\delta(\mathrm{d}s', \mathrm{d}a' | s,a) \right) f(s',a') \right|.
$$

Then, for any $\rho \in \mathscr{P}(\mathcal{S} \times \mathcal{A})$, $p \geqslant 1$, and $k \in \mathbb{N}_{\geqslant 1}$, it holds that:

$$
\| Q^\pi - Q^\pi_k \|_{p,\rho} \leqslant \frac{\gamma(1-\gamma^{k-1})}{(1-\gamma)(1-\gamma^k)} \left\| d^\pi_{\mathcal{Q}_k} \right\|_{p, \eta^{\rho,\pi}_k},
$$

where $\eta^{\rho,\pi}_k \in \mathscr{P}(\mathcal{S} \times \mathcal{A})$ is a probability measure defined for any measurable set $\mathcal{B} \subseteq \mathcal{S} \times \mathcal{A}$ as:

$$
\eta^{\rho,\pi}_k(\mathcal{B}) = \frac{(1-\gamma)(1-\gamma^k)}{\gamma(1-\gamma^{k-1})} \sum_{\substack{i \in \mathbb{N} \\ i \bmod k \neq 0}} \gamma^i \left(\rho (P^\pi)^{i-1} \right) (\mathcal{B}).
$$

Proof. We start from the first equality derived in Lemma A.3, and we apply the $L_p(\rho)$-norm both sides, with $p \geqslant 1$:

$$
\| Q^\pi - Q^\pi_k \|^p_{p,\rho} = \left\| \sum_{\substack{i \in \mathbb{N} \\ i \bmod k \neq 0}} \gamma^i (P^\pi)^{i-1} \left(P^\pi - P^\delta\right) \left(T^\delta\right)^{k-2-(i-1) \bmod k} T^\pi Q^\pi_k \right\|^p_{p,\rho}
$$

$$=\rho\left|\sum_{\substack{i\in\mathbb{N}\\ i \bmod k\neq 0}}\gamma^i (P^\pi)^{i-1}\left(P^\pi - P^\delta\right)\left(T^\delta\right)^{k-2-(i-1) \bmod k} T^\pi Q_k^\pi\right|^p \tag{P.25}$$

$$\leqslant\rho\left|\sum_{\substack{i\in\mathbb{N}\\ i \bmod k\neq 0}}\gamma^i (P^\pi)^{i-1}\sup_{f\in\mathcal{Q}_k}\left|\left(P^\pi - P^\delta\right) f\right|\right|^p \tag{P.26}$$

$$=\left(\frac{\gamma(1-\gamma^{k-1})}{(1-\gamma)(1-\gamma^k)}\right)^p \rho\left|\frac{(1-\gamma)(1-\gamma^k)}{\gamma(1-\gamma^{k-1})}\sum_{\substack{i\in\mathbb{N}\\ i \bmod k\neq 0}}\gamma^i (P^\pi)^{i-1} d^\pi_{\mathcal{Q}_k}\right|^p \tag{P.27}$$

$$\leqslant\left(\frac{\gamma(1-\gamma^{k-1})}{(1-\gamma)(1-\gamma^k)}\right)^p \frac{(1-\gamma)(1-\gamma^k)}{\gamma(1-\gamma^{k-1})}\rho\sum_{\substack{i\in\mathbb{N}\\ i \bmod k\neq 0}}\gamma^i (P^\pi)^{i-1} \left|d^\pi_{\mathcal{Q}_k}\right|^p \tag{P.28}$$

$$=\left(\frac{\gamma(1-\gamma^{k-1})}{(1-\gamma)(1-\gamma^k)}\right)^p \eta_k^{\rho,\pi}\left|d^\pi_{\mathcal{Q}_k}\right|^p \tag{P.29}$$

$$=\left(\frac{\gamma(1-\gamma^{k-1})}{(1-\gamma)(1-\gamma^k)}\right)^p \left\|d^\pi_{\mathcal{Q}_k}\right\|^p_{p,\eta^{\rho,\pi}}. \tag{P.30}$$

where line (P.25) is obtained by the definition of norm, written in the operator form, line (P.26) is obtained by bounding $\left(P^\pi - P^\delta\right)\left(T^\delta\right)^{k-2-(i-1) \bmod k} \leqslant \sup_{f\in\mathcal{Q}_k}\left|\left(P^\pi - P^\delta\right) f\right|$, recalling the definition of $\mathcal{Q}_k$ and that $(i-1) \bmod k \leqslant k-2$ for all $i\in\mathbb{N}$ and $i \bmod k \neq 0$. Then, line (P.27) follows from deriving the normalization constant in order to make the summation $\sum_{\substack{i\in\mathbb{N}\\ i \bmod k\neq 0}}\gamma^i (P^\pi)^{i-1}$ a proper probability distribution. Such a constant can be obtained as follows:

$$\sum_{\substack{i\in\mathbb{N}\\ i \bmod k\neq 0}}\gamma^i = \sum_{i\in\mathbb{N}}\gamma^i - \sum_{i\in\mathbb{N}}\gamma^{ki} = \frac{\gamma(1-\gamma^{k-1})}{(1-\gamma)(1-\gamma^k)}.$$

Line (P.28) is obtained by applying Jensen inequality recalling that $p\geqslant 1$. Finally, line (P.29) derives from the definition of the distribution $\eta_k^{\rho,\pi}$ and line (P.30) from the definition of $L_p(\eta_k^{\rho,\pi})$-norm. □

Lemma A.4. *Let $\mathcal{M}$ be an MDP and $\pi\in\Pi$ be a Markovian stationary policy. Let $f\in\mathscr{B}(\mathcal{S}\times\mathcal{A})$ that is L_f–LC. Then, under Assumptions 2.1 and 2.2, the following statements hold:*

i) $T^\pi f$ is $(L_r+\gamma L_P(L_\pi+1)L_f)$–LC;

ii) $T^\delta f$ is $(L_r+\gamma(L_P+1)L_f)$–LC;

iii) $T^* f$ is $(L_r+\gamma L_P L_f)$–LC.

Proof. Let $f\in\mathscr{B}(\mathcal{S}\times\mathcal{A})$ be L_f-LC. Consider an application of T^π and $(s,a),(\overline{s},\overline{a})\in\mathcal{S}\times\mathcal{A}$:

$$|(T^\pi f)(s,a)-(T^\pi f)(\overline{s},\overline{a})| = \left|r(s,a)+\gamma\int_{\mathcal{S}}\int_{\mathcal{A}}P(\mathrm{d}s'|s,a)\pi(\mathrm{d}a'|s')f(s',a') - r(\overline{s},\overline{a}) - \gamma\int_{\mathcal{S}}\int_{\mathcal{A}}P(\mathrm{d}s'|\overline{s},\overline{a})\pi(\mathrm{d}a'|s')f(s',a')\right|$$

$$\leqslant |r(s,a)-r(\overline{s},\overline{a})| + \gamma\left|\int_{\mathcal{S}}\left(P(\mathrm{d}s'|s,a)-P(\mathrm{d}s'|\overline{s},\overline{a})\right)\int_{\mathcal{A}}\pi(\mathrm{d}a'|s')f(s',a')\right| \tag{P.31}$$

$$\leqslant |r(s,a)-r(\overline{s},\overline{a})| + \gamma(L_\pi+1)L_f \sup_{f:\|f\|_L\leqslant 1}\left|\int_{\mathcal{S}}\left(P(\mathrm{d}s'|s,a)-P(\mathrm{d}s'|\overline{s},\overline{a})\right)f(s')\right| \tag{P.32}$$

$$\leqslant (L_r+\gamma L_P(L_\pi+1)L_f)\, d_{\mathcal{S}\times\mathcal{A}}((s,a),(\overline{s},\overline{a})), \tag{P.33}$$

where line (P.31) follows from triangular inequality, line (P.32) is obtained from observing that the function $g_f(s') = \int_{\mathcal{A}}\pi(\mathrm{d}a'|s')f(s',a')$ is $(L_\pi+1)L_f$–LC, since for any $s,\overline{s}\in\mathcal{S}$:

$$|g_f(s)-g_f(\overline{s})| = \left|\int_{\mathcal{A}}\pi(\mathrm{d}a|s)f(s,a) - \int_{\mathcal{A}}\pi(\mathrm{d}a|\overline{s})f(\overline{s},a)\right|$$

$$= \left|\int_{\mathcal{A}}\pi(\mathrm{d}a|s)f(s,a) - \int_{\mathcal{A}}\pi(\mathrm{d}a|\overline{s})f(\overline{s},a) \pm \int_{\mathcal{A}}\pi(\mathrm{d}a|\overline{s})f(s,a)\right|$$

$$\leqslant \left| \int_{\mathcal{A}} (\pi(\mathrm{d}a|s) - \pi(\mathrm{d}a|\overline{s})) f(s,a) \right| + \left| \int_{\mathcal{A}} \pi(\mathrm{d}a|\overline{s}) (f(\overline{s},a) - f(s,a)) \right|$$
$$\leqslant L_f \sup_{f: \|f\|_L \leqslant 1} \left| \int_{\mathcal{A}} (\pi(\mathrm{d}a|s) - \pi(\mathrm{d}a|\overline{s})) f(a) \right| + \left| \int_{\mathcal{A}} \pi(\mathrm{d}a|\overline{s}) (f(\overline{s},a) - f(s,a)) \right|$$
$$\leqslant L_f L_\pi d_{\mathcal{S}}(s,\overline{s}) + L_f d_{\mathcal{S}}(s,\overline{s}),$$

where we exploited the fact that L_π–LC. Finally, line (P.33) is obtained by recalling that the reward function is L_r–LC and the transition model is L_P–LC. The derivations are analogous for T^δ and T^*. Concerning T^δ we have:

$$\begin{aligned}\left|(T^\delta f)(s,a) - (T^\delta f)(\overline{s},\overline{a})\right| &\leqslant |r(s,a) - r(\overline{s},\overline{a})| + \gamma \left| \int_{\mathcal{S}} \int_{\mathcal{A}} \left(\delta_a(\mathrm{d}a') P(\mathrm{d}s'|s,a) - \delta_{\overline{a}}(\mathrm{d}a') P(\mathrm{d}s'|\overline{s},\overline{a})\right) f(s',a') \right| \\ &\leqslant L_r d_{\mathcal{S}\times\mathcal{A}}((s,a),(\overline{s},\overline{a})) + \gamma \left| \int_{\mathcal{S}} \left(P(\mathrm{d}s'|s,a) - P(\mathrm{d}s'|\overline{s},\overline{a})\right) \int_{\mathcal{A}} \delta_a(\mathrm{d}a') f(s',a') \right| \\ &\quad + \gamma \int_{\mathcal{S}} P(\mathrm{d}s'|\overline{s},\overline{a}) \left| \int_{\mathcal{A}} \left(\delta_a(\mathrm{d}a') - \delta_{\overline{a}}(\mathrm{d}a')\right) f(s',a') \right| \\ &\leqslant (L_r + \gamma L_f L_P + \gamma L_f) d_{\mathcal{S}\times\mathcal{A}}((s,a),(\overline{s},\overline{a})),\end{aligned}$$

where we observed that $\int_{\mathcal{A}} \delta_a(\mathrm{d}a') f(s',a') = f(s',a)$ is L_f–LC and that $\int_{\mathcal{A}} \left|\delta_a(\mathrm{d}a') - \delta_{\overline{a}}(\mathrm{d}a')\right| f(s',a') \leqslant L_f d_{\mathcal{A}}(a,\overline{a}) \leqslant L_f d_{\mathcal{S}\times\mathcal{A}}((s,a),(\overline{a},\overline{a}))$. Finally, considering T^*, we have:

$$\begin{aligned}\left|(T^* f)(s,a) - (T^* f)(\overline{s},\overline{a})\right| &\leqslant |r(s,a) - r(\overline{s},\overline{a})| + \gamma \left| \int_{\mathcal{S}} \left(P(\mathrm{d}s'|s,a) - P(\mathrm{d}s'|\overline{s},\overline{a})\right) \max_{a' \in As} f(s',a') \right| \\ &\leqslant (L_r + \gamma L_f L_P) d_{\mathcal{S}\times\mathcal{A}}((s,a),(\overline{s},\overline{a})),\end{aligned}$$

where we observed that the function $h_f(s') = \max_{a' \in As} f(s',a')$ is L_f–LC, since:

$$\begin{aligned}|h_f(s) - h_f(\overline{s})| &= \left| \max_{a' \in As} f(s,a') - \max_{a' \in As} f(\overline{s},a') \right| \\ &\leqslant \max_{a' \in \mathcal{A}} \left| f(s,a') - f(\overline{s},a') \right| \\ &\leqslant L_f d_{\mathcal{S}}(s,\overline{s}).\end{aligned}$$

□

Lemma A.5. *Let $\mathcal{M}$ be an MDP and $\pi \in \Pi$ be a Markovian stationary policy. Then, under Assumptions 2.1 and 2.2, if $\gamma \max\{L_P + 1, L_P(L_\pi + 1)\} < 1$, the functions $f \in \mathcal{Q}_k$ are $L_{\mathcal{Q}_k}$–LC, where:*

$$L_{\mathcal{Q}_k} \leqslant \frac{L_r}{1 - \gamma \max\{L_P + 1, L_P(L_\pi + 1)\}}. \tag{16}$$

Furthermore, for all $(s,a) \in \mathcal{S} \times \mathcal{A}$ it holds that:

$$d_{\mathcal{Q}_k}(s,a) \leqslant L_{\mathcal{Q}_k} \mathcal{W}_1 \left(P^\pi(\cdot|s,a), P^\delta(\cdot|s,a)\right). \tag{17}$$

Proof. First of all consider the action-value function of the k–persistent MDP Q_k^π, which is the fixed point of the operator T_k^π that decomposes into $(T^\delta)^{k-1} T^\pi$ according to Theorem 3.1. It follows that for any $f \in \mathscr{B}(\mathcal{S} \times \mathcal{A})$ we have:

$$Q_k^\pi = \lim_{j \to +\infty} (T_k^\pi)^j f = \lim_{j \to +\infty} \left((T^\delta)^{k-1} T^\pi\right)^j f.$$

We now want to bound the Lipschitz constant of Q_k^π. To this purpose, let us first compute the Lipschitz constant of $T_k^\pi f = ((T^\delta)^{k-1} T^\pi) f$ for $f \in \mathscr{B}(\mathcal{S} \times \mathcal{A})$ being an L_f–LC function. From Lemma A.4 we can bound the Lipschitz constant a_h of $(T^\delta)^h T^\pi f$ for $h \in \{0, \dots k-1\}$, leading to the sequence:

$$a_h = \begin{cases} L_r + \gamma L_P (L_\pi + 1) L_f & \text{if } h = 0 \\ L_r + \gamma (L_P + 1) a_{h-1} & \text{if } h \in \{1, \dots k-1\} \end{cases}.$$

Thus, the Lipschitz constant of $((T^\delta)^{k-1} T^\pi) f$ is a_{k-1}. By unrolling the recursion we have:

$$a_{k-1} = L_r \sum_{i=0}^{k-1} \gamma^i (L_P + 1)^i + \gamma^k L_P (L_\pi + 1)(L_P + 1)^{k-1} L_f = L_r \frac{1 - \gamma^k (L_P + 1)^k}{1 - \gamma(L_P + 1)} + \gamma^k L_P (L_\pi + 1)(L_P + 1)^{k-1} L_f.$$

Let us now consider the sequence b_j of the Lipschitz constants of $(T_k^\pi)^j f$ for $j \in \mathbb{N}$:

$$b_j = \begin{cases} L_f & \text{if } j=0 \\ L_r \frac{1-\gamma^k (L_P+1)^k}{1-\gamma(L_P+1)} + \gamma^k L_P (L_\pi+1)(L_P+1)^{k-1} b_{j-1} & \text{if } j \in \mathbb{N}_{\geqslant 1} \end{cases}.$$

The sequence b_j converges to a finite limit as long as $\gamma^k L_P (L_\pi+1)(L_P+1)^{k-1} < 1$. In such case, the limit b_∞ can be computed solving the fixed point equation:

$$b_\infty = L_r \frac{1-\gamma^k (L_P+1)^k}{1-\gamma(L_P+1)} + \gamma^k L_P (L_\pi+1)(L_P+1)^{k-1} b_\infty \quad \Longrightarrow \quad b_\infty = \frac{L_r \left(1-\gamma^k (L_P+1)^k\right)}{(1-\gamma(L_P+1))(1-\gamma^k L_P (L_\pi+1)(L_P+1)^{k-1})}.$$

Thus, b_∞ represents the Lipschitz constant of Q_k^π. It is worth noting that when setting $k=1$ we recover the Lipschitz constant of the Q^π as in (Rachelson & Lagoudakis, 2010). To get a bound that is independent on k we define $L = \max\{L_P(L_\pi+1), L_P+1\}$, assuming that $\gamma L < 1$ so that:

$$b_\infty = \frac{L_r \left(1-\gamma^k (L_P+1)^k\right)}{(1-\gamma(L_P+1))(1-\gamma^k L_P (L_\pi+1)(L_P+1)^{k-1})} \leqslant \frac{L_r}{1-\gamma L},$$

having observed that $\frac{1-\gamma^k (L_P+1)^k}{1-\gamma(L_P+1)} \leqslant \frac{1-\gamma^k L^k}{1-\gamma L}$. Thus, we conclude that Q_k^π is also $\frac{L_r}{1-\gamma L}$–LC for any $k \in \mathbb{N}_{\geqslant 1}$. Consider now the application of the operator T^π to Q_k^π, we have that the corresponding Lipschitz constant can be bounded by:

$$L_{T^\pi Q_k^\pi} \leqslant L_r + \gamma L_P (L_\pi+1) \frac{L_r}{1-\gamma L} \leqslant L_r + \gamma L \frac{L_r}{1-\gamma L} = \frac{L_r}{1-\gamma L}. \tag{P.34}$$

A similar derivation holds for the application of T^δ. As a consequence, any arbitrary sequence of applications of T^π and T^δ to Q_k^π generates a sequence of $\frac{L_r}{1-\gamma L}$–LC functions. Even more so for the functions in the set $\mathcal{Q}_k = \{ (T^\delta)^{k-2-l} T^\pi Q_k^\pi : l \in \{0,\dots,k-2\}\}$. As a consequence, we can rephrase the dissimilarity term $d_{\mathcal{Q}_k}^\pi(s,a)$ as a Kantorovich distance:

$$\begin{aligned} d_{\mathcal{Q}_k}^\pi(s,a) &= \sup_{f \in \mathcal{Q}_k} \left| \int_{\mathcal{S}} \int_{\mathcal{A}} \left(P^\pi(\mathrm{d}s',\mathrm{d}a'|s,a) - P^\delta(\mathrm{d}s',\mathrm{d}a'|s,a) \right) f(s',a') \right| \\ &\leqslant L_{\mathcal{Q}_k} \sup_{f: \|f\|_L \leqslant 1} \left| \int_{\mathcal{S}} \int_{\mathcal{A}} \left(P^\pi(\mathrm{d}s',\mathrm{d}a'|s,a) - P^\delta(\mathrm{d}s',\mathrm{d}a'|s,a) \right) f(s',a') \right| \\ &= L_{\mathcal{Q}_k} \mathcal{W}_1 \left(P^\pi(\cdot|s,a), P^\delta(\cdot|s,a) \right). \end{aligned}$$

□

Theorem 4.2. *Let $\mathcal{M}$ be an MDP and $\pi \in \Pi$ be a Markovian stationary policy. Under Assumptions 2.1, 2.2, and 4.1, if $\gamma \max\{L_P+1, L_P(1+L_\pi)\} < 1$ and if $\rho(s,a) = \rho_{\mathcal{S}}(s)\pi(a|s)$ with $\rho_{\mathcal{S}} \in \mathscr{P}(\mathcal{S})$, then for any $k \in \mathbb{N}_{\geqslant 1}$:*

$$\left\| d_{\mathcal{Q}_k}^\pi \right\|_{p,\eta_k^{\rho,\pi}} \leqslant L_{\mathcal{Q}_k} \left[(L_\pi+1) L_T + \sigma_p \right].$$

where $\sigma_p^p = \sup_{s \in \mathcal{S}} \int_{\mathcal{A}} \int_{\mathcal{A}} d_{\mathcal{A}}(a,a')^p \pi(\mathrm{d}a|s)\pi(\mathrm{d}a'|s)$, and $L_{\mathcal{Q}_k} = \frac{L_r}{1-\gamma \max\{L_P+1, L_P(1+L_\pi)\}}$.

Proof. Let us now consider the dissimilarity term in norm:

$$\begin{aligned} \|d_{\mathcal{Q}_k}^\pi\|_{p,\eta_k^{\rho,\pi}}^p &= \int_{\mathcal{S}} \int_{\mathcal{A}} \eta_k^{\rho,\pi}(\mathrm{d}s,\mathrm{d}a) \left| \sup_{f \in \mathcal{Q}_k} \left| \int_{\mathcal{S}} \int_{\mathcal{A}} \left(P^\pi(\mathrm{d}s',\mathrm{d}a'|s,a) - P^\delta(\mathrm{d}s',\mathrm{d}a'|s,a) \right) f(s',a') \right| \right|^p \\ &\leqslant L_{\mathcal{Q}_k}^p \int_{\mathcal{S}} \int_{\mathcal{A}} \eta_k^{\rho,\pi}(\mathrm{d}s,\mathrm{d}a) \left| \sup_{f: \|f\|_L \leqslant 1} \left| \int_{\mathcal{S}} \int_{\mathcal{A}} \left(P^\pi(\mathrm{d}s',\mathrm{d}a'|s,a) - P^\delta(\mathrm{d}s',\mathrm{d}a'|s,a) \right) f(s',a') \right| \right|^p, \end{aligned}$$

where the inequality follows from Lemma A.5. We now consider the inner term and perform the following algebraic manipulations:

$$\begin{aligned} &\sup_{f: \|f\|_L \leqslant 1} \left| \int_{\mathcal{S}} \int_{\mathcal{A}} \left(P^\pi(\mathrm{d}s',\mathrm{d}a'|s,a) - P^\delta(\mathrm{d}s',\mathrm{d}a'|s,a) \right) f(s',a') \right| \\ &\quad = \sup_{f: \|f\|_L \leqslant 1} \left| \int_{\mathcal{S}} \int_{\mathcal{A}} P(\mathrm{d}s'|s,a)\pi(\mathrm{d}a'|s') f(s',a') - \int_{\mathcal{S}} \int_{\mathcal{A}} P(\mathrm{d}s'|s,a)\delta_a(\mathrm{d}a') f(s',a') \right. \\ &\quad\quad \left. \pm \int_{\mathcal{S}} \int_{\mathcal{A}} \delta_s(\mathrm{d}s')\pi(\mathrm{d}a'|s') \pm \int_{\mathcal{S}} \int_{\mathcal{A}} \delta_s(\mathrm{d}s')\delta_a(\mathrm{d}a') f(s',a') \right| \end{aligned}$$

$$
\begin{aligned}
&\leqslant \sup_{f:\|f\|_L\leqslant 1}\left|\int_{\mathcal{S}}\left(P(\mathrm{d}s'|s,a)-\delta_s(\mathrm{d}s')\right)\int_{\mathcal{A}}\pi(\mathrm{d}a'|s')f(s',a')\right| \\
&\quad+\sup_{f:\|f\|_L\leqslant 1}\left|\int_{\mathcal{S}}\left(P(\mathrm{d}s'|s,a)-\delta_s(\mathrm{d}s')\right)\int_{\mathcal{A}}\delta_a(\mathrm{d}a')f(s',a')\right| \\
&\quad+\sup_{f:\|f\|_L\leqslant 1}\left|\int_{\mathcal{S}}\delta_s(\mathrm{d}s')\int_{\mathcal{A}}\left(\pi(\mathrm{d}a'|s')-\delta_a(\mathrm{d}a')\right)f(s',a')\right|.
\end{aligned}
$$

We now consider the first two terms:

$$
\begin{aligned}
&\sup_{f:\|f\|_L\leqslant 1}\left|\int_{\mathcal{S}}\left(P(\mathrm{d}s'|s,a)-\delta_s(\mathrm{d}s')\right)\int_{\mathcal{A}}\pi(\mathrm{d}a'|s')f(s',a')\right|+\sup_{f:\|f\|_L\leqslant 1}\left|\int_{\mathcal{S}}\left(P(\mathrm{d}s'|s,a)-\delta_s(\mathrm{d}s')\right)\int_{\mathcal{A}}\delta_a(\mathrm{d}a')f(s',a')\right| \\
&\qquad\leqslant (L_\pi+1)\mathcal{W}_1(P(\cdot|s,a),\delta_s) \qquad \text{(P.35)}\\
&\qquad\leqslant (L_\pi+1)L_T,
\end{aligned}
$$

where line (P.35) follows from observing that the function $g_f(s')=\int_{\mathcal{A}}\pi(\mathrm{d}a'|s')f(s',a')$ is L_π-LC, and function $h_f(s')=\int_{\mathcal{A}}\delta_a(\mathrm{d}a')f(s',a')=f(s',a)$ is 1-LC. Moreover, under Assumption 4.1, we have that $\mathcal{W}_1(P(\cdot|s,a),\delta_s)\leqslant L_T$. Let us now focus on the third term:

$$
\begin{aligned}
&\sup_{f:\|f\|_L\leqslant 1}\left|\int_{\mathcal{S}}\delta_s(\mathrm{d}s')\int_{\mathcal{A}}\left(\pi(\mathrm{d}a'|s')-\delta_a(\mathrm{d}a')\right)f(s',a')\right|=\sup_{f:\|f\|_L\leqslant 1}\left|\int_{\mathcal{A}}\left(\pi(\mathrm{d}a'|s)-\delta_a(\mathrm{d}a')\right)f(s,a')\right| \\
&\qquad=\sup_{f:\|f\|_L\leqslant 1}\left|\int_{\mathcal{A}}\left(\pi(\mathrm{d}a'|s)-\delta_a(\mathrm{d}a')\right)f(a')\right| \qquad \text{(P.36)}\\
&\qquad=\sup_{f:\|f\|_L\leqslant 1}\left|\int_{\mathcal{A}}\left(\int_{\mathcal{A}}\pi(\mathrm{d}a''|s)\delta_{a'}(\mathrm{d}a'')-\delta_a(\mathrm{d}a')\right)f(a')\right| \qquad \text{(P.37)}\\
&\qquad=\sup_{f:\|f\|_L\leqslant 1}\left|\int_{\mathcal{A}}\pi(\mathrm{d}a''|s)\int_{\mathcal{A}}\left(\delta_{a''}(\mathrm{d}a')-\delta_a(\mathrm{d}a')\right)f(a')\right| \qquad \text{(P.38)}\\
&\qquad\leqslant\int_{\mathcal{A}}\pi(\mathrm{d}a''|s)\sup_{f:\|f\|_L\leqslant 1}\left|\int_{\mathcal{A}}\left(\delta_{a''}(\mathrm{d}a')-\delta_a(\mathrm{d}a')\right)f(a')\right| \qquad \text{(P.39)}\\
&\qquad=\int_{\mathcal{A}}\pi(\mathrm{d}a''|s)d_{\mathcal{A}}(a,a''), \qquad \text{(P.40)}
\end{aligned}
$$

where line (P.36) follows from observing that the dependence on s for function f can be neglected because of the supremum, line (P.37) is obtained from the equality $\pi(\mathrm{d}a'|s)=\int_{\mathcal{A}}\pi(\mathrm{d}a''|s)\delta_{a'}(\mathrm{d}a'')$, line (P.38) derives from moving the integral over a'' outside and recalling that $\delta_{a''}(\mathrm{d}a')=\delta_{a'}(\mathrm{d}a'')$, line (P.39) comes from Jensen inequality. Finally, line (P.40) is obtained from the definition of Kantorovich distance between Dirac deltas. Now, we take the expectation w.r.t. $\eta_k^{\rho,\pi}$. Recalling that $\rho(s,a)=\rho_{\mathcal{S}}(s)\pi(a|s)$ it follows that the same decomposition holds for $\eta_k^{\rho,\pi}(s,a)=\eta_{k,\mathcal{S}}^{\rho,\pi}(s)\pi(a|s)$. Consequently, exploiting the above equation, we have:

$$
\begin{aligned}
\int_{\mathcal{S}}\eta_{k,\mathcal{S}}^{\rho,\pi}(\mathrm{d}s)\int_{\mathcal{A}}\pi(\mathrm{d}a|s)\left|\int_{\mathcal{A}}\pi(\mathrm{d}a''|s)d_{\mathcal{A}}(a,a'')\right|^p &\leqslant\int_{\mathcal{S}}(\eta_k^{\rho,\pi})_{\mathcal{S}}(\mathrm{d}s)\int_{\mathcal{A}}\pi(\mathrm{d}a|s)\int_{\mathcal{A}}\pi(\mathrm{d}a''|s)d_{\mathcal{A}}(a,a'')^p \\
&\leqslant\sup_{s\in\mathcal{S}}\int_{\mathcal{A}}\int_{\mathcal{A}}\pi(\mathrm{d}a|s)\pi(\mathrm{d}a''|s)d_{\mathcal{A}}(a,a'')^p=\sigma_p^p,
\end{aligned}
$$

where the first inequality follows from an application of Jensen inequality. An application of Minkowski inequality on the norm $\left\|d_{\mathcal{Q}_k}^{\pi}\right\|_{p,\eta_k^{\rho,\pi}}$ concludes the proof. □

Lemma A.6. *If $\mathcal{A}=\mathbb{R}^{d_{\mathcal{A}}}$, $d_{\mathcal{A}}(\mathbf{a},\mathbf{a}')=\|\mathbf{a}-\mathbf{a}'\|_2$, then it holds that $\sigma_2^2\leqslant 2\sup_{s\in\mathcal{S}}\mathbb{V}\mathrm{ar}[A]$, with $A\sim\pi(\cdot|s)$.*

Proof. Let $s\in\mathcal{S}$ and define the mean–action in state s as:

$$
\overline{\mathbf{a}}(s)=\int_{\mathcal{A}}\mathbf{a}\pi(\mathrm{d}\mathbf{a}|s).
$$

Thus, we have:

$$
\begin{aligned}
\sigma_2^2&=\sup_{s\in\mathcal{S}}\int_{\mathcal{A}}\int_{\mathcal{A}}\|\mathbf{a}-\mathbf{a}'\|_2^2\pi(\mathrm{d}\mathbf{a}|s)\pi(\mathrm{d}\mathbf{a}'|s) \\
&=\sup_{s\in\mathcal{S}}\int_{\mathcal{A}}\int_{\mathcal{A}}\|\mathbf{a}-\mathbf{a}'\pm\overline{\mathbf{a}}(s)\|_2^2\pi(\mathrm{d}\mathbf{a}|s)\pi(\mathrm{d}\mathbf{a}'|s)
\end{aligned}
$$

$$\leqslant \sup_{s\in\mathcal{S}} \int_{\mathcal{A}}\int_{\mathcal{A}} \|\mathbf{a}-\overline{\mathbf{a}}(s)\|_2^2 \pi(\mathrm{d}\mathbf{a}|s)\pi(\mathrm{d}\mathbf{a}'|s) + \sup_{s\in\mathcal{S}} \int_{\mathcal{A}}\int_{\mathcal{A}} \|\mathbf{a}'-\overline{\mathbf{a}}(s)\|_2^2 \pi(\mathrm{d}\mathbf{a}|s)\pi(\mathrm{d}\mathbf{a}'|s)$$
$$= \sup_{s\in\mathcal{S}} \int_{\mathcal{A}} \|\mathbf{a}-\overline{\mathbf{a}}(s)\|_2^2 \pi(\mathrm{d}\mathbf{a}|s) + \sup_{s\in\mathcal{S}} \int_{\mathcal{A}} \|\mathbf{a}'-\overline{\mathbf{a}}(s)\|_2^2 \pi(\mathrm{d}\mathbf{a}'|s)$$
$$= 2\sup_{s\in\mathcal{S}} \int_{\mathcal{A}} \|\mathbf{a}-\overline{\mathbf{a}}(s)\|_2^2 \pi(\mathrm{d}\mathbf{a}|s) = 2\sup_{s\in\mathcal{S}} \mathbb{V}\mathrm{ar}[A].$$

□

Remark A.1 (On the choice of $d_{\mathcal{A}}$ when $|\mathcal{A}|<+\infty$)**.** *When the action space $\mathcal{A}$ is finite and it is a subset of a metric space (e.g., $\mathbb{R}^{d_{\mathcal{A}}}$) we can employ the same metric as $d_{\mathcal{A}}$. Otherwise, we use the* discrete metric $d_{\mathcal{A}}(a,a') = \mathbb{1}\{a \neq a'\}$.

A.3. Proofs of Section 5

Proposition A.1. *Assuming that the evaluation of the estimated Q-function in a state action pair has computational complexity $\mathcal{O}(1)$, the computational complexity of J iterations of PFQI(k) run with a dataset $\mathcal{D}$ of n samples, neglecting the cost of the regression, is given by:*

$$\mathcal{O}\left(Jn\left(1+\frac{|\mathcal{A}|-1}{k}\right)\right).$$

Proof. Let us consider an iteration $j=0,\dots,J-1$. If $j \bmod k = 0$, we perform an application of $\widehat{T}^*$ which requires to perform $n|\mathcal{A}|$ evaluations of the next-state value function in order to compute the maximum over the actions. On the contrary, when $j \bmod k \neq 0$, we perform an application of $\widehat{T}^\delta$ which requires just n evaluations, since the next-state value function is evaluated in the persistent action only. By the definition of PFQI(k), J must be an integer multiple of the persistence k. Recalling that a single evaluation of the approximate Q-function is $\mathcal{O}(1)$, we have that the overall complexity is:

$$\mathcal{O}\left(\sum_{j\in\{0,\dots,J-1\} \wedge j \bmod k=0} n|\mathcal{A}| + \sum_{j\in\{0,\dots,J-1\} \wedge j \bmod k\neq 0} n\right) = \mathcal{O}\left(\frac{J}{k}n|\mathcal{A}| + \frac{J(k-1)}{k}n\right) = \mathcal{O}\left(Jn\left(1+\frac{|\mathcal{A}|-1}{k}\right)\right).$$

□

Theorem 5.1 (Error Propagation for PFQI(k))**.** *Let $p\geqslant 1$, $k\in\mathbb{N}_{\geqslant 1}$, $J\in\mathbb{N}_{\geqslant 1}$ with $J \bmod k = 0$ and $\rho\in\mathscr{P}(\mathcal{S}\times\mathcal{A})$. Then for any sequence $(Q^{(j)})_{j=0}^J \subset \mathcal{F}$ uniformly bounded by $Q_{\max} \leqslant \frac{R_{\max}}{1-\gamma}$, the corresponding $(\epsilon^{(j)})_{j=0}^{J-1}$ defined in Equation (12) and for any $r\in[0,1]$ and $q\in[1,+\infty]$ it holds that:*

$$\left\|Q_k^* - Q_k^{\pi^{(J)}}\right\|_{p,\rho} \leqslant \frac{2\gamma^k}{(1-\gamma)(1-\gamma^k)}\left[\frac{2}{1-\gamma}\gamma^{\frac{J}{p}} R_{\max} + C_{\mathrm{VI},\rho,\nu}^{\frac{1}{2p}}(J,r,q)\mathcal{E}^{\frac{1}{2p}}(\epsilon^{(0)},\dots,\epsilon^{(J-1)};r,q)\right].$$

The expression of $C_{\mathrm{VI},\rho,\nu}(J;r,q)$ and $\mathcal{E}(\cdot;r,q)$ can be found in Appendix A.3.

Before proving the main result, we need to introduce a variation of the *concentrability* coefficients (Antos et al., 2008; Farahmand, 2011) to account for action persistence.

Definition A.1 (Persistent Expected Concentrability)**.** *Let $\rho,\nu\in\mathscr{P}(\mathcal{S}\times\mathcal{A})$, $L\in\mathbb{N}_{\geqslant 1}$, and an arbitrary sequence of stationary policies $(\pi^{(l)})_{l=1}^L$. Let $k\in\mathbb{N}_{\geqslant 1}$ be the persistence. For any $m_1,m_2,m_3\in\mathbb{N}_{\geqslant 1}$ and $q\in[1,+\infty]$, we define:*

$$c_{\mathrm{VI}_1,k,q,\rho,\nu}(m_1,m_2,m_3;\pi) = \mathbb{E}\left[\left|\frac{\mathrm{d}\left(\rho(P_k^\pi)^{m_1}(P_k^{\pi_k^*})^{m_2}(P^\delta)^{m_3}\right)}{\mathrm{d}\nu}(S,A)\right|^{\frac{q}{q-1}}\right]^{\frac{q-1}{q}},$$

$$c_{\mathrm{VI}_2,k,q,\rho,\nu}(m_1,m_2;(\pi^{(l)})_{l=1}^L) = \mathbb{E}\left[\left|\frac{\mathrm{d}\left(\rho(P_k^{\pi^{(L)}})^{m_1}P_k^{\pi^{(L-1)}}\dots P_k^{\pi^{(1)}}(P^\delta)^{m_2}\right)}{\mathrm{d}\nu}(S,A)\right|^{\frac{q}{q-1}}\right]^{\frac{q-1}{q}},$$

with $(S,A)\sim\nu$. If $\rho(P_k^\pi)^{m_1}(P_k^{\pi_k^*})^{m_2}(P^\delta)^{m_3}$ (resp. $\rho(P_k^{\pi^{(L)}})^{m_1}P_k^{\pi^{(L-1)}}\dots P_k^{\pi^{(1)}}(P^\delta)^{m_2}$) is not absolutely continuous w.r.t. to ν, then we take $c_{\mathrm{VI}_1,\rho,\nu}(m_1,m_2,m_3;\pi,k)=+\infty$ (resp. $c_{\mathrm{VI}_2,\rho,\nu}(m_1,m_2;(\pi^{(l)})_{l=1}^L,k)=+\infty$).

This definition is a generalization of that provided in Farahmand (2011), that can be recovered by setting $k=1$, $q=2$, $m_3=0$ for the first coefficient and $m_2=0$ for the second coefficient..

Proof. The proof follows most of the steps of Theorem 3.4 of Farahmand (2011). We start by deriving a bound relating $Q^* - Q^{(J)}$ to $(\epsilon^{(j)})_{j=0}^{J-1}$. To this purpose, let us first define the cumulative error over k iterations for every $j \bmod k = 0$:

$$\epsilon_k^{(j)} = T_k^* Q^{(j)} - Q^{(j+k)}. \tag{P.41}$$

Let us denote with π_k^* one of the optimal policies of the k-persistent MDP $\mathcal{M}_k$. We have:

$$Q_k^* - Q^{(j+k)} = T_k^{\pi_k^*} Q_k^* - T_k^{\pi_k^*} Q^{(j)} + T_k^{\pi_k^*} Q^{(j)} - T_k^* Q^{(j)} + \epsilon_k^{(j)} \leqslant \gamma^k P_k^{\pi_k^*} (Q_k^* - Q^{(j)}) + \epsilon_k^{(j)},$$
$$Q_k^* - Q^{(j+k)} = T_k^* Q_k^* - T_k^{\pi^{(j)}} Q^* + T_k^{\pi^{(j)}} Q^* - T_k^* Q^{(j)} + \epsilon_k^{(j)} \geqslant \gamma^k P_k^{\pi^{(j)}} (Q_k^* - Q^{(j)}) + \epsilon_k^{(j)},$$

where we exploited the fact that $T_k^* Q^{(j)} \geqslant T_k^{\pi_k^*} Q^{(j)}$, the definition of greedy policy $\pi^{(j)}$ that implies that $T_k^{\pi^{(j)}} Q^{(j)} = T_k^* Q^{(j)}$ and the definition of $\epsilon_k^{(j)}$. By unrolling the expression derived above, we have that for every $J \bmod k = 0$:

$$\begin{aligned} Q_k^* - Q^{(J)} &\leqslant \sum_{h=0}^{\frac{J}{k}-1} \gamma^{J-k(h+1)} \left(P_k^{\pi_k^*} \right)^{\frac{J}{k}-h-1} \epsilon_k^{(j)} + \gamma^J \left(P_k^{\pi_k^*} \right)^{\frac{J}{k}} (Q_k^* - Q^{(0)}) \\ Q_k^* - Q^{(J)} &\geqslant \sum_{h=0}^{\frac{J}{k}-1} \gamma^{J-k(h+1)} \left(P_k^{\pi^{(J-k)}} P_k^{\pi^{(J-2k)}} \dots P_k^{\pi^{(k(h+1))}} \right) \epsilon_k^{(j)} + \gamma^J \left(P_k^{\pi^{(J)}} P_k^{\pi^{(J-k)}} \dots P_k^{\pi^{(k)}} \right) (Q_k^* - Q^{(0)}). \end{aligned} \tag{P.42}$$

We now provide the following bound relating the difference $Q_k^* - Q_k^{\pi^{(J)}}$ to the difference $Q_k^* - Q^{(J)}$:

$$\begin{aligned} Q_k^* - Q_k^{\pi^{(J)}} &= T_k^{\pi_k^*} Q_k^* - T_k^{\pi_k^*} Q^{(J)} + T_k^{\pi_k^*} Q^{(J)} - T_k^* Q^{(J)} + T_k^* Q^{(J)} - T_k^{\pi^{(J)}} Q_k^{\pi^{(J)}} \\ &\leqslant T_k^{\pi_k^*} Q_k^* - T_k^{\pi_k^*} Q^{(J)} + T_k^* Q^{(J)} - T_k^{\pi^{(J)}} Q_k^{\pi^{(J)}} \\ &= \gamma^k P_k^{\pi_k^*} (Q^* - Q^{(J)}) + \gamma^k P_k^{\pi^{(J)}} (Q^{(J)} - Q_k^{\pi^{(J)}}) \\ &= \gamma^k P_k^{\pi_k^*} (Q^* - Q^{(J)}) + \gamma^k P_k^{\pi^{(J)}} (Q^{(J)} - Q_k^* + Q_k^* - Q_k^{\pi^{(J)}}), \end{aligned}$$

where we exploited the fact that $T_k^* Q^{(J)} \geqslant T_k^{\pi_k^*} Q^{(J)}$ and observed that $T_k^* Q^{(J)} = T_k^{\pi^{(J)}} Q^{(J)}$. By using Lemma 4.2 of Munos (2007) we can derive:

$$Q_k^* - Q_k^{\pi^{(J)}} \leqslant \gamma^k \left(\mathrm{Id} - \gamma^k P_k^{\pi^{(J)}} \right)^{-1} \left(P_k^{\pi_k^*} - P_k^{\pi^{(J)}} \right) (Q^* - Q^{(J)}). \tag{P.43}$$

By plugging Equation (P.42) into Equation (P.43):

$$\begin{aligned} Q_k^* - Q_k^{\pi^{(J)}} \leqslant \gamma^k \left(\mathrm{Id} - \gamma^k P_k^{\pi^{(J)}} \right)^{-1} \Bigg[& \sum_{h=0}^{\frac{J}{k}-1} \gamma^{J-k(h+1)} \left(\left(P_k^{\pi_k^*} \right)^{\frac{J}{k}-h} - \left(P_k^{\pi^{(J)}} P_k^{\pi^{(J-k)}} P_k^{\pi^{(J-2k)}} \dots P_k^{\pi^{(k(h+1))}} \right) \right) \epsilon_k^{(j)} \\ & + \gamma^J \left(\left(P_k^{\pi_k^*} \right)^{\frac{J}{k}+1} - \left(P_k^{\pi^{(J)}} P_k^{\pi^{(J)}} P_k^{\pi^{(J-k)}} \dots P_k^{\pi^{(k)}} \right) \right) (Q_k^* - Q^{(0)}) \Bigg]. \end{aligned} \tag{P.44}$$

Before proceeding, we need to relate the cumulative errors $\epsilon_k^{(j)}$ to the single-step errors $\epsilon^{(j)}$:

$$\begin{aligned} \epsilon_k^{(j)} &= T_k^* Q^{(j)} - Q^{(j+k)} \\ &= (T^\delta)^{k-1} T^* Q^{(j)} - (T^\delta)^{k-1} Q^{(j+1)} + (T^\delta)^{k-1} Q^{(j+1)} - Q^{(j+k)} \\ &= \gamma^{k-1} (P^\delta)^{k-1} \left(T^* Q^{(j)} - Q^{(j+1)} \right) + (T^\delta)^{k-1} Q^{(j+1)} - Q^{(j+k)} \\ &= \gamma^{k-1} (P^\delta)^{k-1} \epsilon^{(j)} + (T^\delta)^{k-1} Q^{(j+1)} - Q^{(j+k)}. \end{aligned}$$

Let us now consider the remaining term $(T^\delta)^{k-1} Q^{(j+1)} - Q^{(j+k)}$:

$$\begin{aligned} (T^\delta)^{k-1} Q^{(j+1)} - Q^{(j+k)} &= (T^\delta)^{k-1} Q^{(j+1)} - (T^\delta)^{k-2} Q^{(j+2)} + (T^\delta)^{k-2} Q^{(j+2)} - Q^{(j+k)} \\ &= \gamma^{k-2} (P^\delta)^{k-2} \left(T^\delta Q^{(j+1)} - Q^{(j+2)} \right) + (T^\delta)^{k-2} Q^{(j+2)} - Q^{(j+k)} \end{aligned}$$

$$= \gamma^{k-2} (P^\delta)^{k-2} \epsilon^{(j+1)} + (T^\delta)^{k-2} Q^{(j+2)} - Q^{(j+k)}$$
$$= \sum_{l=2}^{k} \gamma^{k-l} (P^\delta)^{k-l} \epsilon^{(j+l-1)},$$

where the last step is obtained by unrolling the recursion. Putting all together, we get:

$$\epsilon_k^{(j)} = \sum_{l=1}^{k} \gamma^{k-l} (P^\delta)^{k-l} \epsilon^{(j+l-1)}. \tag{P.45}$$

Consequently, we can rewrite Equation (P.44) as follows:

$$\begin{aligned}
Q_k^* - Q_k^{\pi^{(J)}} &\leqslant \gamma^k \left(\mathrm{Id} - \gamma^k P_k^{\pi^{(J)}}\right)^{-1} \left[\sum_{h=0}^{\frac{J}{k}-1} \gamma^{J-k(h+1)} \left(\left(P_k^{\pi_k^*}\right)^{\frac{J}{k}-h} - \left(P_k^{\pi^{(J)}} P_k^{\pi^{(J-k)}} P_k^{\pi^{(J-2k)}} \dots P_k^{\pi^{(k(h+1))}} \right) \right) \right. \\
&\quad \left. \times \sum_{l=1}^{k} \gamma^{k-l} (P^\delta)^{k-l} \epsilon^{(j+l-1)} + \gamma^J \left(\left(P_k^{\pi_k^*}\right)^{\frac{J}{k}+1} - \left(P_k^{\pi^{(J)}} P_k^{\pi^{(J)}} P_k^{\pi^{(J-k)}} \dots P_k^{\pi^{(k)}} \right) \right) (Q_k^* - Q^{(0)}) \right] \\
&= \gamma^k \left(\mathrm{Id} - \gamma^k P_k^{\pi^{(J)}}\right)^{-1} \left[\sum_{h=0}^{\frac{J}{k}-1} \sum_{l=1}^{k} \gamma^{J-kh-l} \left(\left(P_k^{\pi_k^*}\right)^{\frac{J}{k}-h} - \left(P_k^{\pi^{(J)}} P_k^{\pi^{(J-k)}} P_k^{\pi^{(J-2k)}} \dots P_k^{\pi^{(k(h+1))}} \right) \right) \right. \qquad \text{(P.46)} \\
&\quad \left. \times (P^\delta)^{k-l} \epsilon^{(j+l-1)} + \gamma^J \left(\left(P_k^{\pi_k^*}\right)^{\frac{J}{k}+1} - \left(P_k^{\pi^{(J)}} P_k^{\pi^{(J)}} P_k^{\pi^{(J-k)}} \dots P_k^{\pi^{(k)}} \right) \right) (Q_k^* - Q^{(0)}) \right] \\
&= \gamma^k \left(\mathrm{Id} - \gamma^k P_k^{\pi^{(J)}}\right)^{-1} \left[\sum_{j=0}^{J-1} \gamma^{J-j-1} \left(\left(P_k^{\pi_k^*}\right)^{\frac{J}{k}-j \text{ div } k} - \left(P_k^{\pi^{(J)}} P_k^{\pi^{(J-k)}} P_k^{\pi^{(J-2k)}} \dots P_k^{\pi^{(J-k(j \text{ div } k+1))}} \right) \right) \right. \\
&\quad \left. \times (P^\delta)^{k-j \bmod k-1} \epsilon^{(j)} + \gamma^J \left(\left(P_k^{\pi_k^*}\right)^{\frac{J}{k}+1} - \left(P_k^{\pi^{(J)}} P_k^{\pi^{(J)}} P_k^{\pi^{(J-k)}} \dots P_k^{\pi^{(k)}} \right) \right) (Q_k^* - Q^{(0)}) \right] \qquad \text{(P.47)} \\
&\leqslant \gamma^k \left(\mathrm{Id} - \gamma^k P_k^{\pi^{(J)}}\right)^{-1} \left[\sum_{j=0}^{J-1} \gamma^{J-j-1} \left(\left(P_k^{\pi_k^*}\right)^{\frac{J}{k}-j \text{ div } k} + \left(P_k^{\pi^{(J)}} P_k^{\pi^{(J-k)}} P_k^{\pi^{(J-2k)}} \dots P_k^{\pi^{(J-k(j \text{ div } k+1))}} \right) \right) \right. \\
&\quad \left. \times (P^\delta)^{k-j \bmod k-1} \left|\epsilon^{(j)}\right| + \gamma^J \left(\left(P_k^{\pi_k^*}\right)^{\frac{J}{k}+1} + \left(P_k^{\pi^{(J)}} P_k^{\pi^{(J)}} P_k^{\pi^{(J-k)}} \dots P_k^{\pi^{(k)}} \right) \right) \left|Q_k^* - Q^{(0)}\right| \right], \qquad \text{(P.48)}
\end{aligned}$$

where line (P.46) derives from rearranging the two summations, line (P.47) is obtained from a redefinition of the indexes. Specifically, we observed that $h = j \text{ div } k$, $j+1 = kh + l$, and $l = j \bmod k + 1$. Finally, line (P.48) is obtained by applying the absolute value to the right hand side and using Jensen inequality. We now introduce the following terms:

$$A_j = \begin{cases} \frac{1-\gamma^k}{2} \left(\mathrm{Id} - \gamma^k P_k^{\pi^{(J)}}\right)^{-1} \left(\left(P_k^{\pi_k^*}\right)^{\frac{J}{k}-j \text{ div } k} + \left(P_k^{\pi^{(J)}} P_k^{\pi^{(J-k)}} P_k^{\pi^{(J-2k)}} \dots P_k^{\pi^{(J-k(j \text{ div } k+1))}} \right) \right) (P^\delta)^{k-j \bmod k-1} & \text{if } 0 \leqslant j < J \\ \frac{1-\gamma^k}{2} \left(\mathrm{Id} - \gamma^k P_k^{\pi^{(J)}}\right)^{-1} \left(\left(P_k^{\pi_k^*}\right)^{\frac{J}{k}+1} + \left(P_k^{\pi^{(J)}} P_k^{\pi^{(J)}} P_k^{\pi^{(J-k)}} \dots P_k^{\pi^{(k)}} \right) \right) & \text{if } j = J \end{cases}. \tag{P.49}$$

Let us recall the definition of α_j as in Farahmand (2011):

$$\alpha_j = \begin{cases} \frac{(1-\gamma)\gamma^{J-j-1}}{1-\gamma^{J+1}} & \text{if } 0 \leqslant j < J \\ \frac{(1-\gamma)\gamma^J}{1-\gamma^{J+1}} & \text{if } j = J \end{cases}. \tag{P.50}$$

Recalling that $\left|Q_k^* - Q^{(0)}\right| \leqslant Q_{\max} + \frac{R_{\max}}{1-\gamma} \leqslant \frac{2R_{\max}}{1-\gamma}$ and applying Jensen inequality we get to the inequality:

$$Q_k^* - Q_k^{\pi^{(J)}} \leqslant \frac{2\gamma^k (1-\gamma^{J+1})}{(1-\gamma^k)(1-\gamma)} \left[\sum_{j=0}^{J-1} \alpha_j A_j \left|\epsilon^{(j)}\right| + \alpha_J \frac{2R_{\max}}{1-\gamma} \mathbf{1} \right],$$

where $\mathbf{1}$ denotes the constant function on $\mathcal{S} \times \mathcal{A}$ with value 1. Taking the $L_p(\rho)$–norm both sides, recalling that $\sum_{j=1}^{J} \alpha_j = 1$ and that the terms A_j are positive linear operators $A_j : \mathscr{B}(\mathcal{S} \times \mathcal{A}) \to \mathscr{B}(\mathcal{S} \times \mathcal{A})$ such that $A_j \mathbf{1} = \mathbf{1}$. Thus, by Lemma 12 of Antos et al. (2008), we can apply Jensen inequality twice (once w.r.t. α_j and once w.r.t. A_j), getting:

$$\left\|Q_k^* - Q_k^{\pi^{(J)}}\right\|_{p,\rho}^p \leqslant \left(\frac{2\gamma^k(1-\gamma^{J+1})}{(1-\gamma^k)(1-\gamma)}\right)^p \rho\left[\sum_{j=0}^{J-1} \alpha_j A_j \left|\epsilon^{(j)}\right|^p + \alpha_J \left(\frac{2R_{\max}}{1-\gamma}\right)^p \mathbf{1}\right].$$

Consider now the individual terms $\rho A_j \left|\epsilon^{(j)}\right|^p$ for $0 \leqslant j < J$. By the properties of the Neumann series we have:

$$\begin{aligned}
\rho A_j \left|\epsilon^{(j)}\right|^p &= \frac{1-\gamma^k}{2}\rho\left(\mathrm{Id} - \gamma^k P_k^{\pi^{(J)}}\right)^{-1}\left(\left(P_k^{\pi_k^*}\right)^{\frac{J}{k} - j \text{ div } k} + \left(P_k^{\pi^{(J)}} P_k^{\pi^{(J-k)}} P_k^{\pi^{(J-2k)}} \dots P_k^{\pi^{(J-k(j \text{ div } k+1))}}\right)\right) \\
&\quad \times (P^\delta)^{k-j \bmod k-1}\left|\epsilon^{(j)}\right|^p \\
&= \frac{1-\gamma^k}{2}\rho\left[\sum_{m=0}^{+\infty} \gamma^{km}\left(\left(P_k^{\pi^{(J)}}\right)^m \left(P_k^{\pi_k^*}\right)^{\frac{J}{k} - j \text{ div } k} + \left(\left(P_k^{\pi^{(J)}}\right)^{m+1} P_k^{\pi^{(J-k)}} P_k^{\pi^{(J-2k)}} \dots P_k^{\pi^{(J-k(j \text{ div } k))}}\right)\right)\right] \\
&\quad \times (P^\delta)^{k-j \bmod k-1}\left|\epsilon^{(j)}\right|^p.
\end{aligned}$$

We now aim at introducing the concentrability coefficients and for this purpose, we employ the following inequality. For any measurable function $f \in (\mathcal{X}) \to \mathbb{R}$, and the probability measures $\mu_1, \mu_2 \in \mathscr{P}(\mathcal{X})$ such that μ_2 is absolutely continuous w.r.t. μ_1, we have the following Hölder inequality, for any $q \in [1, +\infty]$:

$$\int_{\mathcal{X}} f \mathrm{d}\mu_1 \leqslant \left(\int_{\mathcal{X}} \left|\frac{\mathrm{d}\mu_1}{\mathrm{d}\mu_2}\right|^{\frac{q}{q-1}} \mathrm{d}\mu_2\right)^{\frac{q-1}{q}} \left(\int_{\mathcal{X}} |f|^q \mathrm{d}\mu_2\right)^{\frac{1}{q}}. \tag{P.51}$$

We now focus on a single term $\rho\left(P_k^{\pi^{(J)}}\right)^m \left(P_k^{\pi_k^*}\right)^{\frac{J}{k} - j \text{ div } k} \left|\epsilon^{(j)}\right|^p$ and we apply the above inequality:

$$\begin{aligned}
\rho\left(P_k^{\pi^{(J)}}\right)^m \left(P_k^{\pi_k^*}\right)^{\frac{J}{k} - j \text{ div } k} (P^\delta)^{k-j \bmod k-1}\left|\epsilon^{(j)}\right|^p &\leqslant \left(\int_{\mathcal{S}\times\mathcal{A}} \left|\frac{\mathrm{d}\rho\left(P_k^{\pi^{(J)}}\right)^m \left(P_k^{\pi_k^*}\right)^{\frac{J}{k} - j \text{ div } k} (P^\delta)^{k-j \bmod k-1}}{\mathrm{d}\nu}\right|^{\frac{q}{q-1}} \mathrm{d}\nu\right)^{\frac{q-1}{q}} \\
&\quad \times \left(\int_{\mathcal{S}\times\mathcal{A}} \left|\epsilon^{(j)}\right|^{pq} \mathrm{d}\nu\right)^{\frac{1}{q}} \\
&= c_{\mathrm{VI}_1,k,q,\rho,\nu}\left(m, \frac{J}{k} - j \text{ div } k, k-j \bmod k-1; \pi^{(J)}\right)\left\|\epsilon^{(j)}\right\|_{pq,\nu}^p.
\end{aligned}$$

Proceeding in an analogous way for the remaining terms, we get to the expression:

$$\begin{aligned}
\left\|Q_k^* - Q_k^{\pi^{(J)}}\right\|_{p,\rho}^p \leqslant \left(\frac{2\gamma^k(1-\gamma^{J+1})}{(1-\gamma^k)(1-\gamma)}\right)^p &\left[\frac{1-\gamma^k}{2}\sum_{j=0}^{J-1}\sum_{m=0}^{+\infty} \gamma^{km}\left(c_{\mathrm{VI}_1,k,q,\rho,\nu}\left(m, \frac{J}{k} - j \text{ div } k, k-j \bmod k-1; \pi^{(J)}\right)\right.\right. \\
&\left.\left. + c_{\mathrm{VI}_2,k,q,\rho,\nu}\left(m+1, k-j \bmod k-1; \{\pi^{(J-lk)}\}_{l=1}^{j \text{ div } k}\right)\right)\left\|\epsilon^{(j)}\right\|_{pq,\nu}^p + \alpha_J\left(\frac{2R_{\max}}{1-\gamma}\right)^p\right].
\end{aligned}$$

To separate the concentrability coefficients and the approximation errors, we apply Hölder inequality with $s \in [1, +\infty]$:

$$\sum_{j=0}^{J} a_j b_j \leqslant \left(\sum_{j=0}^{J} |a_j|^s\right)^{\frac{1}{s}} \left(|b_j|^{\frac{s}{s-1}}\right)^{\frac{s-1}{s}}. \tag{P.52}$$

Let $r \in [0,1]$, we set $a_j = \alpha_j^r \left\|\epsilon^{(j)}\right\|_{pq,\nu}^p$ and $b_j = \alpha_j^{1-r} \frac{1-\gamma^k}{2} \sum_{j=0}^{J-1} \sum_{m=0}^{+\infty} \gamma^{km}\left(c_{\mathrm{VI}_1,k,q,\rho,\nu}\left(m, \frac{J}{k} - j \text{ div } k, k-j \bmod k-1; \pi^{(J)}\right) + c_{\mathrm{VI}_2,k,q,\rho,\nu}\left(m+1, k-j \bmod k-1; \{\pi^{(J-lk)}\}_{l=1}^{j \text{ div } k}\right)\right)$. The application of Hölder inequality leads to:

$$\left\|Q_k^* - Q_k^{\pi^{(J)}}\right\|_{p,\rho}^p \leqslant \left(\frac{2\gamma^k(1-\gamma^{J+1})}{(1-\gamma^k)(1-\gamma)}\right)^p \frac{1-\gamma^k}{2}\left[\sum_{j=0}^{J-1} \alpha_j^{\frac{s(1-r)}{s-1}} \left(\sum_{m=0}^{+\infty} \gamma^{km}\left(c_{\mathrm{VI}_1,k,q,\rho,\nu}\left(m, \frac{J}{k} - j \text{ div } k, k-j \bmod k-1; \pi^{(J)}\right)\right.\right.\right.$$

$$+c_{\mathrm{VI}_2,k,q,\rho,\nu}\left(m+1,k-j \bmod k-1;\{\pi^{(J-lk)}\}_{l=1}^{j \text{ div } k}\right)\Bigg)\Bigg)^{\frac{s}{s-1}}\Bigg]^{\frac{s-1}{s}}\left[\sum_{j=0}^{J-1}\alpha_j^{sr}\left\|\epsilon^{(j)}\right\|_{pq,\nu}^{sp}\right]^{\frac{1}{s}}$$

$$+\left(\frac{2\gamma^k(1-\gamma^{J+1})}{(1-\gamma^k)(1-\gamma)}\right)^p\alpha_J\left(\frac{2R_{\max}}{1-\gamma}\right)^p.$$

Since the policies $(\pi^{(J-lk)})_{l=1}^{j \text{ div } k}$ are not known, we define the following quantity by taking the supremum over any sequence of policies:

$$C_{\mathrm{VI},\rho,\nu}(J;r,s,q)=\left(\frac{1-\gamma^k}{2}\right)^s\sup_{\pi_0,\dots,\pi_J\in\Pi}\sum_{j=0}^{J-1}\alpha_j^{\frac{s(1-r)}{s-1}}\left(\sum_{m=0}^{+\infty}\gamma^{km}\left(c_{\mathrm{VI}_1,k,q,\rho,\nu}\left(m,\frac{J}{k}-j \text{ div } k,k-j \bmod k-1;\pi_J\right)\right.\right.$$
$$\left.\left.+c_{\mathrm{VI}_2,k,q,\rho,\nu}\left(m+1,k-j \bmod k-1;\{\pi_l\}_{l=1}^{j \text{ div } k}\right)\right)\right)^{\frac{s}{s-1}}. \tag{P.53}$$

Moreover, we define the following term that embeds all the terms related to the approximation error:

$$\mathcal{E}(\epsilon^{(0)},\dots,\epsilon^{(J-1)};r,s,q)=\sum_{j=0}^{J-1}\alpha_j^{sr}\left\|\epsilon^{(j)}\right\|_{pq,\nu}^{sp}. \tag{P.54}$$

Observing that $\frac{1-\gamma}{1-\gamma^{J+1}}\leqslant 1$ and $1-\gamma^{J-1}\leqslant 1$, we can put all together and taking the p–th root and recalling that the inequality holds for all $q\in[1,+\infty]$, $r\in[0,1]$, and $s\in[1,+\infty]$:

$$\left\|Q_k^*-Q_k^{\pi^{(J)}}\right\|_{p,\rho}\leqslant\frac{2\gamma^k}{(1-\gamma^k)(1-\gamma)}\left[\inf_{\substack{q\in[1,+\infty]\\ r\in[0,1]\\ s\in[1,+\infty]}}C_{\mathrm{VI},\rho,\nu}(J;r,s,q)^{\frac{s-1}{ps}}\mathcal{E}(\epsilon^{(0)},\dots,\epsilon^{(J-1)};r,s,q)^{\frac{1}{ps}}+\gamma^{\frac{J}{p}}\frac{2R_{\max}}{1-\gamma}\right].$$

The statement is simplified by taking $s=2$.

□

A.4. Proofs of Section 6

Lemma 6.1. *Let $Q\in\mathscr{B}(\mathcal{S}\times\mathcal{A})$ and π be a greedy policy w.r.t. Q. Let $J^\rho=\int\rho(\mathrm{d}s)V(s)$, with $V(s)=\max_{a\in\mathcal{A}}Q(s,a)$ for all $s\in\mathcal{S}$. Then, for any $k\in\mathbb{N}_{\geqslant 1}$, it holds that:*

$$J_k^{\rho,\pi}\geqslant J^\rho-\frac{1}{1-\gamma^k}\left\|T_k^*Q-Q\right\|_{1,\eta^{\rho,\pi}}, \tag{14}$$

where $\eta^{\rho,\pi}=(1-\gamma^k)\rho\pi\left(\mathrm{Id}-\gamma^k P_k^\pi\right)^{-1}$, is the γ-discounted stationary distribution induced by policy π and distribution ρ in MDP $\mathcal{M}_k$.

Proof. We start by providing the following equality, recalling that $T_k^*Q=T_k^\pi Q$, being π the greedy policy w.r.t. Q:

$$\begin{aligned}Q_k^\pi-Q&=T_k^\pi Q_k^\pi-T_k^\pi Q+T_k^*Q-Q\\&=\gamma^k P_k^\pi(Q_k^\pi-Q)+T_k^*Q-Q\\&=\left(\mathrm{Id}-\gamma^k P_k^\pi\right)^{-1}(T_k^*Q-Q),\end{aligned}$$

where the last equality follows from the properties of the Neumann series. We take the expectation w.r.t. to the distribution $\rho\pi$ both sides. For the left hand side we have:

$$J_k^{\rho,\pi}-J^\rho=\rho\pi Q_k^\pi-\rho\pi Q.$$

Concerning the right hand side, instead, we have:

$$\rho\pi\left(\mathrm{Id}-\gamma^k P_k^\pi\right)^{-1}(T_k^*Q-Q)=\frac{1}{1-\gamma^k}\eta^{\rho,\pi}(T_k^*Q-Q),$$

where we introduced the γ–discounted stationary distribution (Sutton et al., 1999a) after normalization. Putting all together, we can derive the following inequality:

$$J_k^{\rho,\pi}-J^\rho=\frac{1}{1-\gamma^k}\eta^{\rho,\pi}(T_k^*Q-Q)$$

$$\geqslant -\frac{1}{1-\gamma^k}\eta^{\rho,\pi}\left|T_k^* Q - Q\right|$$
$$= -\frac{1}{1-\gamma^k}\left\|T_k^* Q - Q\right\|_{1,\eta^{\rho,\pi}}.$$

□

B. Details on Bounding the Performance Loss (Section 4)

In this appendix, we report some additional material that is referenced in Section 4, concerning the performance loss due to the usage of action persistence.

B.1. Discussion on the Persistence Bound (Theorem 4.1)

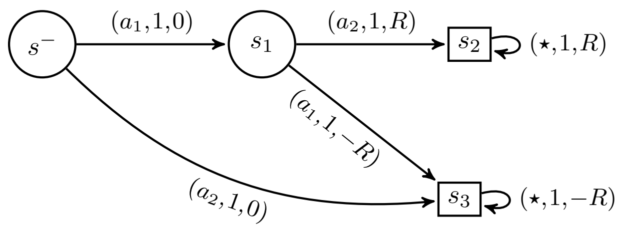

Figure 4. The MDP counter-example of Proposition B.1, where $R>0$. Each arrow connecting two states s and s' is labeled with the 3-tuple $(a,P(s'|s,a),r(s,a))$; the symbol $\star$ denotes any action in $\mathcal{A}$. While the optimal policy in the original MDP starting in s^- can avoid negative rewards by executing an action sequence of the kind $(a_1,a_2,\dots)$, every policy in the k-persistent MDP, with $k\in\mathbb{N}_{\geqslant 2}$, inevitably ends in the negative terminal state, as the only possible action sequences are of the kind $(a_1,a_1,\dots)$ and $(a_2,a_2,\dots)$.

We start with a negative result, showing that with no structure it is possible to make the bound of Theorem 4.1 vacuous, and thus, independent from k.

Proposition B.1. *For any MDP $\mathcal{M}$ and $k\in\mathbb{N}_{\geqslant 2}$ it holds that:*

$$V_k^*(s) \geqslant V^*(s) - \frac{2\gamma R_{\max}}{1-\gamma}, \quad \forall s\in\mathcal{S}. \tag{18}$$

Furthermore, there exists an MDP $\mathcal{M}^-$ (Figure 4) and a state $s^-\in\mathcal{S}$ such that the bound holds with equality for all $k\in\mathbb{N}_{\geqslant 2}$.

Proof. First of all, we recall that $V^*(s)-V_k^*(s)\geqslant 0$ since we cannot increase performance when executing a policy with a persistence k. Let π^* an optimal policy on the MDP $\mathcal{M}$, we observe that for all $s\in\mathcal{S}$:

$$V^*(s) - V_k^*(s) \leqslant V^{\pi^*}(s) - V_k^{\pi^*}(s), \tag{P.55}$$

since $V^{\pi^*}(s)=V^*(s)$ and $V_k^*(s)\geqslant V_k^{\pi^*}(s)$. Let us now consider the corresponding Q-functions $Q^{\pi^*}(s,a)$ and $Q_k^{\pi^*}(s,a)$. Recalling that they are the fixed points of the Bellman operators T^{π^*} and $T_k^{\pi^*}$ we have:

$$\begin{aligned}
Q^{\pi^*} - Q_k^{\pi^*} &= T^{\pi^*}Q^{\pi^*} - T_k^{\pi^*}Q_k^{\pi^*} \\
&= r + \gamma P^{\pi} Q^{\pi^*} - r_k - \gamma^k P_k^{\pi} Q_k^{\pi^*} \\
&= r + \gamma P^{\pi} Q^{\pi^*} - \sum_{i=0}^{k-1}\gamma^i \left(P^{\delta}\right)^i r - \gamma^k P_k^{\pi} Q_k^{\pi^*} \\
&= \gamma P^{\pi} Q^{\pi^*} - \sum_{i=1}^{k-1}\gamma^i \left(P^{\delta}\right)^i r - \gamma^k P_k^{\pi} Q_k^{\pi^*},
\end{aligned}$$

where we exploited the definitions of the Bellman expectation operators in the k-persistent MDP. As a consequence, we have that for all

$(s,a)\in\mathcal{S}\times\mathcal{A}$:

$$Q^{\pi^*}(s,a)-Q^{\pi^*}(s,a)\leqslant\gamma\frac{R_{\max}}{1-\gamma}+R_{\max}\sum_{i=1}^{k-1}\gamma^i+\gamma^k\frac{R_{\max}}{1-\gamma}$$
$$=\gamma\frac{R_{\max}}{1-\gamma}+R_{\max}\frac{\gamma(1-\gamma^{k-1})}{1-\gamma}+\gamma^k\frac{R_{\max}}{1-\gamma}=\frac{2\gamma R_{\max}}{1-\gamma},$$

where we considered the following facts that hold for all $(s,a)\in\mathcal{S}\times\mathcal{A}$: $\left(P^\pi Q^{\pi^*}\right)(s,a)\leqslant\frac{R_{\max}}{1-\gamma}$, $\left(\left(P^\delta\right)^i r\right)(s,a)\leqslant R_{\max}$, and $\left(P_k^\pi Q_k^{\pi^*}\right)\leqslant R_{\max}$. The result follows, by simply observing that $V^{\pi^*}(s)-V_k^{\pi^*}(s)=\mathbb{E}\left[Q^{\pi^*}(s,A)-Q^{\pi^*}(s,A)\right]$, where $A\sim\pi^*(\cdot|s)$.

We now prove that the bound is tight for the MDP of Figure 4. From inspection, we observe that the optimal policy must reach the terminal state s_2 yielding the positive reward $R>0$. Thus the optimal policy plays action a_1 in state s^- and action a_2 in state s_1, generating a value function $V^*(s^-)=\frac{\gamma R}{1-\gamma}$. Let us now consider the 2-persistent MDP $\mathcal{M}_2^-$. Whichever action is played in state s^- it is going to be persisted for the subsequent decision epoch and, consequently, we will end up in state s_3, yielding the negative reward $-R<0$. Thus, the optimal value function will be $V_2^*(s^-)=-\frac{\gamma R}{1-\gamma}$. Clearly, the same rationale holds for any persistence $k\in\mathbb{N}_{\geqslant 3}$. □

The quantity $\frac{2\gamma R_{\max}}{1-\gamma}$ is the maximum performance that we can lose if we perform the same action at decision epoch $t=0$ and then we follow an arbitrary policy thereafter.

B.2. On using divergences other than the Kantorovich

The Persistence Bound presented in Theorem 4.1 is defined in terms of the dissimilarity index $d^\pi_{\mathcal{Q}_k}$ which depends on the set of functions $\mathcal{Q}_k$ defined in terms of the k-persistent Q-function Q_k^π and in terms of the Bellman operators T^π and T^δ. Clearly, this bound is meaningful when it yields a value that is smaller than $\frac{2\gamma R_{\max}}{1-\gamma}$ that we already know to be the maximum performance degradation we experience when executing policy π with persistence (Proposition B.1). Therefore, for any meaningful choice of $\mathcal{Q}_k$, we require that, at least for $k=2$, the following condition to hold:

$$\frac{\gamma(1-\gamma^{k-1})}{(1-\gamma)(1-\gamma^k)}\left\|d^\pi_{\mathcal{Q}_k}\right\|_{p,\eta_k^{\rho,\pi}}\Bigg|_{k=2}=\frac{\gamma}{(1-\gamma^2)}\left\|d^\pi_{\mathcal{Q}_2}\right\|_{p,\eta_2^{\rho,\pi}}<\frac{2\gamma R_{\max}}{1-\gamma}. \tag{19}$$

If we require no additional regularity conditions on the MDP, we can only exploit the fact that all functions $f\in\mathcal{Q}_k$ are uniformly bounded by $\frac{R_{\max}}{1-\gamma}$, reducing $d^\pi_{\mathcal{Q}_k}$ to the total variation distance between P^π and P^δ:

$$d^\pi_{\mathcal{Q}_k}(s,a)\leqslant\frac{R_{\max}}{1-\gamma}\sup_{f:\|f\|_\infty\leqslant 1}\left|\int_{\mathcal{S}}\int_{\mathcal{A}}\left(P^\pi(\mathrm{d}s',\mathrm{d}a'|s,a)-P^\delta(\mathrm{d}s',\mathrm{d}a'|s,a)\right)f(s',a')\right|=\frac{2R_{\max}}{1-\gamma}d^\pi_{\mathrm{TV}}(s,a). \tag{20}$$

We restrict our discussion to deterministic policies and, for this purpose, we denote with $\pi(s)\in\mathcal{A}$ the action prescribed by policy π in the state $s\in\mathcal{S}$. Thus, the total variation distance as follows:

$$d^\pi_{\mathrm{TV}}(s,a)=\frac{1}{2}\int_{\mathcal{S}}\int_{\mathcal{A}}\left|P^\pi(\mathrm{d}s',\mathrm{d}a'|s,a)-P^\delta(\mathrm{d}s',a'|s,a)\right|$$
$$=\frac{1}{2}\int_{\mathcal{S}}P(\mathrm{d}s'|s,a)\int_{\mathcal{A}}\left|\pi(\mathrm{d}a'|s')-\delta_a(\mathrm{d}a')\right|$$
$$=\frac{1}{2}\int_{\mathcal{S}}P(\mathrm{d}s'|s,a)\int_{\mathcal{A}}\left|\delta_{\pi(s')}(a')-\delta_{\pi(s)}(\mathrm{d}a')\right|$$
$$=\int_{\mathcal{S}}P(\mathrm{d}s'|s,a)\mathbb{1}_{\{\pi(s)\neq\pi(s')\}},$$

where $\mathbb{1}_{\mathcal{X}}$ denotes the indicator function for the measurable set $\mathcal{X}$. Consequently, we can derive for the norm:

$$\left\|d^\pi_{\mathcal{Q}_2}\right\|^p_{p,\eta_2^{\rho,\pi}}\leqslant\frac{2R_{\max}}{1-\gamma}\int_{\mathcal{S}}\int_{\mathcal{A}}\eta_k^{\rho,\pi}(\mathrm{d}s,\mathrm{d}a)\left|\int_{\mathcal{S}}P(\mathrm{d}s'|s,a)\mathbb{1}_{\{\pi(s)\neq\pi(s')\}}\right|^p$$
$$\leqslant\frac{2R_{\max}}{1-\gamma}\int_{\mathcal{S}}\int_{\mathcal{A}}\eta_k^{\rho,\pi}(\mathrm{d}s,\mathrm{d}a)\int_{\mathcal{S}}P(\mathrm{d}s'|s,a)\left|\mathbb{1}_{\{\pi(s)\neq\pi(s')\}}\right|^p$$

$$= \frac{2R_{\max}}{1-\gamma} \int_{\mathcal{S}} \int_{\mathcal{A}} \eta_k^{\rho,\pi}(\mathrm{d}s, \mathrm{d}a) \int_{\mathcal{S}} P(\mathrm{d}s'|s,a) \mathbb{1}_{\{\pi(s) \neq \pi(s')\}}.$$

Thus, such term depend on the expected fraction of state-next-state pairs such that their policies prescribe different actions. Consequently, considering the condition at Equation (19), we have that it must be fulfilled:

$$\int_{\mathcal{S}} \int_{\mathcal{A}} \eta_k^{\rho,\pi}(\mathrm{d}s, \mathrm{d}a) \int_{\mathcal{S}} P(\mathrm{d}s'|s,a) \mathbb{1}_{\{\pi(s) \neq \pi(s')\}} \leqslant 1 - \gamma^2.$$

However, if for every state-next-state pair the prescribed actions are different (even if very similar in some metric space), the left hand side would be 1 and the inequality never satisfied. To embed the notion of closeness of actions we need to resort to distance metrics different from the total variation (e.g., the Kantorovich). These considerations can be extended to the case of stochastic policies.

B.3. Time–Lipschitz Continuity for dynamical systems

We now draw a connection between the rate at which a dynamical system evolves and the L_T constant of Assumption 4.1. Consider a continuous-time dynamical system having $\mathcal{S} = \mathbb{R}^{d_{\mathcal{S}}}$ and $\mathcal{A} = \mathbb{R}^{d_{\mathcal{A}}}$ governed by the law $\dot{\mathbf{s}}(t) = \mathbf{f}(\mathbf{s}(t), \mathbf{a}(t))$ such that $\sup_{\mathbf{s} \in \mathcal{S}, \mathbf{a} \in \mathcal{A}} \|\mathbf{f}(\mathbf{s}, \mathbf{a})\| \leqslant F < +\infty$. Suppose to control the system with a discrete time step $\Delta t_0 > 0$, inducing an MDP with transition model $P_{\Delta t_0}$. Using a norm $\|\cdot\|$, Assumption 4.1 becomes:

$$\begin{aligned} \mathcal{W}_1(P_{\Delta t_0}(\cdot|\mathbf{s}, \mathbf{a}), \delta_{\mathbf{s}}) &= \|\mathbf{s}(t + \Delta t_0) - \mathbf{s}(t)\| \\ &= \left\| \int_t^{t+\Delta t_0} \dot{\mathbf{s}}(\mathrm{d}t) \right\| \leqslant F \Delta t_0. \end{aligned}$$

Thus, the Time Lipschitz constant L_T depends on: i) how fast the dynamical system evolves (F); ii) the duration of the control time step (Δt_0).

B.4. Discussion on Conditions of Theorem 4.2

In order to bound the dissimilarity term $\left\| d^{\pi}_{\mathcal{Q}_k} \right\|_{p, \eta_k^{\rho,\pi}}$ we require in Theorem 4.2 that $\max\{L_P + 1, L_P(1 + L_\pi)\} < \frac{1}{\gamma}$. This condition can be decomposed in the two conditions: (i) $L_P + 1 < \frac{1}{\gamma}$ and ii) $L_P(1 + L_\pi) < \frac{1}{\gamma}$. While (ii) inherits from the Lipschitz MDP literature with Wasserstein metric (Rachelson & Lagoudakis, 2010), condition i) is typical of action persistence. In principle, we could replace Wasserstein with Total Variation, getting less restrictive conditions (Munos & Szepesvári, 2008, Section 7) but this would rule out deterministic systems. Moreover, the Lipschitz constants are a bound, derived to separate the effects of π and P, as commonly done in the literature. Tighter bounds can be obtanied if we consider the Lipschitz constants of the joint transition models P^π and P^δ. Indeed, lookning at the proof of Lemma A.4 we immediately figure out that:

$$L_{P^\pi} \leqslant L_P(L_\pi + 1), \qquad L_{P^\delta} \leqslant L_P + 1. \tag{21}$$

To clarify the point, consider the following deterministic dynamical linear system with $\mathcal{S} = \mathbb{R}^{d_{\mathcal{S}}}$ controlled via a deterministic linear policy with $\mathcal{A} = \mathbb{R}^{d_{\mathcal{A}}}$:

$$\begin{aligned} \mathbf{s}_{t+1} &= \mathbf{A}\mathbf{s}_t + \mathbf{B}\mathbf{a}_t, \\ \mathbf{a}_t &= \mathbf{K}\mathbf{s}_t, \end{aligned}$$

where $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{K}$ are properly sized matrices. Let us now compute L_{P^π} and L_{P^δ} and the corresponding bounds of Equation (21). To this purpose we use as metric $d_{\mathcal{S} \times \mathcal{A}}((\mathbf{s}, \mathbf{a}), (\overline{\mathbf{s}}, \overline{\mathbf{a}})) = \|\mathbf{s} - \overline{\mathbf{s}}\| + \|\mathbf{a} - \overline{\mathbf{a}}\|$:

$$\begin{aligned} \mathcal{W}_1(P^\pi(\cdot|\mathbf{s}, \mathbf{a}), P^\pi(\cdot|\overline{\mathbf{s}}, \overline{\mathbf{a}})) &\leqslant \|\mathbf{A}(\mathbf{s} - \overline{\mathbf{s}}) + \mathbf{B}(\mathbf{a} - \overline{\mathbf{a}})\| + \|\mathbf{K}\mathbf{A}(\mathbf{s} - \overline{\mathbf{s}}) + \mathbf{K}\mathbf{B}(\mathbf{a} - \overline{\mathbf{a}})\| \\ &\leqslant (\|\mathbf{K}\mathbf{A}\| + \|\mathbf{A}\|) \|\mathbf{s} - \overline{\mathbf{s}}\| + (\|\mathbf{K}\mathbf{B}\| + \|\mathbf{B}\|) \|\mathbf{a} - \overline{\mathbf{a}}\|, \end{aligned}$$

$$\begin{aligned} \mathcal{W}_1\left(P^\delta(\cdot|\mathbf{s}, \mathbf{a}), P^\delta(\cdot|\overline{\mathbf{s}}, \overline{\mathbf{a}})\right) &\leqslant \|\mathbf{A}(\mathbf{s} - \overline{\mathbf{s}}) + \mathbf{B}(\mathbf{a} - \overline{\mathbf{a}})\| + \|\mathbf{a} - \overline{\mathbf{a}}\| \\ &\leqslant \|\mathbf{A}\| \|\mathbf{s} - \overline{\mathbf{s}}\| + (\|\mathbf{B}\| + 1) \|\mathbf{a} - \overline{\mathbf{a}}\|, \end{aligned}$$

leading to $L_{P^\pi} \leqslant \max\{\|\mathbf{K}\mathbf{A}\| + \|\mathbf{A}\|, \|\mathbf{K}\mathbf{B}\| + \|\mathbf{B}\|\}$ and $L_{P^\delta} \leqslant \max\{\|\mathbf{A}\|, \|\mathbf{B}\| + 1\}$. If instead, we compute the corresponding bounds of Equation (21), we have:

$$\mathcal{W}_1\left(P(\cdot|\mathbf{s},\mathbf{a}),P(\cdot|\overline{\mathbf{s}},\overline{\mathbf{a}})\right)\leqslant\|\mathbf{A}(\mathbf{s}-\overline{\mathbf{s}})+\mathbf{B}(\mathbf{a}-\overline{\mathbf{a}})\|\leqslant\|\mathbf{A}\|\,\|\mathbf{s}-\overline{\mathbf{s}}\|+\|\mathbf{B}\|\,\|\mathbf{a}-\overline{\mathbf{a}}\|,$$

$$\mathcal{W}_1\left(\pi(\cdot|\mathbf{s}),\pi(\cdot|\overline{\mathbf{s}})\right)\leqslant\|\mathbf{K}(\mathbf{s}-\overline{\mathbf{s}})\|\leqslant\|\mathbf{K}\|\,\|\mathbf{s}-\overline{\mathbf{s}}\|,$$

leading to $L_P\leqslant\max\{\|\mathbf{A}\|,\|\mathbf{B}\|\}$ and $L_\pi\leqslant\|\mathbf{K}\|$ and, consequently, $L_P(L_{\pi+1})\leqslant\max\{\|\mathbf{A}\|,\|\mathbf{B}\|\}(\|\mathbf{K}\|+1)$ and $L_P+1\leqslant\max\{\|\mathbf{A}\|,\|\mathbf{B}\|\}+1$. Clearly, these latter results induce more restrictive conditions for certain values of $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{K}$. Nevertheless, we believe that the bounds of Equation (21) are unavoidable in the general case.

C. Details on Persistence Selection (Section 6)

In this appendix, we illustrate some details behind the simplifications of Lemma 6.1 to get the persistence selection index B_k.

C.1. Change of Distribution

We discuss intuitively the effects of replacing the distribution $\eta^{\rho,\pi}$ with the sampling distribution ν. To this purpose, we consider the particular case in which ν is the γ-discounted stationary distribution obtained by running a sampling policy u in the environment and using the same ρ as initial state distribution. Therefore, we can state:

$$\eta^{\rho,\pi}=(1-\gamma^k)\rho\pi\left(\mathrm{Id}-\gamma^k P_k^\pi\right)^{-1}=(1-\gamma^k)\sum_{i=0}^{\infty}\gamma^{ki}\rho\pi\left(P_k^\pi\right)^i,$$

$$\nu=(1-\gamma)\rho\pi\left(\mathrm{Id}-\gamma P^u\right)^{-1}=(1-\gamma)\sum_{i=0}^{\infty}\gamma^{i}\rho u\left(P^u\right)^i.$$

There are two main differences between $\eta^{\rho,\pi}$ and ν. First, $\eta^{\rho,\pi}$ a discounted stationary distribution in the k-persistent MDP, while ν is the sampling distribution and thus, it is defined in the original (1-persistent) MDP. Second, while $\eta^{\rho,\pi}$ comes from the execution of the policy π obtained after a certain number iterations of learning, ν is derived by the execution of the sampling policy u. To decouple the effects stated above, let us define the following auxiliary discounted stationary distributions:

$$\eta_1^{\rho,\pi}=(1-\gamma)\rho\pi\left(\mathrm{Id}-\gamma P^\pi\right)^{-1},$$

$$\nu_k=(1-\gamma^k)\rho\pi\left(\mathrm{Id}-\gamma^k P_k^u\right)^{-1}.$$

Thus, $\eta_1^{\rho,\pi}$ is obtained by executing policy π in the original (1-persistent) MDP, while ν_k comes from the execution of u in the k-persistent MDP. Therefore, we can provide the following two decomposition of $\left\|\frac{\eta^{\rho,\pi}}{\nu}\right\|_\infty$:

$$\left\|\frac{\eta^{\rho,\pi}}{\nu}\right\|_\infty=\left\|\frac{\eta^{\rho,\pi}}{\eta_1^{\rho,\pi}}\frac{\eta_1^{\rho,\pi}}{\nu}\right\|_\infty\leqslant\left\|\frac{\eta^{\rho,\pi}}{\eta_1^{\rho,\pi}}\right\|_\infty\left\|\frac{\eta_1^{\rho,\pi}}{\nu}\right\|_\infty,$$

$$\left\|\frac{\eta^{\rho,\pi}}{\nu}\right\|_\infty=\left\|\frac{\eta^{\rho,\pi}}{\nu_k}\frac{\nu_k}{\nu}\right\|_\infty\leqslant\left\|\frac{\eta^{\rho,\pi}}{\nu_k}\right\|_\infty\left\|\frac{\nu_k}{\nu}\right\|_\infty.$$

Therefore, looking at the first decomposition, we observe that in order to keep $\left\|\frac{\eta^{\rho,\pi}}{\nu}\right\|_\infty$ small we can require the following two conditions. First, executing the same policy π at persistence k and 1 must induce similar discounted stationary distributions, i.e., $\left\|\frac{\eta^{\rho,\pi}}{\eta_1^{\rho,\pi}}\right\|_\infty\simeq 1$. This is a condition related to persistence only and connected, in some sense, to the regularity conditions employed in Section 4 to bound the loss induced by action persistence. Second, executing policy π or policy u in the same 1-persistent MDP must induce similar γ-discounted stationary distributions, i.e., $\left\|\frac{\eta_1^{\rho,\pi}}{\nu}\right\|_\infty\simeq 1$. This condition, instead, depends on the similarity between policies π and u and on the properties of the transition model. Clearly, an analogous rationale holds when focusing on the second decomposition. We leave as future work the derivation of more formal conditions to bound the magnitude of $\left\|\frac{\eta^{\rho,\pi}}{\nu}\right\|_\infty$.

C.2. Estimating the Expected Bellman Residual

Once we have an approximation $\widetilde{Q}_k$ of $T_k^* Q$ obtained with the regressor Reg, we can proceed to the decomposition, thanks to the triangular inequality:

$$\|T_k^* Q - Q\|_{1,\nu} \leqslant \left\|\widetilde{Q}_k - Q\right\|_{1,\nu} + \left\|T_k^* Q - \widetilde{Q}_k\right\|_{1,\nu}. \tag{22}$$

As discussed in Farahmand & Szepesvári (2011), simply using $\left\|\widetilde{Q}_k - Q\right\|_{1,\nu}$ as a proxy for $\|T_k^* Q - Q\|_{1,\nu}$ might be overlay optimistic. To overcome this problem we must prevent the underestimation of the expected Bellman residual. The idea proposed in Farahmand & Szepesvári (2011) consists in replacing the regression error $\left\|T_k^* Q - \widetilde{Q}_k\right\|_{1,\nu}$ with a high–probability bound $b_{k,\mathcal{G}}$, depending on the functional space $\mathcal{G}$ of the chosen regressor Reg. Clearly, we have the new problem of getting a meaningful bound $b_{k,\mathcal{G}}$. This issue is treated in Section 7.4 of Farahmand & Szepesvári (2011). If $\mathcal{G}$ is a *small* functional space, i.e., with finite pseudo–dimension, we can employ a standard learning theory bound (Györfi et al., 2002). Since for the persistence selection we employ the same functional space $\mathcal{G}$ and the same number of samples m for all persistences $k \in \mathcal{K}$, the value of such a bound will not depend on k and, therefore, it can be neglected in the optimization process. We stress that our goal is to provide a practical method to have an idea on which is a reasonable persistence to employ.

D. Details on Experimental Evaluation (Section 8)

In this appendix, we report the details about our experimental setting (Appendix D.1), together with additional plots (Appendix D.2) and an experiment investigating the effect of the batch size when using persistence (Appendix D.4).

D.1. Experimental Setting

Table 2 reports the parameters of the experimental setting, which are described in the following.

Infrastructure The experiments have been run on a machine with two CPUs Intel(R) Xeon(R) CPU E7-8880 v4 @ 2.20GHz (22 cores, 44 thread, 55 MB cache) and 128 GB RAM.

Environments The implementation of the environments are the ones provided in Open AI Gym (Brockman et al., 2016) `https://gym.openai.com/envs/`.

Action Spaces For the environments with finite action space, we collect samples with a uniform policy over $\mathcal{A}$; whereas for the environments with a continuous action space, we perform a discretization, reported in the column "Action space", and we employ the uniform policy over the resulting finite action space.

Sample Collection Samples are collected in the base MDP at persistence 1, although for some of them the uniform policy is executed at a higher persistence, k_{sampling}, reported in the column "Sampling Persistence". Using a persistence greater than 1 to generate samples has been fundamental in some cases (e.g., Mountain Car) to get a better exploration of the environment and improving the learning performances.[8]

Number of Iterations In order to perform a complete application of a k-Persisted Bellman Operator in the PFQI algorithm, we need k iterations, so the total number of iterations needed to complete the training must be an integer multiple of k. In order to compare the resulting performances, we chose the persistences as a range of powers of 2. The total number of iterations J is selected empirically so that the estimated Q-function has reached convergence for all tested persistences.

Time Discretization Every environment has its own way to deal with time discretization. In some cases, in order to make the benefits of persistence evident, we needed to reduce the base control timestep of the environment w.r.t. to the original implementation. We report in the column "Original timestep" ($\Delta t_{\text{original}}$) the control timestep in the original implementation of the environment, while the base time step (Δt_0) is obtained as a fraction of $\Delta t_{\text{original}}$. The reduction of the timestep by a factor $m = \Delta t_{\text{original}} / \Delta t_0$ results in an extension of the horizon of the same factor, hence there is a greater number of rewards to sum, with the consequent need of a larger discount factor to maintain the same "effective horizon". Thus, the new horizon H (resp. discount factor γ) can be determined starting from the original horizon H_{original} (resp. original discount factor γ_{original}) as:

$$H = m H_{\text{original}}, \qquad \gamma = (\gamma_{\text{original}})^{\frac{1}{m}}, \qquad \text{where} \quad m = \frac{\Delta t_{\text{original}}}{\Delta t_0}.$$

Regressor Hyperparameters We used the class *ExtraTreesRegressor* in the *scikit-learn* library (Pedregosa et al., 2011) with the following parameters: n_estimators = 100, min_samples_split = 5, and min_samples_leaf = 2.

[8] When considering a sampling persistence $k_{\text{sampling}} > 1$, we record in the dataset all the intermediate repeated actions, so that the tuples (S_t, A_t, S'_t, R_t) are transitions of the base MDP $\mathcal{M}$.

Table 2. Parameters of the experimental setting, used for the PFQI(k) experiments.

Environment	Action space $\mathcal{A}$	Sampling Persistence k_{sampling}	Original timestep $\Delta t_{\text{original}}$ (sec)	Factor $m = \Delta t_{\text{original}} / \Delta t_0$	Original Horizon H_{original}	Original Discount factor γ_{original}	Batch size n	Iterations J
Cartpole	$\{-1,1\}$	1	0.02	4	128	0.99	400	512
Mountain Car	$\{-1,0,1\}$	8	1	2	128	0.99	20	256
Lunar Lander	{Nop, left, main, right}	1	0.02	1	256	0.99	100	256
Pendulum	$\{-2,0,2\}$	1	0.05	1	256	0.99	100	64
Acrobot	$\{-1,0,1\}$	4	0.2	4	128	0.99	200	512
Swimmer	$\{-1,0,1\}^2$	1	2 (frame-skip)	2	128	0.99	100	128
Hopper	$\{-1,0,1\}^3$	1	1 (frame-skip)	2	128	0.99	100	128
Walker 2D	$\{-1,0,1\}^9$	1	1 (frame-skip)	2	128	0.99	100	128

D.2. Additional Plots

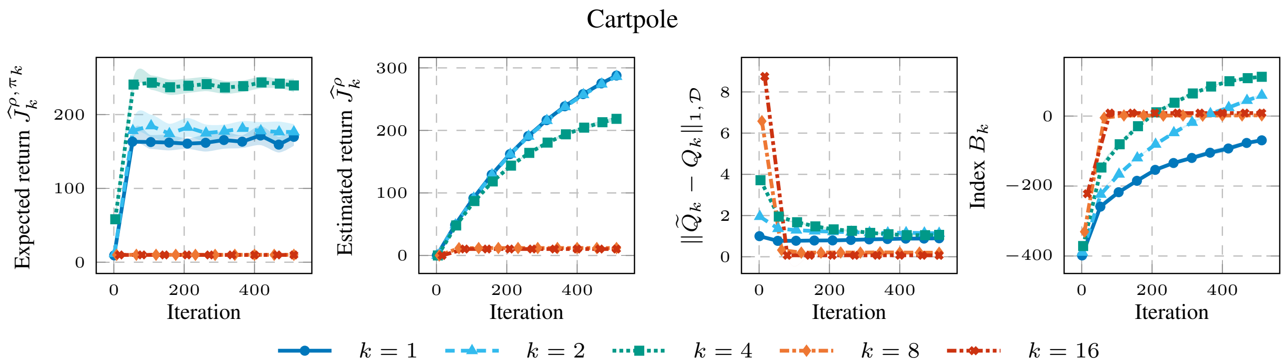

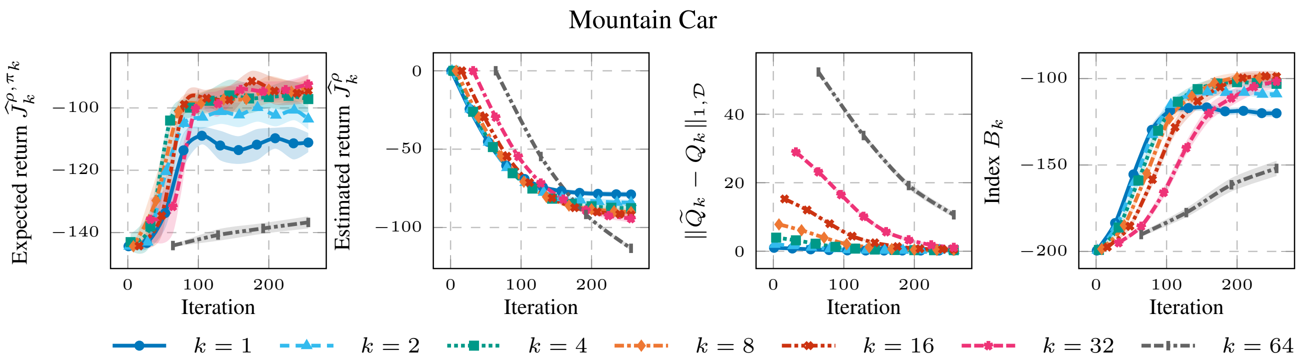

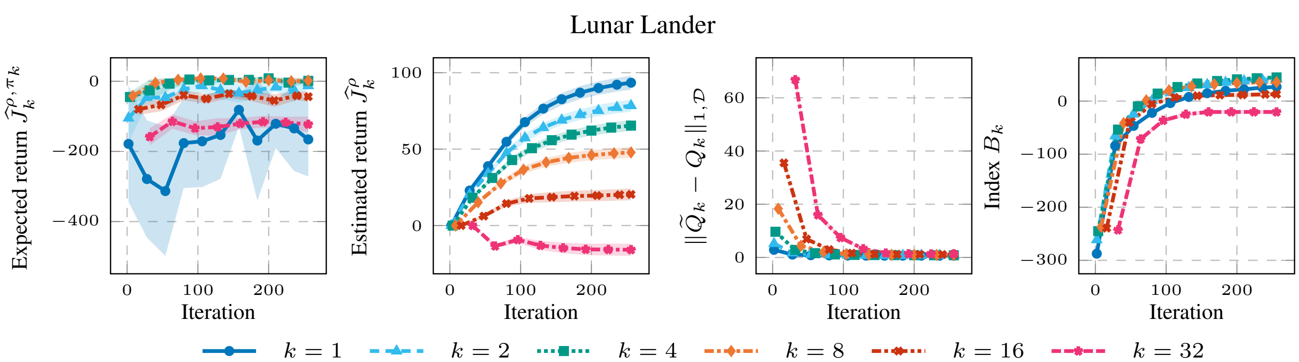

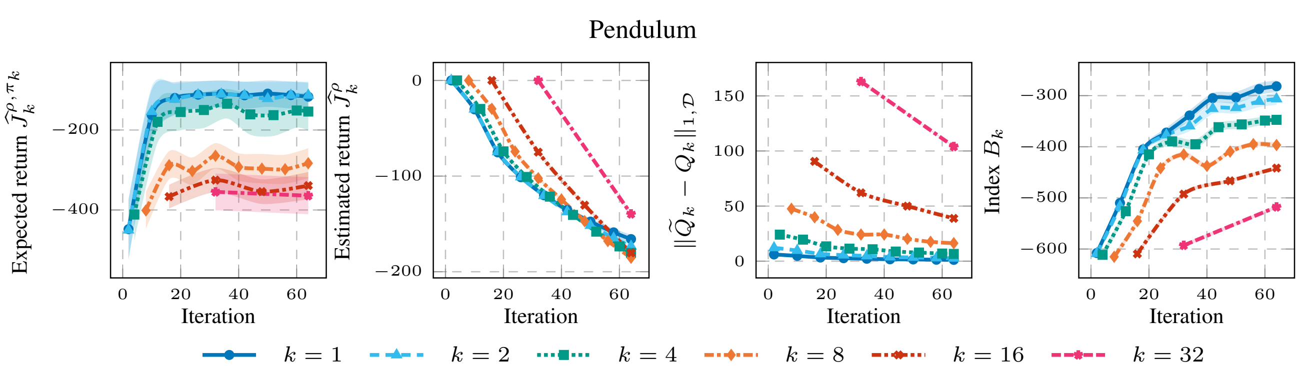

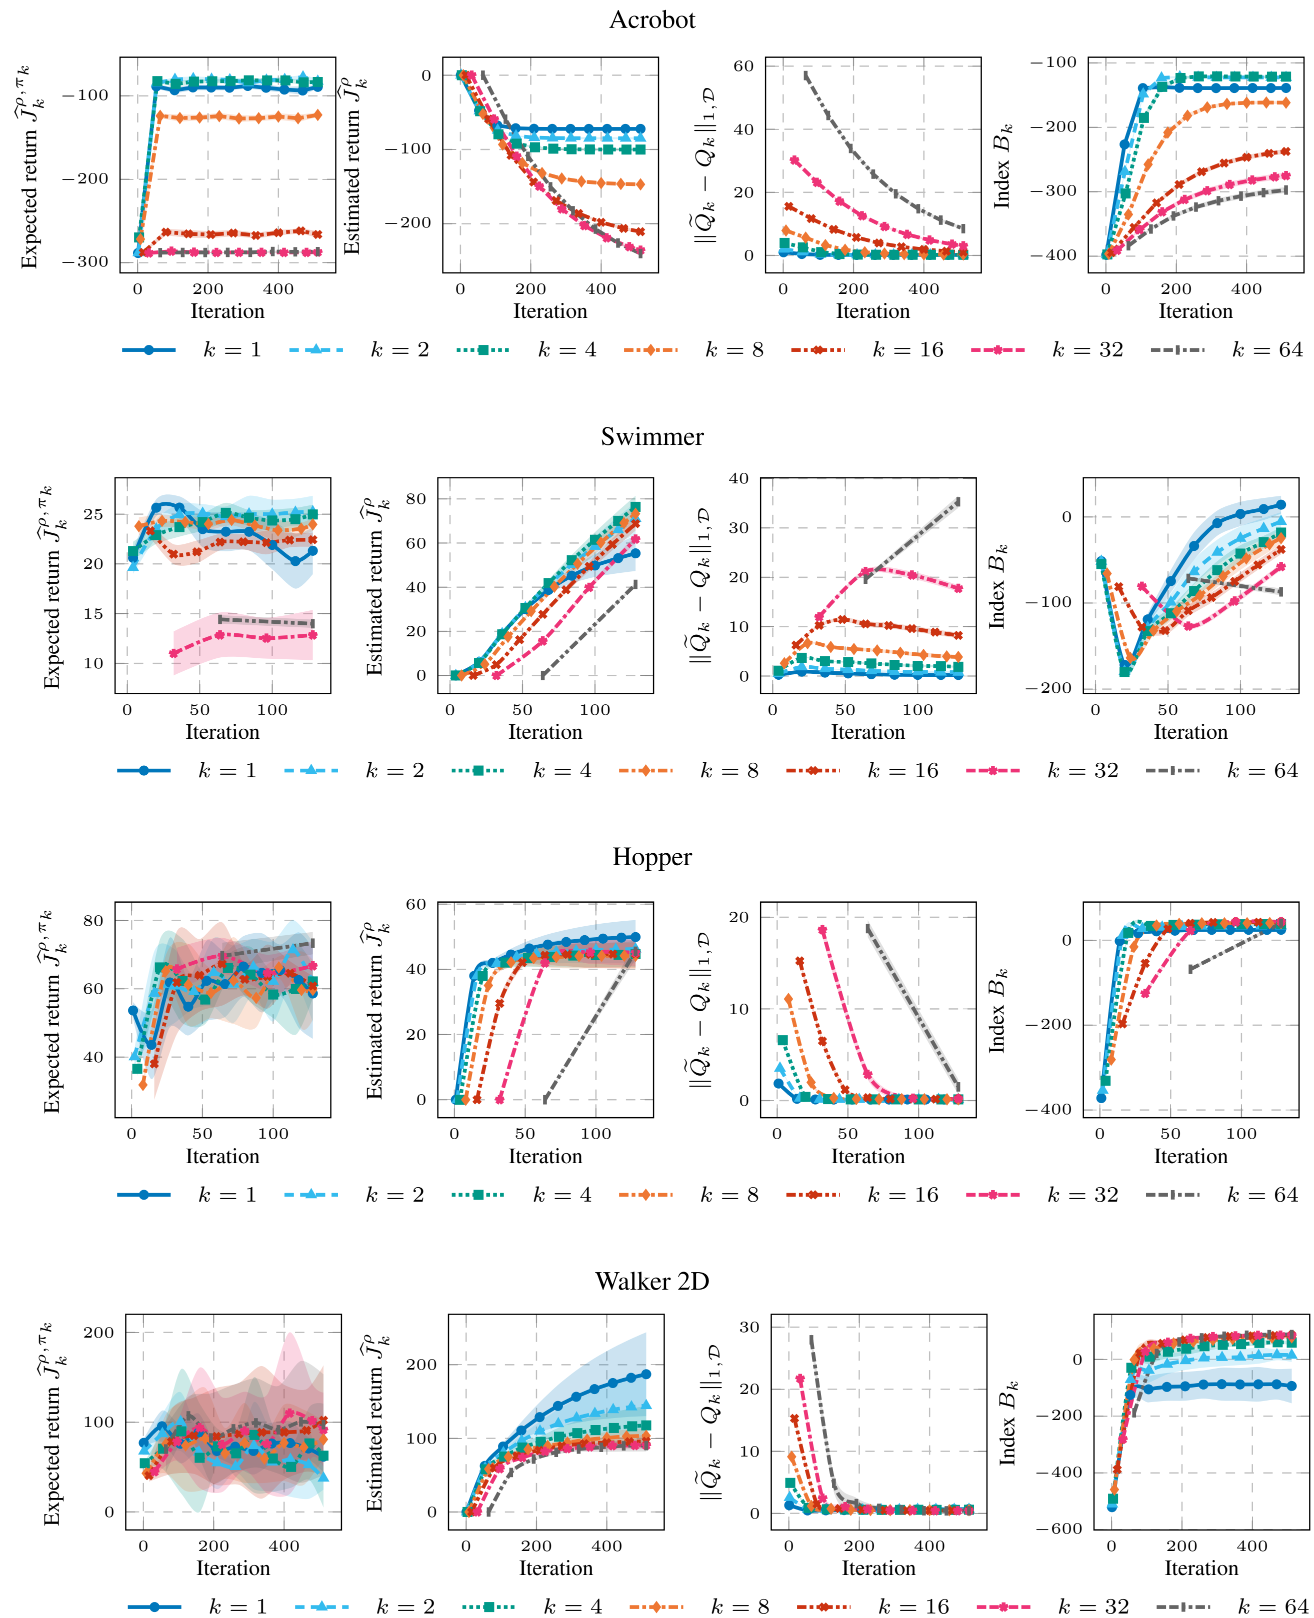

Figure 5. Expected return $\widehat{J}_k^{\rho,\pi_k}$, estimated return $\widehat{J}_k^{\rho}$, estimated expected Bellman residual $\|\widetilde{Q}_k - Q_k\|_{1,\mathcal{D}}$, and persistence selection index B_k for the different experiments as a function of the number of iterations for different persistences. 20 runs, 95 % c.i.

D.3. PFQI with Neural Network as regressor

In the previous experiments we employed extra-trees as regressor to run PFQI. In this appendix, we investigate the effect of employing a neural network as regressor. More specifically, we consider a two-layer network with 64 neurons each and ReLU activation. Figure 6 and Table 3 show the results. The experimental setting is identical to that presented in Appendix D.1. Although the performances are overall lower compared to the case of extremely randomized trees, we notice the same trade-off in the choice of the persistence.

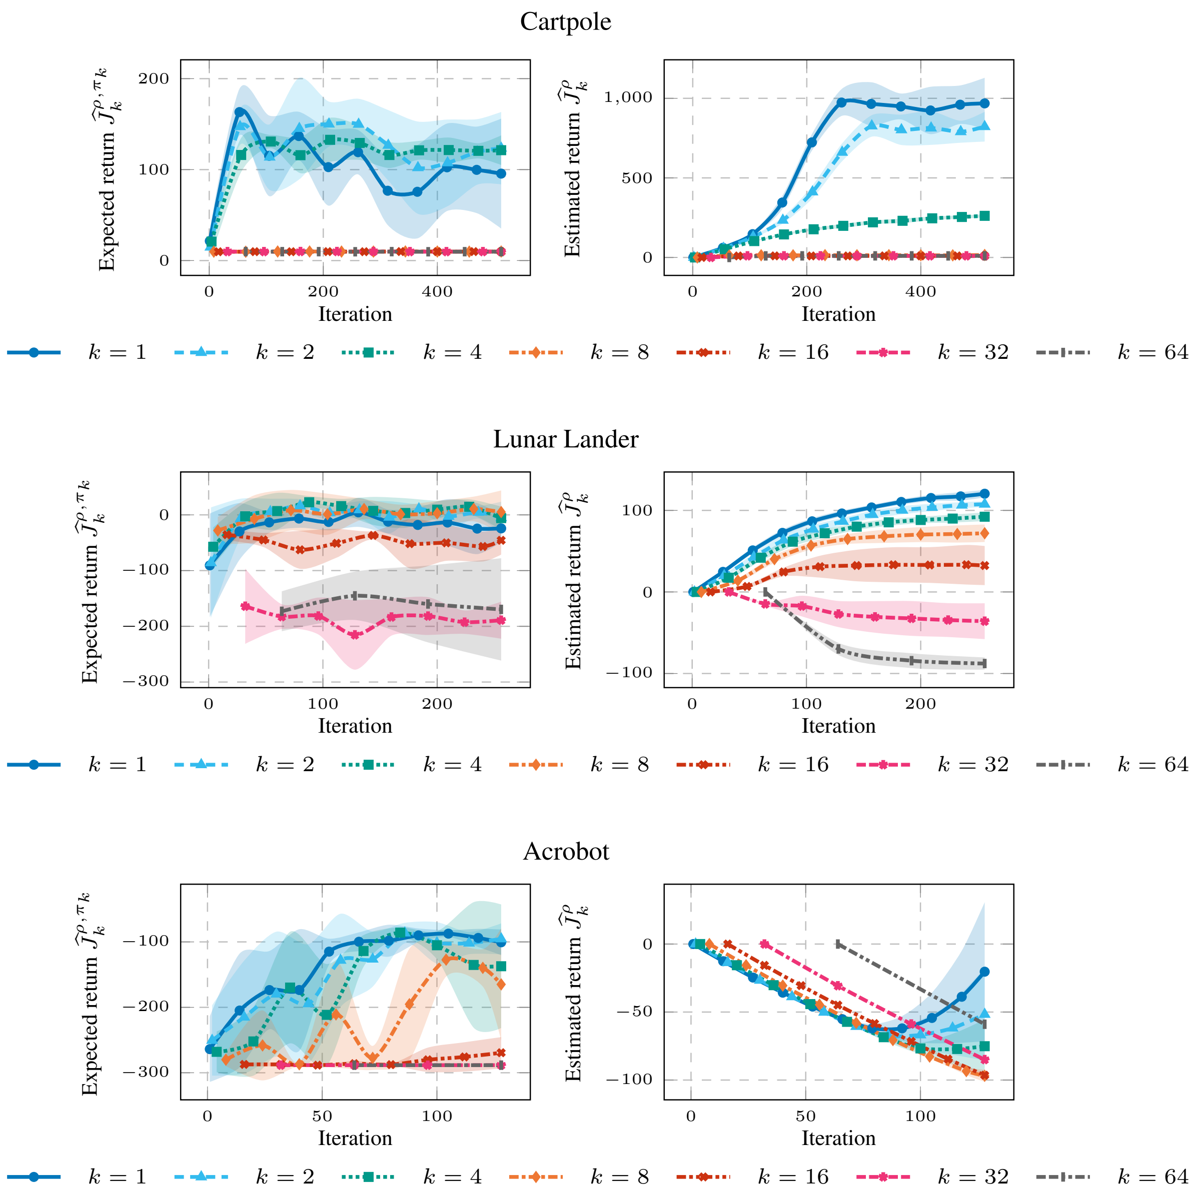

Figure 6. Expected return $\widehat{J}_k^{\rho,\pi_k}$ and estimated return $\widehat{J}_k^{\rho}$ for the different experiments with neural network, as a function of the number of iterations for different persistences. 20 runs, 95 % c.i.

D.4. Performance Dependence on Batch Size

In previous experiments we assumed we could choose the batch size, however, in real contexts this is not always allowed. In PFQI, lower batch sizes increase the estimation error, but the effect can change according to the used persistence. We wanted to investigate how the batch size influences the performance of PFQI policies for different persistences. Therefore,

Table 3. Results of PFQI execution in different environments and persistences with neural network as regressor. For each persistence k, we report the sample mean and the standard deviation of the estimated return of the last policy $\widehat{J}_k^{\rho,\pi_k}$. For each environment, the persistence with the highest average performance and the ones that are not statistically significantly different from that one (Welch's t-test with $p<0.05$) are in bold.

Environment	Expected return at persistence k ($\widehat{J}_k^{\rho,\pi_k}$, mean ± std)						
	$k=1$	$k=2$	$k=4$	$k=8$	$k=16$	$k=32$	$k=64$
Cartpole	$\mathbf{95.6 \pm 21.8}$	$\mathbf{123.6 \pm 14.4}$	$\mathbf{121.4 \pm 5.9}$	10.0 ± 0.1	9.7 ± 0.0	9.8 ± 0.1	9.8 ± 0.0
LunarLander	-24.3 ± 8.8	$\mathbf{-5.3 \pm 10.4}$	$\mathbf{-5.9 \pm 7.4}$	$\mathbf{5.0 \pm 14.0}$	-45.7 ± 9.2	-189.0 ± 12.0	-169.7 ± 33.1
Acrobot	$\mathbf{-100.6 \pm 7.1}$	$\mathbf{-95.4 \pm 8.5}$	$\mathbf{-137.2 \pm 34.0}$	-164.9 ± 30.9	-269.4 ± 8.6	-288.4 ± 0.0	-288.4 ± 0.0

we run PFQI on the Trading environment (described below) changing the number of sampled trajectories. As it can be noticed in Figure 7, if the batch size is small $(10, 50, 100)$, higher persistences $(2, 4, 8)$ results in better performances, while, with persistence 1, performance decreases with the iterations. In particular, with 50 trajectories, we can notice how all persistences except from 1 obtain a positive gain.

FX Trading Environment Description This environment simulates trading on a foreign exchange market. Trader's own currency is USD and it can be traded with EUR. The trader can be in three different position w.r.t. the foreign currency: long, short or flat, indicated, respectively, with $1, -1, 0$. Short selling is possible, i.e., the agent can sell a stock it does not own. At each timestep the agent can choose its next position with its action a_t. The exchange rate at time t is p_t, and the reward is equal to $R_t = a_t(p_t - p_{t-1}) - f|a_t - a_{t-1}|$, where the first term is the profit or loss given by the action a_t, and the second term represents the transaction costs, where f is a proportionality constant set to $4 \cdot 10^{-5}$. A timestep corresponds to 1 minute, an episode corresponds to a work day and it is composed by 1170 steps. It is assumed that at each time-step the trader goes long or short of the same unitary amount, thus the profits are not re-invested (and similarly for the losses), which means that the return is the sum of all the daily rewards (with a discount factor equal to 0.9999). The state consists of the last 60 minutes of price differences with the first price of the day ($p_t - p_0$), with the addition of the previous portfolio position as well as the fraction of time remaining until the end of the episode. For our experiments we sampled randomly daily episodes from a window of 64 work days of 2017, evaluating the performances on the last 20 days of the window.

Regressor Hyperparameters We used the class *ExtraTreesRegressor* in the *scikit-learn* library (Pedregosa et al., 2011) with the following parameters: n_estimators = 10, min_samples_split = 2, and min_samples_leaf = 2.

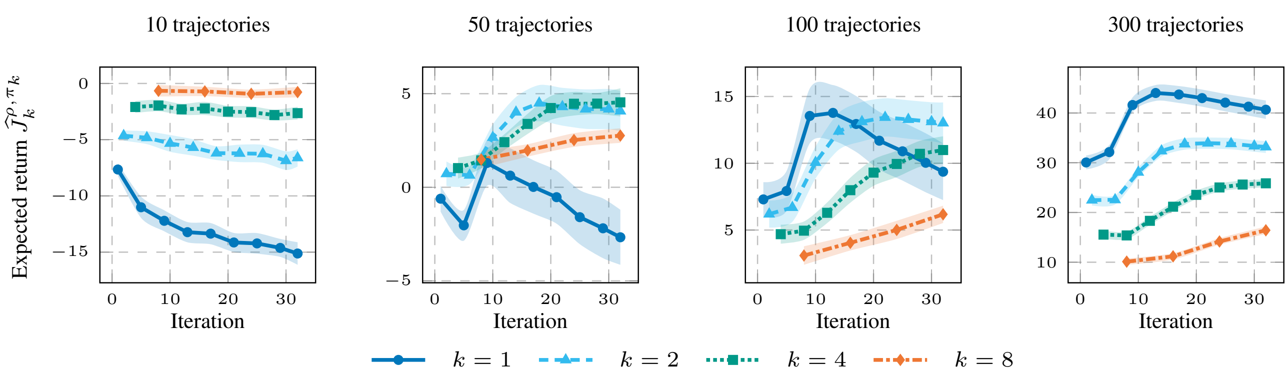

Figure 7. Performances for each persistence along the iterations, with different numbers of trajectories. 10 runs, 95% c.i.

E. Preliminary Results on Open Questions (Section 9)

In this appendix, we report some preliminary results related to the first two open questions about action persistence we presented in Section 9.

E.1. Improving Exploration with Persistence

As we already mentioned, action persistence might have an effect on the exploration properties of distribution ν used to collect samples. To avoid this phenomenon, in this work, we assumed to feed PFQI(k) with the same dataset collected in the base MDP $\mathcal{M}$, independently on which target persistence k we are interested in. In this appendix, we want to briefly analyze what happens when we feed standard FQI with a dataset collected by executing the same policy (e.g., the uniform policy over $\mathcal{A}$) in the k–persistent MDP $\mathcal{M}_k$,[9] in order to estimate the corresponding k–persistence action-value function Q^*_k. In this way, for each persistence k we have a different sampling distribution ν_k, but, being the dataset $\mathcal{D}_k \sim \nu_k$ collected in $\mathcal{M}_k$, we can apply standard FQI to estimate Q^*_k. Refer to Figure 8 for a graphical comparison between PFQI(k) executed in the base MDP and FQI executed in the k–persistent MDP.

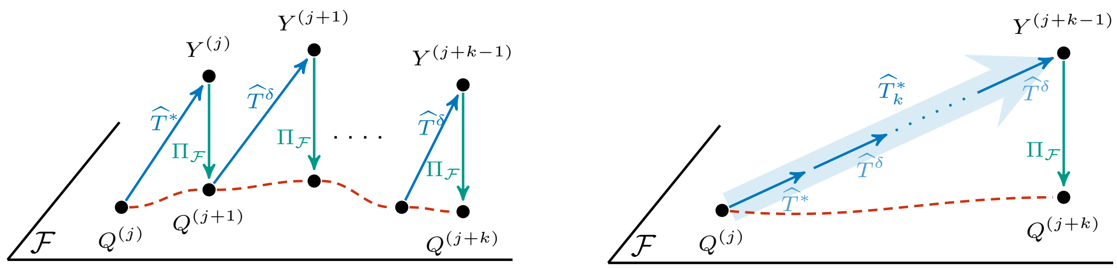

(a) PFQI(k) on $\mathcal{M}$ (b) FQI on $\mathcal{M}_k$

Figure 8. Illustration of (a) PFQI(k) executed in the base MDP $\mathcal{M}$ and (b) the standard FQI executed in the k-persistent MDP $\mathcal{M}_k$.

When we compare the performances of the policies obtained with different persistence levels learned starting with a dataset $\mathcal{D}_k \sim \nu_k$, we should consider two different effects: i) how training samples are generated (i.e., the sampling distribution ν_k, which changes for every persistence k); ii) how they affect the learning process in FQI. Unfortunately, in this setting we are not able to separate the two effects.

Our goal, in this appendix, is to compare for different values of $k \in \mathcal{K} = \{1, 2, \dots 64\}$ the performance of PFQI(k) and the performance of FQI run on the k–persistent MDP $\mathcal{M}_k$. The experimental setting is the same as in Appendix D, apart from the "sampling persistence" which is set to 1 also for the Mountain Car environment. In Figure 9, we show the performance at the end of training of the policies obtained with PFQI(k), the one derived with FQI on $\mathcal{M}_k$, and the uniform policy over the action space. First of all, we observe that when $k=1$, executing FQI on $\mathcal{M}_1$ is in all regards equivalent to executing PFQI(1) on $\mathcal{M}$, since PFQI(1) is FQI and $\mathcal{M}_1$ is $\mathcal{M}$. We can see that in the Cartpole environment, fixing a value of $k \in \mathcal{K}$, there is no significant difference in the performances obtained with PFQI(k) and FQI on $\mathcal{M}_k$. The behavior is significantly different when considering Mountain Car. Indeed, we notice that only FQI on $\mathcal{M}_k$ is able to learn a policy that reaches the goal for some specific values of $k \in \mathcal{K}$. We can justify this behavior with the fact that by collecting samples at a persistence k, like in FQI on $\mathcal{M}_k$, the exploration properties of the sampling distribution change, as we can see from the line "Uniform policy". If the input dataset contains no trajectory reaching the goal, our algorithms cannot solve the task. This is why PFQI(k), that uses persistence 1 to collect the samples, is unable to learn at all.[10]

This experiment gives a preliminary hint on how action persistence can affect exploration. More in general, we wonder which are the necessary characteristics of the environment such that the same sampling policy (e.g., the uniform policy over $\mathcal{A}$) allows to perform a better exploration. More formally, we ask ourselves how the persistence affects the entropy of the stationary distribution induced by the sampling policy.

[9] This procedure generates a different dataset compared to the case in which we use a "sampling persistence" $k_{\text{sampling}} > 1$, as illustrated in Appendix D.1. Indeed, in this case we do not record in the dataset the intermediate repeated actions, since we want a dataset of transition of the k–persistent MDP $\mathcal{M}_k$.

[10] Recall that in our main experiments (Appendix D.1), we had to employ for the Mountain Car a "sampling persistence" $k_{\text{sampling}} = 8$. Indeed, for $k_{\text{sampling}} \in \{1, 2, 4\}$ the uniform policy is unable to reach the goal, while for $k_{\text{sampling}} = 8$ it allows reaching the goal in the 6% of the times on average.

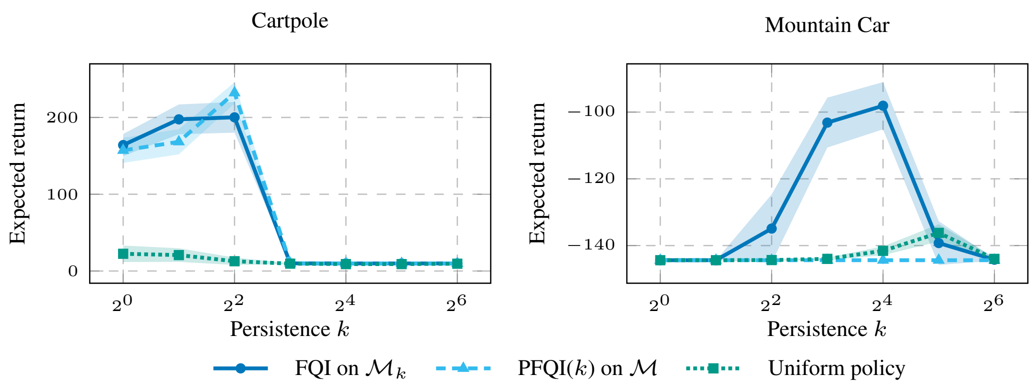

Figure 9. Performance of the policies learned with FQI on $\mathcal{M}_k$, PFQI(k) on $\mathcal{M}$ and the one of the uniform policies for different values of the persistence $k \in \mathcal{K}$. 10 runs. 95% c.i.

Table 4. Results of PFQI execution of the policy π_k learned with the k–persistent operator in the k'–persistent MDP $\mathcal{M}_{k'}$ in the Cartpole experiment. For each k, we report the sample mean and the standard deviation of the estimated return of the last policy $\widehat{J}^{\rho,\pi_k}_{k'}$. For each k, the persistence k' with the highest average performance and the ones k' that are not statistically significantly different from that one (Welch's t-test with $p<0.05$) are in bold.

	$k'=1$	$k'=2$	$k'=4$	$k'=8$	$k'=16$	$k'=32$	$k'=64$	$k'=128$	$k'=256$
$k=1$	$\mathbf{172.0 \pm 6.8}$	$\mathbf{174.1 \pm 6.5}$	113.0 ± 5.3	9.8 ± 0.0	9.7 ± 0.0	9.7 ± 0.1	9.8 ± 0.0	9.7 ± 0.0	9.7 ± 0.0
$k=2$	$\mathbf{178.4 \pm 6.7}$	$\mathbf{182.2 \pm 7.2}$	151.6 ± 5.1	9.9 ± 0.0	9.8 ± 0.0	9.8 ± 0.0	9.8 ± 0.0	9.8 ± 0.0	9.8 ± 0.0
$k=4$	276.2 ± 3.8	$\mathbf{287.3 \pm 1.1}$	237.0 ± 5.4	10.0 ± 0.0	9.8 ± 0.0	9.8 ± 0.0	9.9 ± 0.0	9.8 ± 0.0	9.9 ± 0.0
$k=8$	$\mathbf{284.3 \pm 1.6}$	$\mathbf{281.4 \pm 3.0}$	211.5 ± 4.0	10.0 ± 0.0	9.8 ± 0.0	9.8 ± 0.0	9.8 ± 0.0	9.8 ± 0.0	9.9 ± 0.0
$k=16$	$\mathbf{285.9 \pm 1.1}$	$\mathbf{282.9 \pm 2.6}$	223.5 ± 3.2	10.0 ± 0.0	9.9 ± 0.0	9.8 ± 0.0	9.9 ± 0.0	9.9 ± 0.0	9.8 ± 0.0
$k=32$	$\mathbf{285.7 \pm 1.3}$	$\mathbf{283.6 \pm 2.7}$	222.2 ± 3.6	10.0 ± 0.0	9.9 ± 0.0	9.9 ± 0.0	9.8 ± 0.0	9.9 ± 0.0	9.9 ± 0.0
$k=64$	$\mathbf{283.6 \pm 2.3}$	$\mathbf{284.1 \pm 2.0}$	225.5 ± 4.4	10.0 ± 0.0	9.9 ± 0.0	9.8 ± 0.0	9.9 ± 0.0	9.8 ± 0.0	9.9 ± 0.0
$k=128$	$\mathbf{282.9 \pm 2.2}$	$\mathbf{282.5 \pm 3.1}$	221.9 ± 4.7	10.0 ± 0.0	9.8 ± 0.0	9.9 ± 0.0	9.9 ± 0.0	9.9 ± 0.0	9.9 ± 0.0
$k=256$	$\mathbf{282.5 \pm 2.3}$	$\mathbf{283.4 \pm 2.4}$	224.3 ± 3.9	10.0 ± 0.0	9.9 ± 0.0	9.9 ± 0.0	9.9 ± 0.0	9.9 ± 0.0	9.9 ± 0.0

E.2. Learn in $\mathcal{M}_k$ and execute in $\mathcal{M}_{k'}$

In this appendix, we empirically analyze what happens when a policy is learned by PFQI with a certain persistence level k and executed later on with a different persistence level $k' \neq k$. We consider an experiment on the Cartpole environment, in the same setting as Appendix D. We run PFQI(k) for $k \in \mathcal{K} = \{1, 2, \dots, 256\}$ and then for each k we execute policy π_k (i.e., the policy learned by applying the k–persistent operator) in the k'–persistent MDP $\mathcal{M}_{k'}$ for $k' \in \mathcal{K}$. The results are shown in Table 4. Thus, for each pair (k, k'), Table 4 shows the sample mean and the sample standard deviation over 20 runs of the expected return of policy π_k in MDP $\mathcal{M}_k$, i.e., $J^{\rho,\pi_k}_{k'}$. First of all, let us observe that the diagonal of Table 4 corresponds to the first row of Table 1 (apart from the randomness due to the evaluation). If we take a row k, i.e., we fix the persistence of the operator, we notice that, in the majority of the cases, the persistence k' of the MDP yielding the best performance is smaller than k. Moreover, even if we learn a policy with the operator at a given persistence k and we see that such a policy displays a poor performance in the k–persistent MDP (e.g., for $k \geqslant 8$), when we reduce the persistence, the performance of that policy seems to improve.

Figure 10 compares for different values of k, determining the persistence of the operator, the performance of the policy π_k when we execute it in $\mathcal{M}_k$ and the performance of π_k in the MDP $\mathcal{M}_{(k')^*}$, where $(k')^* \in \operatorname{argmax}_{k' \in \mathcal{K}} \widehat{J}^{\rho,\pi_k}_{k'}$. We clearly see that suitably selecting the persistence k' of the MDP in which we will deploy the policy, allows reaching higher performances.

The question we wonder is whether this behavior is a property of the Cartpole environment or is a general phenomenon that we expect to occur in environments with certain characteristics. If so, which are those characteristics? Furthermore, when

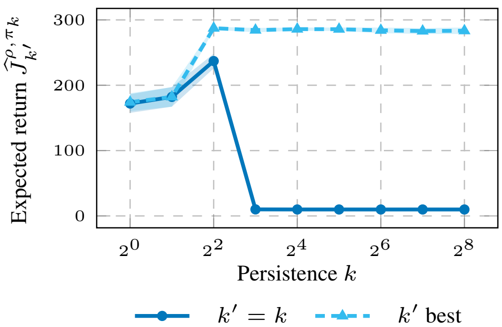

Figure 10. Performance of the policies π_k for $k \in \mathcal{K}$ comparing when they are executed in $\mathcal{M}_k$ and when they are executed in $\mathcal{M}_{(k')^*}$. 20 runs, 95% c.i.

we allow executing π_k in $\mathcal{M}_{k'}$ we should rephrase the persistence selection problem (Equation (13)) as follows:

$$k^*, (k')^* \in \operatorname*{arg\,max}_{k,k' \in \mathcal{K}} J^{\rho,\pi_k}_{k'}, \quad \rho \in \mathscr{P}(\mathcal{S}). \tag{23}$$

Similarly to the case of Equation (13), we cannot directly solve the problem if we are not allowed to interact with the environment. Is it possible to extend Lemma 6.1 and the subsequent heuristic simplifications to get a usable index $B_{k,k'}$ similar to Equation (15)?