

Configurable Markov Decision Processes

Alberto Maria Metelli*

Mirco Mutti*

Marcello Restelli

ALBERTOMARIA.METELLI@POLIMI.IT

MIRCO.MUTTI@MAIL.POLIMI.IT

MARCELLO.RESTELLI@POLIMI.IT

Department of Electronics, Information and Bioengineering, Politecnico di Milano, Milan, Italy

Abstract

In many real-world problems, there is the possibility to configure, to a limited extent, some environmental parameters to improve the performance of a learning agent. In this paper, we propose a novel framework, Configurable Markov Decision Processes (Conf-MDPs), to model this new type of interaction with the environment. Furthermore, we provide a new learning algorithm, Safe Policy-Model Iteration (SPMI), to jointly and adaptively optimize the policy and the environment configuration. After having introduced our approach, we present the experimental evaluation in two explicative problems to show the benefits of the environment configurability on the performance of the learned policy.

1. Introduction

Markov Decision Processes (MDPs) (Puterman, 2014) are a popular formalism to model sequential decision-making problems. Solving an MDP means to find a *policy*, i.e., a prescription of actions, that maximizes a given utility function. Typically, the environment dynamics is assumed to be fixed, unknown and out of the control of the agent. However, there exist several real-world scenarios in which the environment is *partially controllable* and, therefore, it might be beneficial to configure some of its features in order to select the most convenient MDP to solve. For instance, a human car driver has at her/his disposal a number of possible vehicle configurations she/he can act on (e.g., seasonal tires, stability and vehicle attitude, engine model, automatic speed control, parking aid system) to improve the driving style or quicken the process of learning a good driving policy. Another example is the interaction between a student and an automatic teaching system: the teaching model can be tailored to improve the student's learning experience (e.g., increasing or decreasing the difficulties of the questions or the speed at which the concepts are presented). It is worth noting that the active entity in the configuration process might be the agent itself or an external supervisor guiding the learning process. In a more abstract sense, the possibility to act on the environmental parameters can have essentially two benefits: i) it allows improving the agent performance; ii) it may allow to speed up the learning process.

In this paper, we propose a framework to model a *Configurable Markov Decision Process* (Conf-MDP), i.e., an MDP in which the environment can be configured to a limited extent. In principle, any of the Conf-MDP's parameters can be tuned, but we restrict our attention to the transition model and we focus to the problem of identifying the environment that allows achieving the highest performance possible. At an intuitive level, there exists a tight connection between environment and policy: variations of the environment induce modifications of the optimal policy. Furthermore, even for the same task, in presence of

*. Equal contribution.

agents having access to different policy spaces, the optimal environment might be different.¹ The spirit of this work is to investigate and exercise the tight connection between policy and model, pursuing the goal of improving the final policy performance.

After having introduced the definition of Conf-MDP, we propose a method to *jointly* and *adaptively* optimize the policy and the transition model, named *Safe Policy-Model Iteration* (SPMI). The algorithm adopts a safe learning approach (Pirodda et al., 2013) based on the maximization of a lower bound on the guaranteed performance improvement, yielding a sequence of model-policy pairs with monotonically increasing performance. The safe learning perspective makes our approach suitable for critical applications where performance degradation during learning is not allowed (e.g., industrial scenarios where extensive exploration of the policy space might damage the machinery).

In the standard Reinforcement Learning (RL) framework (Sutton and Barto, 1998), the usage of a lower bound to guide the choice of the policy has been first introduced by Conservative Policy Iteration (CPI) (Kakade and Langford, 2002), improved by Safe Policy Iteration (SPI) (Pirodda et al., 2013) and subsequently exploited by Ghavamzadeh et al. (2016); Abbasi-Yadkori et al. (2016); Papini et al. (2017). These methods revealed their potential thanks to the preference towards small policy updates, preventing from moving in a single step too far away from the current policy and avoiding premature convergence to suboptimal policies. A similar rationale is at the basis of Relative Entropy Policy Search (REPS) (Peters et al., 2010), and, more recently, Trust Region Policy Optimization (TRPO) (Schulman et al., 2015) and Proximal Policy Optimization (PPO) (Schulman et al., 2017). In order to introduce our framework and highlight its benefits, we limit our analysis to the scenario in which the model space (and the policy space) is known. However, when the model space is unknown, we could resort to a sample-based version of SPMI, which could be derived by adapting that of SPI (Pirodda et al., 2013).

We start in Section 2 by recalling some basic notions about MDPs and providing the definition of Conf-MDP. In Section 3 we derive the performance improvement bound and we outline the main features of SPMI (Section 4).² Then, we present the experimental evaluation (Section 5) in two explicative domains, representing simple abstractions of the main application of Conf-MDPs, with the purpose of showing how configuring the transition model can be beneficial for the final policy performance. A detailed analysis of the Conf-MDP framework, along with a more extensive theoretical and empirical study of SPMI can be found in (Metelli et al., 2018).

2. Preliminaries

A discrete-time Markov Decision Process (MDP) (Puterman, 2014) is defined as $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, R, \gamma, \mu)$ where \mathcal{S} is the state space, \mathcal{A} is the action space, $P(s'|s, a)$ is a Markovian transition model that defines the conditional distribution of the next state s' given the current state s and the current action a , $\gamma \in (0, 1)$ is the discount factor, $R(s, a) \in [0, 1]$ is the reward for performing action a in state s and μ is the distribution of the initial state. A

1. In general, a modification of the environment (e.g., changing the configuration of a car) is more expensive and more constrained w.r.t. to a modification of the policy.
 2. The proofs of all the lemmas and theorems can be found in Appendix B.

policy $\pi(a|s)$ defines the probability distribution of an action a given the current state s . We now formalize the Configurable Markov Decision Process (Conf-MDP).

Definition 2.1 A Configurable Markov Decision Process is a tuple $\mathcal{CM} = (\mathcal{S}, \mathcal{A}, R, \gamma, \mu, \mathcal{P}, \Pi)$ where $(\mathcal{S}, \mathcal{A}, R, \gamma, \mu)$ is an MDP without the transition model and \mathcal{P} and Π are the model and policy spaces.

More specifically, Π is the set of policies the agent has access to and \mathcal{P} is the set of available environment configurations (transition models). The performance of a model-policy pair $(P, \pi) \in \mathcal{P} \times \Pi$ is evaluated through the *expected return*, i.e., the expected discounted sum of the rewards collected along a trajectory

$$J_\mu^{P,\pi} = \frac{1}{1-\gamma} \int_{\mathcal{S}} d_\mu^{P,\pi}(s) \int_{\mathcal{A}} \pi(a|s) R(s, a) da ds, \quad (1)$$

where $d_\mu^{P,\pi}$ is the γ -discounted state distribution (Sutton et al., 2000), defined recursively as $d_\mu^{P,\pi}(s) = (1-\gamma)\mu(s) + \gamma \int_{\mathcal{S}} d_\mu^{P,\pi}(s') P^\pi(s'|s) ds'$. We can also define the γ -discounted state-action distribution as $\delta_\mu^{P,\pi}(s, a) = \pi(a|s) d_\mu^{P,\pi}(s)$. While solving an MDP consists in finding a policy π^* that maximizes $J_\mu^{P,\pi}$ under the given fixed environment P , solving a Conf-MDP consists in finding a model-policy pair (P^*, π^*) such that $P^*, \pi^* = \arg \max_{P \in \mathcal{P}, \pi \in \Pi} J_\mu^{\pi, P}$.

For control purposes, the state-action value function, or *Q-function*, is introduced as the expected return starting from a state s and performing action a :

$$Q^{P,\pi}(s, a) = R(s, a) + \gamma \int_{\mathcal{S}} P(s'|s, a) V^{P,\pi}(s') ds'. \quad (2)$$

For learning the transition model we introduce the state-action-next-state value function or *U-function*, defined as the expected return starting from the state s , performing action a and landing to state s' :

$$U^{P,\pi}(s, a, s') = R(s, a) + \gamma V^{P,\pi}(s'), \quad (3)$$

where $V^{P,\pi}$ is the state value function or *V-function*. These three functions are tightly connected due to the trivial relations: $V^{P,\pi}(s) = \int_{\mathcal{A}} \pi(a|s) Q^{P,\pi}(s, a) da$ and $Q^{P,\pi}(s, a) = \int_{\mathcal{S}} P(s'|s, a) U^{P,\pi}(s, a, s') ds'$. Furthermore, we define the *policy advantage function* $A^{P,\pi}(s, a) = Q^{P,\pi}(s, a) - V^{P,\pi}(s)$ that quantifies how much an action is better than the others and the *model advantage function* $A^{P,\pi}(s, a, s') = U^{P,\pi}(s, a, s') - Q^{P,\pi}(s, a)$ that quantifies how much the next state is better than the other ones.

In order to evaluate the *one-step improvement* in performance attained by a new policy π' or model P' when the current policy is π and the current model is P , we introduce the *relative advantage functions* (Kakade and Langford, 2002): $A_{P,\pi}^{P,\pi'}(s) = \int_{\mathcal{A}} \pi'(a|s) A^{P,\pi}(s, a) da$, $A_{P,\pi}^{P',\pi}(s, a) = \int_{\mathcal{S}} P'(s'|s, a) A^{P,\pi}(s, a, s') ds'$, and the corresponding expected values under the γ -discounted distributions: $\mathbb{A}_{P,\pi,\mu}^{P,\pi'} = \int_{\mathcal{S}} d_\mu^{P,\pi}(s) A_{P,\pi}^{P,\pi'}(s) ds$ and $\mathbb{A}_{P,\pi,\mu}^{P',\pi} = \int_{\mathcal{S}} \int_{\mathcal{A}} \delta_\mu^{P,\pi}(s, a) A_{P,\pi}^{P',\pi}(s, a) ds da$.

3. Performance Improvement

The goal of this section is to provide a lower bound to the performance improvement obtained by moving from a model-policy pair (P, π) to another pair (P', π') .³ We start providing a bound for the difference of γ -discounted distributions under different model-policy pairs.

Proposition 3.1 *Let (P, π) and (P', π') be two model-policy pairs, the ℓ^1 -norm of the difference between the γ -discounted state distributions can be upper bounded as:*

$$\left\| d_\mu^{P', \pi'} - d_\mu^{P, \pi} \right\|_1 \leq \frac{\gamma}{1 - \gamma} \left(D_{\mathbb{E}}^{\pi', \pi} + D_{\mathbb{E}}^{P', P} \right),$$

where $D_{\mathbb{E}}^{\pi', \pi} = \mathbb{E}_{s \sim d_\mu^{P, \pi}} \left\| \pi'(\cdot|s) - \pi(\cdot|s) \right\|_1$ and $D_{\mathbb{E}}^{P', P} = \mathbb{E}_{(s, a) \sim \delta_\mu^{P, \pi}} \left\| P'(\cdot|s, a) - P(\cdot|s, a) \right\|_1$.

We now exploit the previous result to obtain a lower bound on the performance improvement as an effect of the policy and model updates.

Theorem 3.1 (Decoupled Bound) *The performance improvement of model-policy pair (P', π') over (P, π) can be lower bounded as:*

$$\underbrace{J_\mu^{P', \pi'} - J_\mu^{P, \pi}}_{\text{performance improvement}} \geq B(P', \pi') = \underbrace{\frac{\mathbb{A}_{P, \pi, \mu}^{P', \pi} + \mathbb{A}_{P, \pi, \mu}^{P, \pi'}}{1 - \gamma}}_{\text{advantage}} - \underbrace{\frac{\gamma \Delta Q^{P, \pi} D}{2(1 - \gamma)^2}}_{\text{dissimilarity penalization}},$$

where D is a dissimilarity term defined as $D = D_{\mathbb{E}}^{\pi', \pi} (D_\infty^{\pi', \pi} + D_\infty^{P', P}) + D_{\mathbb{E}}^{P', P} (D_\infty^{\pi', \pi} + \gamma D_\infty^{P', P})$ and $D_\infty^{\pi', \pi} = \sup_{s \in \mathcal{S}} \left\| \pi'(\cdot|s) - \pi(\cdot|s) \right\|_1$, $D_\infty^{P', P} = \sup_{s \in \mathcal{S}, a \in \mathcal{A}} \left\| P'(\cdot|s, a) - P(\cdot|s, a) \right\|_1$ and $\Delta Q^{P, \pi} = \sup_{s, s' \in \mathcal{S}, a, a' \in \mathcal{A}} \left| Q^{P, \pi}(s', a') - Q^{P, \pi}(s, a) \right|$.

The bound is composed of two terms, like in (Kakade and Langford, 2002; Pirotta et al., 2013): the first term, *advantage*, represents how much gain in performance can be locally obtained by moving from (P, π) to (P', π') , whereas the second term, *dissimilarity penalization*, discourages updates towards model-policy pairs that are too far away.

4. Safe Policy Model Iteration

In this section, we present a new safe learning algorithm for the Conf-MDP framework, *Safe Policy-Model Iteration* (SPMI), inspired by (Pirotta et al., 2013), capable of learning the policy and the model simultaneously.⁴

4.1 The Algorithm

Following (Pirotta et al., 2013), we define the policy and model improvement update rules: $\pi' = \alpha \bar{\pi} + (1 - \alpha) \pi$, $P' = \beta \bar{P} + (1 - \beta) P$, where $\alpha, \beta \in [0, 1]$, $\bar{\pi} \in \Pi$ and $\bar{P} \in \mathcal{P}$ are the target policy and the target model respectively. Extending the rationale of (Pirotta et al., 2013) to our context, we aim to determine the values of α and β which jointly maximize the *decoupled bound* (Theorem 3.1). In the following we will abbreviate $B(P', \pi')$ with $B(\alpha, \beta)$.

3. The proofs of the reported propositions along with other related results can be found in Appendix B.

4. In the context of Conf-MDPs we believe that knowing the model of the configurable part of the environment is a reasonable requirement.

Table 1: The four possible optimal (α, β) pairs, the optimal pair is the one yielding the maximum bound value (all values are clipped in $[0, 1]$).

$\beta^* = 0$	$\alpha_0^* = \frac{(1-\gamma)A_{P,\pi,\mu}^{P,\bar{\pi}}}{\gamma\Delta Q^{P,\pi}D_{\infty}^{\bar{\pi},\pi}D_E^{\bar{\pi},\pi}}$	$\beta^* = 1$	$\alpha_1^* = \alpha_0^* - \frac{1}{2}\left(\frac{D_E^{\bar{P},P}}{D_{\infty}^{\bar{\pi},\pi}} + \frac{D_{\infty}^{\bar{P},P}}{D_E^{\bar{\pi},\pi}}\right)$
$\alpha^* = 0$	$\beta_0^* = \frac{(1-\gamma)A_{P,\pi,\mu}^{\bar{P},\pi}}{\gamma^2\Delta Q^{P,\pi}D_{\infty}^{\bar{P},P}D_E^{\bar{P},P}}$	$\alpha^* = 1$	$\beta_1^* = \beta_0^* - \frac{1}{2\gamma}\left(\frac{D_E^{\bar{\pi},\pi}}{D_{\infty}^{\bar{P},P}} + \frac{D_{\infty}^{\bar{\pi},\pi}}{D_E^{\bar{P},P}}\right)$

Theorem 4.1 For any $\bar{\pi} \in \Pi$ and $\bar{P} \in \mathcal{P}$, the decoupled bound is optimized for $\alpha^*, \beta^* = \arg \max_{\alpha, \beta} \{B(\alpha, \beta) : (\alpha, \beta) \in \mathcal{V}\}$, where $\mathcal{V} = \{(\alpha_0^*, 0), (\alpha_1^*, 1), (0, \beta_0^*), (1, \beta_1^*)\}$ and the values of $\alpha_0^*, \alpha_1^*, \beta_0^*$ and β_1^* are reported in Table 1.

The theorem expresses the fact that the optimal (α, β) pair lies on the boundary of $[0, 1] \times [0, 1]$, i.e., either one between policy and model is moved and the other is kept unchanged or one is moved and the other is set to target. Algorithm 1 reports the basic structure of SPMI. The procedures *PolicyChooser* and *ModelChooser* are designated for selecting the target policy and model (see Section 4.3).

4.2 Policy and Model Spaces

The selection of the target policy and model is a rather crucial component of the algorithm since the quality of the updates largely depends on it. To effectively adopt a target selection strategy we have to know which are the degrees of freedom on the policy and model spaces. Focusing on the model space first, it is easy to discriminate two macro-classes in real-world scenarios. In some cases, there are almost no constraints on the direction in which to update the model. In other cases, only a limited model portion, typically a set of parameters inducing transition probabilities, can be accessed. While we can naturally design the first scenario as an *unconstrained* model space, to represent the second scenario we limit the model space to the convex hull $\text{co}(\mathbf{P})$, where \mathbf{P} is a set of extreme (or vertex) models. Since only the convex combination coefficients can be controlled, we refer to the latter as a *parametric* model space. It is noteworthy that we can symmetrically extend the dichotomy to the policy space, although the need to limit the agent on the direction of policy updates is less significant in our perspective.

4.3 Target Choice

To deal with unconstrained spaces, it is quite natural to adopt the target selection strategy presented in (Pirootta et al., 2013), by introducing the concept of greedy model as $P^+(\cdot|s, a) \in \arg \max_{s' \in \mathcal{S}} U^{P,\pi}(s, a, s')$, i.e., the model that maximizes the relative advantage in each state-action pair. At each step, the greedy policy and model w.r.t. the $Q^{P,\pi}$ and $U^{P,\pi}$ are selected as targets. When we are not free to choose the greedy model, like in the parametric setting,

Algorithm 1 Safe Policy Model Iteration

```

initialize  $\pi_0, P_0$ .
for  $i = 0, 1, 2, \dots$  until  $\epsilon$ -convergence do
     $\bar{\pi}_i = \text{PolicyChooser}(\pi_i)$ 
     $\bar{P}_i = \text{ModelChooser}(P_i)$ 
     $\mathcal{V}_i = \{(\alpha_{0,i}^*, 0), (\alpha_{1,i}^*, 1), (0, \beta_{0,i}^*), (1, \beta_{1,i}^*)\}$ 
     $\alpha_i^*, \beta_i^* = \arg \max_{\alpha, \beta} \{B(\alpha, \beta) : (\alpha, \beta) \in \mathcal{V}_i\}$ 
     $\pi_{i+1} = \alpha_i^* \bar{\pi}_i + (1 - \alpha_i^*) \pi_i$ 
     $P_{i+1} = \beta_i^* \bar{P}_i + (1 - \beta_i^*) P_i$ 
end for
    
```




Figure 1: Expected return for the Student-Teacher domain for different update strategies and target selections.

we select the vertex model that maximizes the expected relative advantage (*greedy choice*). The greedy strategy is based on local information and is not guaranteed to provide a model-policy pair maximizing the bound. However, testing all the model-policy pairs is highly inefficient in the presence of large model-policy spaces. A reasonable compromise is to select, as a target, the model that yields the maximum bound value between the greedy target $\bar{P}_i \in \arg \max_{P \in \mathcal{P}} \mathbb{A}_{P_i, \pi, \mu}^{P, \pi}$ and the previous target \bar{P}_{i-1} (the same procedure can be employed for the policy). This procedure, named *persistent choice*, effectively avoids the oscillating behavior, common with the greedy choice.

5. Experimental Evaluation

The goal of this section is to show the benefits of configuring the environment while the policy learning goes on. The experiments are conducted on two explicative domains: the Student-Teacher domain (unconstrained model space) the Racetrack Simulator (parametric model space). We compare different *target choices* (greedy and persistent, see Section 4.3) and different *update strategies*. Specifically, SPMI, that adaptively updates policy and model, is compared with some alternative model learning approaches: SPMI-alt(ernated) in which model and policy updates are forced to be alternated, SPMI-sup that uses a looser bound, obtained from Theorem 3.1 by replacing $D_{\mathbb{E}}^{\star', \star}$ with $D_{\infty}^{\star', \star}$, SPI+SMI⁵ that optimizes policy and model in sequence and SMI+SPI that does the opposite.

5.1 Student-Teacher domain

The Student-Teacher domain is a simple model of concept learning, inspired by (Rafferty et al., 2011). A student (agent) learns to perform consistent assignments to literals as a result of the statements (e.g., “A+C=3”) provided by an automatic teacher (environment, e.g., online platform). The student has a limited policy space as she/he can only change the values of the literals by a finite quantity, but it is possible to configure the difficulty of the teacher’s statements, selecting the number of literals in the statement (see Appendix D.1).

We consider the illustrative example in which there are two binary literals, and the student can change only one literal at a time. In Figure 1, we show the behavior of the different update strategies starting from a uniform initialization. We first notice (Figure 1

5. SMI (Safe Model Iteration) is SPMI without policy updates.




Figure 2: Expected return and coefficient of the high speed stability vertex model for different update strategies in track T1 and expected return for track T2.

(a)) that learning both policy and model is convenient since the performance of SPMI at convergence is higher than the one of SPI (only policy learned) and SMI (only model learned). Moreover, we observe that the greedy target selection, while being the best local choice maximizing the advantage, might result in an unstable behavior that slows down the convergence of the algorithm, while the persistent target selection enforces more stability (Figure 1 (b)). Finally, we can see that the joint and adaptive strategy of SPMI outperforms SPMI-sup and SPMI-alt (Figure 1 (c)). Notably, both SPMI and SPMI-sup perform the policy updates and the model updates in sequence, since it is more convenient for the student to learn an almost optimal policy with no intervention on the teacher and then refining the teacher model to gain further reward. For this reason, the alternated model-policy update (SPMI-alt) results in poorer performance in an early stage.

5.2 Racetrack simulator

The Racetrack simulator is a simple abstraction of a car driving problem. The driver (agent) has to optimize a driving policy to run the vehicle on the track, reaching the finish line as fast as possible. During the process, the agent can configure two vehicle settings to improve her/his driving performance: the *vehicle stability* and the *engine boost* (see Appendix D.2).

In Figure 2 (a), we highlight the effectiveness of SPMI w.r.t. all the alternatives on track T1, in which only the vehicle stability can be modified. Notably, comparing SPMI with the sequential approaches, SPI+SPI and SMI+SPI, we can easily deduce that it is not valuable to configure the vehicle stability, i.e., updating the model, while the driving policy is still really rough. In Figure 2 (b), we show the trend of the model coefficient related to high-speed stability. While the optimal configuration results in a mixed model for vehicle stability, SPMI exploits the maximal high-speed stability to efficiently learn the driving policy in an early stage, SPI+SPI and SPMI-greedy, instead, execute a sequence of policy updates and then directly lead the model to the optimal configuration. Figure 2 (c) shows how the previous considerations generalize to an example on a morphologically different track (T2), in which also the engine boost can be configured. The learning process is characterized by a long exploration phase, both in the model and the policy space, in which the driver cannot lead the vehicle to the finish line to collect any reward. Then, we observe a fast growth in expected return when the agent has acquired enough information to reach the finish line.

6. Discussion and Conclusions

In this paper, we proposed a novel framework (Conf-MDP), to model a set of real-world scenarios in which the environment dynamics can be partially modified, and a safe learning approach (SPMI), to jointly and adaptively learn model and policy. Conf-MDPs allow modelling many relevant sequential-decision making problems that we believe cannot be effectively addressed using traditional frameworks.

Why not a unique agent? Representing the environment configurability in the agent model when the environment is under the control of a supervisor is certainly inappropriate. Even when the environment configuration is carried out by the agent, this approach would require the inclusion of “configuration actions” in the action space to allow the agent to configure the environment directly as a part of the policy optimization. However, in our framework, the environment configuration is performed just once at the beginning of the episode. Moreover, with configuration actions the agent is not really learning a probability distribution on actions, i.e., a policy, but a probability distribution on state-state couples, i.e., a state kernel. This formulation prevents distinguishing, during the process, the effects of the policy from those of the model, making it difficult to finely constrain the configurations, limit the feasible model space, and recovering, a posteriori, the optimal model-policy pair.

Why not a multi-agent system? When there is no supervisor, the agent is the only learning entity and the environment is completely passive. In the presence of a supervisor, it would be misleading to adopt a cooperative multi-agent approach. The supervisor acts externally, at a different level and could be, possibly, totally transparent to the learning agent. Indeed, the supervisor does not operate inside the environment but it is in charge of selecting the most suitable configuration, whereas the agent needs to learn the optimal policy for the given environment.

The second contribution of this paper is the formulation of a safe approach, suitable to manage critical tasks, to solve a learning problem in the context of the newly introduced Conf-MDP framework. To this purpose, we proposed a novel tight lower bound on the performance improvement and an algorithm (SPMI) optimizing this bound to learn the policy and the model configuration simultaneously. We then presented an empirical study to show the effectiveness of SPMI in our context and to uphold the introduction of the Conf-MDP framework.

We believe that this work is only a first step in solving these kinds of problems: many future research directions are open. Clearly, a first extension could tackle the problem from a sample-based perspective, removing the requirement of knowing the full model space. Furthermore, we could consider different learning approaches, like policy search methods, suitable for continuous state-action spaces.

References

Yasin Abbasi-Yadkori, Peter L Bartlett, and Stephen J Wright. A fast and reliable policy improvement algorithm. In *Artificial Intelligence and Statistics*, pages 1338–1346, 2016.

- Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. Constrained policy optimization. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *ICML'17*, pages 22–31, 2017.
- Bruce Lee Bowerman. *Nonstationary Markov decision processes and related topics in nonstationary Markov chains*. PhD thesis, Iowa State University, 1974.
- Thiago P Bueno, Denis D Mauá, Leliane N Barros, and Fabio G Cozman. Modeling markov decision processes with imprecise probabilities using probabilistic logic programming. In *Proceedings of the Tenth International Symposium on Imprecise Probability: Theories and Applications*, pages 49–60, 2017.
- Torpong Cheevaprawatdomrong, Irwin E Schochetman, Robert L Smith, and Alfredo Garcia. Solution and forecast horizons for infinite-horizon nonhomogeneous markov decision processes. *Mathematics of Operations Research*, 32(1):51–72, 2007.
- Kamil Andrzej Ciosek and Shimon Whiteson. Offer: Off-environment reinforcement learning. In *AAAI*, pages 1819–1825, 2017.
- Karina Valdivia Delgado, Leliane Nunes de Barros, Fabio Gagliardi Cozman, and Ricardo Shirota. Representing and solving factored markov decision processes with imprecise probabilities. *Proceedings ISIPTA, Durham, United Kingdom*, pages 169–178, 2009.
- Carlos Florensa, David Held, Markus Wulfmeier, Michael Zhang, and Pieter Abbeel. Reverse curriculum generation for reinforcement learning. In *Conference on Robot Learning*, pages 482–495, 2017.
- Alfredo Garcia and Robert L Smith. Solving nonstationary infinite horizon dynamic optimization problems. *Journal of Mathematical Analysis and Applications*, 244(2): 304–317, 2000.
- Archis Ghatge and Robert L Smith. A linear programming approach to nonstationary infinite-horizon markov decision processes. *Operations Research*, 61(2):413–425, 2013.
- Mohammad Ghavamzadeh, Marek Petrik, and Yinlam Chow. Safe policy improvement by minimizing robust baseline regret. In *Advances in Neural Information Processing Systems*, pages 2298–2306, 2016.
- Robert Givan, Sonia Leach, and Thomas Dean. Bounded parameter markov decision processes. In Sam Steel and Rachid Alami, editors, *Recent Advances in AI Planning*, pages 234–246. Springer Berlin Heidelberg, 1997.
- Moshe Haviv and Ludo Van der Heyden. Perturbation bounds for the stationary probabilities of a finite markov chain. *Advances in Applied Probability*, 16(4):804–818, 1984.
- Wallace J Hopp, James C Bean, and Robert L Smith. A new optimality criterion for nonhomogeneous markov decision processes. *Operations Research*, 35(6):875–883, 1987.
- Sham Kakade and John Langford. Approximately optimal approximate reinforcement learning. In *Proceedings of the 19th International Conference on Machine Learning*, volume 2, pages 267–274, 2002.

- Alberto Maria Metelli, Mirco Mutti, and Marcello Restelli. Configurable Markov decision processes. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 3491–3500, Stockholmsmssan, Stockholm Sweden, 10–15 Jul 2018. PMLR. URL <http://proceedings.mlr.press/v80/metelli18a.html>.
- Yaodong Ni and Zhi-Qiang Liu. Bounded-parameter partially observable markov decision processes. In *Proceedings of the Eighteenth International Conference on International Conference on Automated Planning and Scheduling*, pages 240–247. AAAI Press, 2008.
- Takayuki Osogami. Robust partially observable markov decision process. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *ICML’15*, pages 106–115, 2015.
- Matteo Papini, Matteo Pirotta, and Marcello Restelli. Adaptive batch size for safe policy gradients. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 3594–3603. Curran Associates, Inc., 2017.
- Jan Peters, Katharina Mülling, and Yasemin Altun. Relative entropy policy search. In *AAAI*, pages 1607–1612. Atlanta, 2010.
- Matteo Pirotta, Marcello Restelli, Alessio Pecorino, and Daniele Calandriello. Safe policy iteration. In *Proceedings of the 30th International Conference on International Conference on Machine Learning*, volume 28 of *ICML’13*, pages 307–315, 2013.
- Martin L Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- Anna N Rafferty, Emma Brunskill, Thomas L Griffiths, and Patrick Shafto. Faster teaching by pomdp planning. In *AIED*, pages 280–287. Springer, 2011.
- Jay K Satia and Roy E Lave Jr. Markovian decision processes with uncertain transition probabilities. *Operations Research*, 21(3):728–740, 1973.
- John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *Proceedings of the 32nd International Conference on Machine Learning*, ICML’15, pages 1889–1897, 2015.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge, 1998.
- Richard S Sutton, David A McAllester, Satinder P Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. In *Advances in Neural Information Processing Systems*, pages 1057–1063, 2000.
- Chelsea C White III and Hany K Eldeib. Markov decision processes with imprecise transition probabilities. *Operations Research*, 42(4):739–749, 1994.

Appendix A. Related Works

Although in the standard framework the environment is assumed to be fixed, several exceptions can be found in the literature, especially in the context of Markov Decision Processes with imprecise probabilities (MDPIPs) (Satia and Lave Jr, 1973; White III and Eldeib, 1994; Bueno et al., 2017) and non-stationary environments (Bowerman, 1974; Hopp et al., 1987). In the former case, the transition kernel is known under uncertainty. Therefore, it cannot be specified using a conditional probability distribution, but it must be defined through a set of probability distributions (Delgado et al., 2009). In this context, Bounded-parameter Markov Decision Processes (BMDPs) consider a special case in which upper and lower bounds on transition probabilities are specified (Givan et al., 1997; Ni and Liu, 2008). A common approach is to solve a minimax problem to find a robust policy maximizing the expected return under the worst possible transition model (Osogami, 2015). In non-stationary environments, the transition probabilities (possibly also the reward function) change over time (Bowerman, 1974). Several works tackle the problem of defining an optimality criterion (Hopp et al., 1987) and finding optimal policies in non-stationary environments (Garcia and Smith, 2000; Cheevaprawatdomrong et al., 2007; Ghate and Smith, 2013). Although the environment is no longer fixed, both Markov decision processes with imprecise probabilities and non-stationary Markov decision processes do not admit the possibility to dynamically alter the environmental parameters. The environment can be modified also to favor the learning process. This instance has been previously addressed in (Ciosek and Whiteson, 2017; Florensa et al., 2017), where the transition model and the initial state distribution are altered in order to reach a faster convergence to the optimal policy. However, in both the cases the environment modification is only simulated, while the underlying environment dynamic remains unchanged.

Appendix B. Proofs and Derivations

In this appendix, we provide the proofs of the theorems and lemmas presented in the paper, we will also provide some comments on the intermediate results. Given a model-policy pair (P, π) we indicate with P^π the state kernel function defined as:

$$P^\pi(s'|s) = \int_{\mathcal{A}} \pi(a|s)P(s'|s, a)da. \quad (4)$$

Proposition B.1 *Let (P, π) and (P', π') be two model-policy pairs, the ℓ^1 -norm of the difference between the γ -discounted state distributions can be upper bounded as:*

$$\left\| d_\mu^{P', \pi'} - d_\mu^{P, \pi} \right\|_1 \leq \frac{\gamma}{1 - \gamma} D_{\mathbb{E}}^{P', \pi', P^\pi},$$

where $D_{\mathbb{E}}^{P', \pi', P^\pi} = \mathbb{E}_{s \sim d_\mu^{P, \pi}} \left\| P'^{\pi'}(\cdot|s) - P^\pi(\cdot|s) \right\|_1$.

Proof Exploiting the recursive equation of the γ -discounted state distribution we can write the distributions difference as follows:

$$\begin{aligned}
 d_{\mu}^{P',\pi'}(s) - d_{\mu}^{P,\pi}(s) &= (1 - \gamma)\mu(s) + \gamma \int_{\mathcal{S}} d_{\mu}^{P',\pi'}(s') P'^{\pi'}(s|s') ds' - (1 - \gamma)\mu(s) \\
 &\quad - \gamma \int_{\mathcal{S}} d_{\mu}^{P,\pi}(s') P^{\pi}(s|s') ds' \\
 &= \gamma \int_{\mathcal{S}} d_{\mu}^{P',\pi'}(s') P'^{\pi'}(s|s') ds' - \gamma \int_{\mathcal{S}} d_{\mu}^{P,\pi}(s') P^{\pi}(s|s') ds' \\
 &\quad \pm \gamma \int_{\mathcal{S}} d_{\mu}^{P,\pi}(s') P'^{\pi'}(s|s') ds' \\
 &= \gamma \int_{\mathcal{S}} (d_{\mu}^{P',\pi'}(s') - d_{\mu}^{P,\pi}(s')) P'^{\pi'}(s|s') ds' \\
 &\quad + \gamma \int_{\mathcal{S}} d_{\mu}^{P,\pi}(s') (P'^{\pi'}(s|s') - P^{\pi}(s|s')) ds'. \tag{5}
 \end{aligned}$$

Applying the ℓ^1 -norm to the equation (5) we can state the following:

$$\begin{aligned}
 \left\| d_{\mu}^{P',\pi'} - d_{\mu}^{P,\pi} \right\|_1 &\leq \gamma \int_{\mathcal{S}} \left| \int_{\mathcal{S}} (d_{\mu}^{P',\pi'}(s') - d_{\mu}^{P,\pi}(s')) P'^{\pi'}(s|s') ds' \right| ds \\
 &\quad + \gamma \int_{\mathcal{S}} \left| \int_{\mathcal{S}} d_{\mu}^{P,\pi}(s') (P'^{\pi'}(s|s') - P^{\pi}(s|s')) ds' \right| ds \tag{6}
 \end{aligned}$$

$$\begin{aligned}
 &\leq \gamma \int_{\mathcal{S}} \left| d_{\mu}^{P',\pi'}(s') - d_{\mu}^{P,\pi}(s') \right| \int_{\mathcal{S}} P'^{\pi'}(s|s') ds ds' \\
 &\quad + \gamma \int_{\mathcal{S}} d_{\mu}^{P,\pi}(s') \int_{\mathcal{S}} \left| P'^{\pi'}(s|s') - P^{\pi}(s|s') \right| ds ds' \tag{7}
 \end{aligned}$$

$$\begin{aligned}
 &\leq \gamma \left\| d_{\mu}^{P',\pi'} - d_{\mu}^{P,\pi} \right\|_1 + \gamma \mathbb{E}_{s \sim d_{\mu}^{P,\pi}} \left\| P'^{\pi'}(\cdot|s) - P^{\pi}(\cdot|s) \right\|_1 \tag{8} \\
 &\leq \frac{\gamma}{1 - \gamma} \mathbb{E}_{s \sim d_{\mu}^{P,\pi}} \left\| P'^{\pi'}(\cdot|s) - P^{\pi}(\cdot|s) \right\|_1 = D_{\mathbb{E}}^{P',\pi',P^{\pi}}.
 \end{aligned}$$

In (6) we exploited the subadditivity of the norm $\|x + y\| \leq \|x\| + \|y\|$ and (8) derives from (7) by observing that $\int_{\mathcal{S}} P'^{\pi'}(s|s') ds = 1$. ■

Proposition 3.1 *Let (P, π) and (P', π') be two model-policy pairs, the ℓ^1 -norm of the difference between the γ -discounted state distributions can be upper bounded as:*

$$\left\| d_{\mu}^{P',\pi'} - d_{\mu}^{P,\pi} \right\|_1 \leq \frac{\gamma}{1 - \gamma} \left(D_{\mathbb{E}}^{\pi',\pi} + D_{\mathbb{E}}^{P',P} \right),$$

where $D_{\mathbb{E}}^{\pi',\pi} = \mathbb{E}_{s \sim d_{\mu}^{P,\pi}} \left\| \pi'(\cdot|s) - \pi(\cdot|s) \right\|_1$ and $D_{\mathbb{E}}^{P',P} = \mathbb{E}_{(s,a) \sim \delta_{\mu}^{P,\pi}} \left\| P'(\cdot|s, a) - P(\cdot|s, a) \right\|_1$.

Proof We prove this proposition from the precedent result by decomposing the expression $\mathbb{E}_{s \sim d_{\mu}^{P,\pi}} \left\| P'^{\pi'}(\cdot|s) - P^{\pi}(\cdot|s) \right\|_1$ as follows:

$$\left\| P'^{\pi'}(\cdot|s) - P^{\pi}(\cdot|s) \right\|_1 = \int_{\mathcal{S}} \left| P'^{\pi'}(s'|s) - P^{\pi}(s'|s) \right| ds'$$

$$\begin{aligned}
 &= \int_{\mathcal{S}} \left| \int_{\mathcal{A}} \left(P'(s'|s, a) \pi'(a|s) - P(s'|s, a) \pi(a|s) \right) da \right| ds' \\
 &= \int_{\mathcal{S}} \left| \int_{\mathcal{A}} \left(P'(s'|s, a) \pi'(a|s) - P(s'|s, a) \pi(a|s) \right. \right. \\
 &\quad \left. \left. \pm P'(s'|s, a) \pi(a|s) \right) da \right| ds' \\
 &= \int_{\mathcal{S}} \left| \int_{\mathcal{A}} \left((\pi'(a|s) - \pi(a|s)) P'(s'|s, a) \right. \right. \\
 &\quad \left. \left. + \pi(a|s') (P'(s'|s, a) - P(s'|s, a)) \right) da \right| ds' \\
 &\leq \int_{\mathcal{S}} \int_{\mathcal{A}} \left| \pi'(a|s) - \pi(a|s) \right| P'(s'|s, a) da ds' \\
 &\quad + \int_{\mathcal{S}} \int_{\mathcal{A}} \pi(a|s) \left| P'(s'|s, a) - P(s'|s, a) \right| da ds'.
 \end{aligned}$$

We now take the expectation w.r.t. $d_{\mu}^{P, \pi}$ and exploit the monotonicity property:

$$\begin{aligned}
 &\mathbb{E}_{s \sim d_{\mu}^{P, \pi}} \left\| P'^{\pi'}(\cdot|s) - P^{\pi}(\cdot|s) \right\|_1 \\
 &\leq \int_{\mathcal{S}} d_{\mu}^{P, \pi}(s) \int_{\mathcal{S}} \int_{\mathcal{A}} \left| \pi'(a|s) - \pi(a|s) \right| P'(s'|s, a) da ds ds' \\
 &\quad + \int_{\mathcal{S}} d_{\mu}^{P, \pi}(s) \int_{\mathcal{S}} \int_{\mathcal{A}} \pi(a|s) \left| P'(s'|s, a) - P(s'|s, a) \right| da ds ds' \tag{9}
 \end{aligned}$$

$$\begin{aligned}
 &\leq \int_{\mathcal{S}} d_{\mu}^{P, \pi}(s) \int_{\mathcal{A}} \left| \pi'(a|s) - \pi(a|s) \right| da ds' + \\
 &\quad + \int_{\mathcal{S}} \int_{\mathcal{A}} \delta_{\mu}^{P, \pi}(s, a) \int_{\mathcal{S}} \left| P'(s'|s, a) - P(s'|s, a) \right| ds da ds' \tag{10}
 \end{aligned}$$

$$\begin{aligned}
 &= \mathbb{E}_{s \sim d_{\mu}^{P, \pi}} \left\| \pi'(\cdot|s) - \pi(\cdot|s) \right\|_1 + \mathbb{E}_{(s, a) \sim \delta_{\mu}^{P, \pi}} \left\| P'(\cdot|s, a) - P(\cdot|s, a) \right\|_1 \\
 &= D_{\mathbb{E}}^{\pi', \pi} + D_{\mathbb{E}}^{P', P},
 \end{aligned}$$

where (10) follows from (9) by observing that $d_{\mu}^{P, \pi}(s) \pi(a|s') = \delta_{\mu}^{P, \pi}(s, a)$. ■

It is worth noting that when $P = P'$ the bound resembles Corollary 3.2 in (Pirotta et al., 2013), but it is tighter as:

$$\mathbb{E}_{s \sim d_{\mu}^{P, \pi}} \left\| \pi'(\cdot|s) - \pi(\cdot|s) \right\|_1 \leq \sup_{s \in \mathcal{S}} \left\| \pi'(\cdot|s) - \pi(\cdot|s) \right\|_1,$$

in particular the bound of (Pirotta et al., 2013) might yield a large bound value in case there exist states in which the policies are very different even if those states are rarely visited according to $d_{\mu}^{P, \pi}$. In the context of policy learning, a lower bound employing the same dissimilarity index $D_{\mathbb{E}}^{\pi', \pi}$ in the penalization term has been previously proposed in (Achiam et al., 2017).

We introduce the *coupled relative advantage function*:

$$A_{P,\pi}^{P',\pi'}(s) = \int_{\mathcal{S}} \int_{\mathcal{A}} \pi'(a|s) P'(s'|s, a) \tilde{A}^{P,\pi}(s, a, s') ds' da, \quad (11)$$

where $\tilde{A}^{P,\pi}(s, a, s') = U^{P,\pi}(s, a, s') - V^{P,\pi}(s)$. $A_{P,\pi}^{P',\pi'}$ represents the one-step improvement attained by the new model-policy pair (P', π') over the current one (P, π) , i.e., the local gain in performance yielded by selecting an action with π' and the next state with P' . The corresponding expectation under the γ -discounted distribution is given by: $\mathbb{A}_{P,\pi,\mu}^{P',\pi'} = \int_{\mathcal{S}} d_{\mu}^{P,\pi}(s) A_{P,\pi}^{P',\pi'}(s) ds$. Now, we have all the elements to express the performance improvement in terms of the relative advantage functions and the γ -discounted distributions.

Theorem B.1 *The performance improvement of model-policy pair (P', π') over (P, π) is given by:*

$$J_{\mu}^{P',\pi'} - J_{\mu}^{P,\pi} = \frac{1}{1-\gamma} \int_{\mathcal{S}} d_{\mu}^{P',\pi'}(s) A_{P,\pi}^{P',\pi'}(s) ds.$$

This theorem is the natural extension of the result proposed by Kakade and Langford (2002), but, unfortunately, it cannot be directly exploited in an algorithm as the dependence of $d_{\mu}^{P',\pi'}$ on the candidate model-policy pair (P', π') is highly nonlinear and difficult to treat. We aim to obtain, from this result, a lower bound on $J_{\mu}^{P',\pi'} - J_{\mu}^{P,\pi}$ that can be efficiently computed using information on the current pair (P, π) .

Proof Let us start from the definition of $J_{\mu}^{P',\pi'}$:

$$\begin{aligned} (1-\gamma)J_{\mu}^{P',\pi'} &= \int_{\mathcal{S}} \int_{\mathcal{A}} d_{\mu}^{P',\pi'}(s) \pi'(a|s) R(s, a) da ds \\ &= \int_{\mathcal{S}} d_{\mu}^{P',\pi'}(s) \int_{\mathcal{A}} \pi'(a|s) R(s, a) da ds \pm \int_{\mathcal{S}} d_{\mu}^{P',\pi'}(s) V^{P,\pi}(s) ds \end{aligned} \quad (12)$$

$$= \int_{\mathcal{S}} d_{\mu}^{P',\pi'}(s) \int_{\mathcal{A}} \pi'(a|s) R(s, a) da ds \quad (13)$$

$$\begin{aligned} &+ \int_{\mathcal{S}} \left((1-\gamma)\mu(s') + \gamma \int_{\mathcal{S}} \int_{\mathcal{A}} d_{\mu}^{P',\pi'}(s) \pi'(a|s) P'(s'|s, a) da ds \right) V^{P,\pi}(s') ds' \\ &- \int_{\mathcal{S}} d_{\mu}^{P',\pi'}(s) V^{P,\pi}(s) ds \\ &= \int_{\mathcal{S}} d_{\mu}^{P',\pi'}(s) \left(\int_{\mathcal{A}} \pi'(a|s) \int_{\mathcal{S}} P'(s'|s, a) \left(R(s, a) + \gamma V^{P,\pi}(s') \right) ds' da \right. \\ &\quad \left. - V^{P,\pi}(s) \right) ds + (1-\gamma) \int_{\mathcal{S}} \mu(s') V^{P,\pi}(s') ds' \end{aligned} \quad (14)$$

$$= \int_{\mathcal{S}} d_{\mu}^{P',\pi'}(s) A_{P,\pi}^{P',\pi'}(s) ds + (1-\gamma)J_{\mu}^{P,\pi}, \quad (15)$$

where we have exploited the recursive formulation of $d_{\mu}^{P',\pi'}$ to rewrite (12) into (13) and (15) follows from (14) by observing that $\int_{\mathcal{S}} \mu(s') V^{P,\pi}(s') ds' = J_{\mu}^{P,\pi}$ and using the definition $U^{P,\pi}(s, a, s') = R(s, a) + \gamma V^{P,\pi}(s')$. \blacksquare

Theorem B.2 (Coupled Bound) *The performance improvement of model-policy pair (P', π') over (P, π) can be lower bounded as:*

$$\underbrace{J_\mu^{P', \pi'} - J_\mu^{P, \pi}}_{\text{performance improvement}} \geq \underbrace{\frac{\mathbb{A}_{P, \pi, \mu}^{P', \pi'}}{1 - \gamma}}_{\text{advantage}} - \underbrace{\frac{\gamma \Delta A_{P, \pi}^{P', \pi'} D_{\mathbb{E}}^{P', \pi', P^\pi}}{2(1 - \gamma)^2}}_{\text{dissimilarity penalization}},$$

where $\Delta A_{P, \pi}^{P', \pi'} = \sup_{s, s' \in \mathcal{S}} |A_{P, \pi}^{P', \pi'}(s') - A_{P, \pi}^{P', \pi'}(s)|$.

Proof

Exploiting the bounds on the γ -discounted state distributions difference (Proposition B) we can easily attain the performance improvement bound:

$$\begin{aligned} J_\mu^{P', \pi'} - J_\mu^{P, \pi} &= \frac{1}{1 - \gamma} \int_{\mathcal{S}} d_\mu^{P', \pi'}(s) A_{P, \pi}^{P', \pi'}(s) ds = \\ &= \frac{1}{1 - \gamma} \int_{\mathcal{S}} d_\mu^{P, \pi}(s) A_{P, \pi}^{P', \pi'}(s) ds + \frac{1}{1 - \gamma} \int_{\mathcal{S}} (d_\mu^{P', \pi'}(s) - d_\mu^{P, \pi}(s)) A_{P, \pi}^{P', \pi'}(s) ds \end{aligned} \quad (16)$$

$$\geq \frac{\mathbb{A}_{P, \pi, \mu}^{P', \pi'}}{1 - \gamma} - \frac{1}{1 - \gamma} \left| \int_{\mathcal{S}} (d_\mu^{P', \pi'}(s) - d_\mu^{P, \pi}(s)) A_{P, \pi}^{P', \pi'}(s) ds \right| \quad (17)$$

$$\geq \frac{\mathbb{A}_{P, \pi, \mu}^{P', \pi'}}{1 - \gamma} - \frac{1}{1 - \gamma} \|d_\mu^{P', \pi'} - d_\mu^{P, \pi}\|_1 \frac{\Delta A_{P, \pi}^{P', \pi'}}{2} \quad (18)$$

$$\geq \frac{\mathbb{A}_{P, \pi, \mu}^{P', \pi'}}{1 - \gamma} - \frac{\gamma}{(1 - \gamma)^2} \mathbb{E}_{s \sim d_\mu^{P, \pi}} \|P'^{\pi'} - P^\pi\|_1 \frac{\Delta A_{P, \pi}^{P', \pi'}}{2}, \quad (19)$$

where (17) follows from (16) by observing that $b \geq -|b|$, line (18) follows from (17) by applying Corollary 2.4 of (Haviv and Van der Heyden, 1984) and (19) is obtained by using Corollary B. \blacksquare

The *coupled bound*, however, is not suitable to be used in an algorithm as it does not separate the contribution of the policy and that of the model. In practice, an agent cannot directly update the kernel function P^π since the environment model can only partially be controlled, whereas, in many cases, we can assume a full control on the policy. For this reason, it is convenient to derive a bound in which the policy and model effects are decoupled.

Lemma 1 *The following equality relates the joint relative advantage function and the relative advantage functions:*

$$A_{P, \pi}^{P', \pi'}(s) = A_{P, \pi}^{P, \pi'}(s) + \int_{\mathcal{A}} \pi'(a|s) A_{P, \pi}^{P', \pi}(s, a) da.$$

Proof

$$A_{P, \pi}^{P', \pi'}(s) = \int_{\mathcal{A}} \int_{\mathcal{S}} \pi'(a|s) P'(s'|s, a) U^{P, \pi}(s, a, s') ds' da - V^{P, \pi}(s)$$

$$\begin{aligned}
 &= \int_{\mathcal{A}} \int_{\mathcal{S}} \pi'(a|s) P'(s'|s, a) U^{P, \pi}(s, a, s') ds' da - V^{P, \pi}(s) \\
 &\quad \pm \int_{\mathcal{A}} \int_{\mathcal{S}} \pi'(a|s) P(s'|s, a) U^{P, \pi}(s, a, s') ds' da \\
 &= \int_{\mathcal{A}} \int_{\mathcal{S}} \pi'(a|s) P(s'|s, a) U^{P, \pi}(s, a, s') ds' da - V^{P, \pi}(s) \\
 &\quad + \int_{\mathcal{A}} \int_{\mathcal{S}} \pi'(a|s) (P'(s'|s, a) - P(s'|s, a)) U^{P, \pi}(s, a, s') ds' da \\
 &= \int_{\mathcal{A}} \pi'(a|s) Q^{P, \pi}(s, a) da - V^{P, \pi}(s) \\
 &\quad + \int_{\mathcal{A}} \pi'(a|s) \int_{\mathcal{S}} (P'(s'|s, a) - P(s'|s, a)) U^{P, \pi}(s, a, s') ds' da \tag{20} \\
 &= A_{P, \pi}^{P, \pi'}(s) + \int_{\mathcal{A}} \pi'(a|s) A_{P, \pi}^{P', \pi}(s, a) da, \tag{21}
 \end{aligned}$$

where we have applied in the first addendum (20) the definition of $A_{P, \pi}^{P, \pi'}(s)$ observing that $A_{P, \pi}^{P, \pi'}(s) = \int_{\mathcal{A}} \pi'(a|s) A^{P, \pi}(s, a) da = \int_{\mathcal{A}} \pi'(a|s) (Q^{P, \pi}(s, a) - V^{P, \pi}(s))$ to get (21); and similarly for the second addendum of (20) the fact that:

$$A_{P, \pi}^{P', \pi}(s, a) = \int_{\mathcal{S}} P'(s'|s, a) A^{P, \pi}(s, a, s') ds' = \int_{\mathcal{S}} P'(s'|s, a) (U^{P, \pi}(s, a, s') - Q^{P, \pi}(s, a)) ds'.$$

■

Lemma 2 *The following bound relates the joint relative advantage function and the relative advantage functions:*

$$\left| \mathbb{A}_{P, \pi, \mu}^{P', \pi'} - \left(\mathbb{A}_{P, \pi, \mu}^{P', \pi} + \mathbb{A}_{P, \pi, \mu}^{P, \pi'} \right) \right| \leq \gamma D_{\mathbb{E}}^{\pi', \pi} D_{\infty}^{P', P} \frac{\Delta Q^{P, \pi}}{2},$$

where $D_{\infty}^{P', P} = \sup_{s \in \mathcal{S}, a \in \mathcal{A}} \|P'(\cdot|s, a) - P(\cdot|s, a)\|_1$ and $\Delta Q^{P, \pi} = \sup_{s, s' \in \mathcal{S}, a, a' \in \mathcal{A}} |Q^{P, \pi}(s', a') - Q^{P, \pi}(s, a)|$.

Proof

We can rewrite the expected relative advantage $\mathbb{A}_{P, \pi, \mu}^{P', \pi'}$ exploiting the definition:

$$\begin{aligned}
 \mathbb{A}_{P, \pi, \mu}^{P', \pi'} &= \int_{\mathcal{S}} d_{\mu}^{P, \pi} A_{P, \pi}^{P', \pi'}(s) ds = \\
 &= \int_{\mathcal{S}} d_{\mu}^{P, \pi}(s) \left(A_{P, \pi}^{P, \pi'}(s) + \int_{\mathcal{A}} \pi'(a|s) A_{P, \pi}^{P', \pi}(s, a) da \right) ds \\
 &= \int_{\mathcal{S}} d_{\mu}^{P, \pi}(s) A_{P, \pi}^{P, \pi'}(s) ds + \int_{\mathcal{S}} \int_{\mathcal{A}} d_{\mu}^{P, \pi}(s) \pi(a|s) A_{P, \pi}^{P', \pi}(s, a) dad s \\
 &\quad + \int_{\mathcal{S}} d_{\mu}^{P, \pi}(s) \int_{\mathcal{A}} (\pi'(a|s) - \pi(a|s)) A_{P, \pi}^{P', \pi}(s, a) dad s
 \end{aligned}$$

$$= \mathbb{A}_{P,\pi,\mu}^{P',\pi'} + \mathbb{A}_{P,\pi,\mu}^{P',\pi} + \int_{\mathcal{S}} d_{\mu}^{P,\pi}(s) \int_{\mathcal{A}} (\pi'(a|s) - \pi(a|s)) A_{P,\pi}^{P',\pi}(s, a) \text{d}a \text{d}s. \quad (22)$$

From equation (22) we can straightforwardly state the following inequalities:

$$\begin{aligned} \mathbb{A}_{P,\pi,\mu}^{P',\pi'} &\geq \mathbb{A}_{P,\pi,\mu}^{P,\pi'} + \mathbb{A}_{P,\pi,\mu}^{P',\pi} - \left| \int_{\mathcal{S}} d_{\mu}^{P,\pi}(s) \int_{\mathcal{A}} (\pi'(a|s) - \pi(a|s)) A_{P,\pi}^{P',\pi}(s, a) \text{d}a \text{d}s \right|, \\ \mathbb{A}_{P,\pi,\mu}^{P',\pi'} &\leq \mathbb{A}_{P,\pi,\mu}^{P,\pi'} + \mathbb{A}_{P,\pi,\mu}^{P',\pi} + \left| \int_{\mathcal{S}} d_{\mu}^{P,\pi}(s) \int_{\mathcal{A}} (\pi'(a|s) - \pi(a|s)) A_{P,\pi}^{P',\pi}(s, a) \text{d}a \text{d}s \right|, \end{aligned}$$

then we bound the right hand side:

$$\begin{aligned} \left| \mathbb{A}_{P,\pi,\mu}^{P',\pi'} - (\mathbb{A}_{P,\pi,\mu}^{P,\pi'} + \mathbb{A}_{P,\pi,\mu}^{P',\pi}) \right| &\leq \left| \int_{\mathcal{S}} d_{\mu}^{P,\pi}(s) \int_{\mathcal{A}} (\pi'(a|s) - \pi(a|s)) A_{P,\pi}^{P',\pi}(s, a) \text{d}a \text{d}s \right| \\ &\leq \int_{\mathcal{S}} \int_{\mathcal{A}} d_{\mu}^{P,\pi}(s) \left| (\pi'(a|s) - \pi(a|s)) A_{P,\pi}^{P',\pi}(s, a) \right| \text{d}a \text{d}s \\ &\leq \int_{\mathcal{S}} d_{\mu}^{P,\pi}(s) \int_{\mathcal{A}} \left| \pi'(a|s) - \pi(a|s) \right| \text{d}a \text{d}s \frac{\Delta A_{P,\pi}^{P',\pi}}{2} \\ &\leq \mathbb{E}_{s \sim d_{\mu}^{P,\pi}} \left\| \pi'(\cdot|s) - \pi(\cdot|s) \right\|_1 \frac{\Delta A_{P,\pi}^{P',\pi}}{2} \\ &\leq D_{\mathbb{E}}^{\pi',\pi} \frac{\Delta A_{P,\pi}^{P',\pi}}{2}. \end{aligned}$$

We conclude by bounding the term $\frac{\Delta A_{P,\pi}^{P',\pi}}{2}$:

$$\begin{aligned} \frac{\Delta A_{P,\pi}^{P',\pi}}{2} &\leq \left\| \Delta A_{P,\pi}^{P',\pi} \right\|_{\infty} \leq \\ &\leq \sup_{s \in \mathcal{S}, a \in \mathcal{A}} \int_{\mathcal{S}} \left(P'(s'|s, a) - P(s'|s, a) \right) U^{P,\pi}(s, a, s') \text{d}s' \quad (23) \end{aligned}$$

$$\leq \gamma \sup_{s \in \mathcal{S}, a \in \mathcal{A}} \int_{\mathcal{S}} \left(P'(s'|s, a) - P(s'|s, a) \right) V^{P,\pi}(s') \text{d}s' \quad (24)$$

$$\leq \gamma \sup_{s \in \mathcal{S}, a \in \mathcal{A}} \left\| P'(\cdot|s, a) - P(\cdot|s, a) \right\|_1 \frac{\Delta V^{P,\pi}}{2} \quad (25)$$

$$\leq \gamma D_{\infty}^{P',P} \frac{\Delta Q^{P,\pi}}{2},$$

where (24) follows from (23) by observing that $\int_{\mathcal{S}} (P'(s'|s, a) - P(s'|s, a)) U^{P,\pi}(s, a, s') \text{d}s' = \int_{\mathcal{S}} (P'(s'|s, a) - P(s'|s, a)) (R(s, a) + \gamma V^{P,\pi}(s')) \text{d}s'$ and (25) is obtained by observing that $\Delta V^{P,\pi} \leq \Delta Q^{P,\pi}$. Putting all together we get the lemma. \blacksquare

Lemma 3 *Let (P, π) and (P', π') be two model-policy pairs, it holds that:*

$$\frac{\Delta A_{P,\pi}^{P',\pi'}}{2} \leq (D_{\infty}^{\pi',\pi} + \gamma D_{\infty}^{P',P}) \frac{\Delta Q^{P,\pi}}{2},$$

where $D_{\infty}^{P',P} = \sup_{s \in \mathcal{S}, a \in \mathcal{A}} \|P'(\cdot|s, a) - P(\cdot|s, a)\|_1$, $D_{\infty}^{\pi',\pi} = \sup_{s \in \mathcal{S}} \|\pi'(\cdot|s) - \pi(\cdot|s)\|_1$ and $\Delta Q^{P,\pi} = \sup_{s, s' \in \mathcal{S}, a, a' \in \mathcal{A}} |Q^{P,\pi}(s', a') - Q^{P,\pi}(s, a)|$.

Proof Let us start rewriting the expression of the relative advantage $A_{P,\pi}^{P',\pi'}$:

$$\begin{aligned} A_{P,\pi}^{P',\pi'}(s) &= \int_{\mathcal{A}} \pi'(a|s) \left(R(s, a) + \gamma \int_{\mathcal{S}} P'(s'|s, a) V^{P,\pi}(s') ds' \right) da - V^{P,\pi}(s) \\ &= \int_{\mathcal{A}} \pi'(a|s) \left(R(s, a) + \gamma \int_{\mathcal{S}} P'(s'|s, a) V^{P,\pi}(s') ds' \right) da \\ &\quad - \int_{\mathcal{A}} \pi(a|s) \left(R(s, a) + \gamma \int_{\mathcal{S}} P(s'|s, a) V^{P,\pi}(s') ds' \right) da \\ &= \int_{\mathcal{A}} (\pi'(a|s) - \pi(a|s)) R(s, a) da \\ &\quad + \gamma \int_{\mathcal{A}} \int_{\mathcal{S}} (\pi'(a|s) P'(s'|s, a) - \pi(a|s) P(s'|s, a)) V^{P,\pi}(s') ds da \\ &= \int_{\mathcal{A}} (\pi'(a|s) - \pi(a|s)) R(s, a) da \\ &\quad + \gamma \int_{\mathcal{A}} \int_{\mathcal{S}} (\pi'(a|s) - \pi(a|s)) P(s'|s, a) V^{P,\pi}(s') ds da \\ &\quad + \gamma \int_{\mathcal{A}} \int_{\mathcal{S}} \pi'(a|s) (P'(s'|s, a) - P(s'|s, a)) V^{P,\pi}(s') ds da \\ &= \int_{\mathcal{A}} (\pi'(a|s) - \pi(a|s)) \left(R(s, a) + \gamma \int_{\mathcal{S}} P(s'|s, a) V^{P,\pi}(s') ds \right) da \\ &\quad + \gamma \int_{\mathcal{A}} \int_{\mathcal{S}} \pi'(a|s) (P'(s'|s, a) - P(s'|s, a)) V^{P,\pi}(s') ds da \\ &= \int_{\mathcal{A}} (\pi'(a|s) - \pi(a|s)) Q^{P,\pi}(s, a) da \\ &\quad + \gamma \int_{\mathcal{A}} \int_{\mathcal{S}} \pi'(a|s) (P'(s'|s, a) - P(s'|s, a)) V^{P,\pi}(s') ds da, \end{aligned}$$

then we can obtain the proof straightforwardly, noticing that $\sup_{s \in \mathcal{S}} \Delta Q^{P,\pi}(s, \cdot) \leq \Delta Q^{P,\pi}$ and recalling that $\Delta V^{P,\pi} \leq \Delta Q^{P,\pi}$:

$$\begin{aligned} \frac{\Delta A_{P,\pi}^{P',\pi'}}{2} &\leq \|A_{P,\pi}^{P',\pi'}\|_{\infty} \leq D_{\infty}^{\pi',\pi} \frac{\sup_{s \in \mathcal{S}} \Delta Q^{P,\pi}(s, \cdot)}{2} + \gamma D_{\infty}^{P',P} \frac{\Delta V^{P,\pi}}{2} \leq \\ &\leq (D_{\infty}^{\pi',\pi} + \gamma D_{\infty}^{P',P}) \frac{\Delta Q^{P,\pi}}{2}. \end{aligned}$$

■

Theorem 3.1 (Decoupled Bound) *The performance improvement of model-policy pair (P', π') over (P, π) can be lower bounded as:*

$$\underbrace{J_\mu^{P', \pi'} - J_\mu^{P, \pi}}_{\text{performance improvement}} \geq B(P', \pi') = \underbrace{\frac{\mathbb{A}_{P, \pi, \mu}^{P', \pi'} + \mathbb{A}_{P, \pi, \mu}^{P, \pi'}}{1 - \gamma}}_{\text{advantage}} - \underbrace{\frac{\gamma \Delta Q^{P, \pi} D}{2(1 - \gamma)^2}}_{\text{dissimilarity penalization}},$$

where D is a dissimilarity term defined as $D = D_{\mathbb{E}}^{\pi', \pi} (D_{\infty}^{\pi', \pi} + D_{\infty}^{P', P}) + D_{\mathbb{E}}^{P', P} (D_{\infty}^{\pi', \pi} + \gamma D_{\infty}^{P', P})$ and $D_{\infty}^{\pi', \pi} = \sup_{s \in \mathcal{S}} \|\pi'(\cdot|s) - \pi(\cdot|s)\|_1$, $D_{\infty}^{P', P} = \sup_{s \in \mathcal{S}, a \in \mathcal{A}} \|P'(\cdot|s, a) - P(\cdot|s, a)\|_1$ and $\Delta Q^{P, \pi} = \sup_{s, s' \in \mathcal{S}, a, a' \in \mathcal{A}} |Q^{P, \pi}(s', a') - Q^{P, \pi}(s, a)|$.

Proof Exploiting the bounds on the coupled expected relative advantage $\mathbb{A}_{P, \pi, \mu}^{P', \pi'}$ (Lemma 2), on the γ -discounted state distributions difference (Corollary 3) and on $\Delta A_{P, \pi}^{P', \pi'}$ (Lemma 3), we can state the following:

$$\begin{aligned} J_\mu^{P', \pi'} - J_\mu^{P, \pi} &\geq \frac{\mathbb{A}_{P, \pi, \mu}^{P', \pi'}}{1 - \gamma} - \frac{\|d_\mu^{P', \pi'} - d_\mu^{P, \pi}\|_1 \Delta A_{P, \pi}^{P', \pi'}}{1 - \gamma} \frac{1}{2} \\ &\geq \frac{\mathbb{A}_{P, \pi, \mu}^{P', \pi'} + \mathbb{A}_{P, \pi, \mu}^{P, \pi}}{1 - \gamma} - \frac{\gamma}{(1 - \gamma)^2} (D_{\mathbb{E}}^{\pi', \pi} + D_{\mathbb{E}}^{P', P}) (D_{\infty}^{\pi', \pi} + \gamma D_{\infty}^{P', P}) \frac{\Delta Q^{P, \pi}}{2} \\ &\quad - \frac{\gamma}{1 - \gamma} D_{\mathbb{E}}^{\pi', \pi} D_{\infty}^{P', P} \frac{\Delta Q^{P, \pi}}{2} \\ &\geq \frac{\mathbb{A}_{P, \pi, \mu}^{P', \pi'} + \mathbb{A}_{P, \pi, \mu}^{P, \pi}}{1 - \gamma} - \frac{\gamma}{(1 - \gamma)^2} (D_{\mathbb{E}}^{\pi', \pi} D_{\infty}^{\pi', \pi} + D_{\mathbb{E}}^{\pi', \pi} D_{\infty}^{P', P} \\ &\quad + D_{\infty}^{\pi', \pi} D_{\mathbb{E}}^{P', P} + \gamma D_{\infty}^{P', P} D_{\mathbb{E}}^{P', P}) \frac{\Delta Q^{P, \pi}}{2}. \end{aligned}$$

With a factorization of the last expression we get the result. ■

Theorem 4.1 *For any $\bar{\pi} \in \Pi$ and $\bar{P} \in \mathcal{P}$, the decoupled bound is optimized for $\alpha^*, \beta^* = \arg \max_{\alpha, \beta} \{B(\alpha, \beta) : (\alpha, \beta) \in \mathcal{V}\}$, where $\mathcal{V} = \{(\alpha_0^*, 0), (\alpha_1^*, 1), (0, \beta_0^*), (1, \beta_1^*)\}$ and the values of α_0^* , α_1^* , β_0^* and β_1^* are reported in Table 1.*

Proof Let us write explicitly the update coefficients in the decoupled bound (3.1):

$$\begin{aligned} J_\mu^{P', \pi'} - J_\mu^{P, \pi} &\geq \frac{\alpha \mathbb{A}_{P, \pi, \mu}^{P, \bar{\pi}} + \beta \mathbb{A}_{P, \pi, \mu}^{\bar{P}, \pi}}{1 - \gamma} - \frac{\gamma}{(1 - \gamma)^2} \\ &\quad \times (\alpha^2 D_{\mathbb{E}}^{\bar{\pi}, \pi} D_{\infty}^{\bar{\pi}, \pi} + \alpha \beta D_{\mathbb{E}}^{\bar{\pi}, \pi} D_{\infty}^{\bar{P}, P} + \alpha \beta D_{\infty}^{\bar{\pi}, \pi} D_{\mathbb{E}}^{\bar{P}, P} + \gamma \beta^2 D_{\infty}^{\bar{P}, P} D_{\mathbb{E}}^{\bar{P}, P}) \frac{\Delta Q^{P, \pi}}{2}, \end{aligned}$$

we now take the derivatives w.r.t. α and β to find the stationary points:

$$\begin{aligned}\frac{\partial B}{\partial \alpha} &= \frac{\mathbb{A}_{P,\pi,\mu}^{P,\bar{\pi}}}{1-\gamma} - \frac{\gamma}{(1-\gamma)^2} (2\alpha D_{\mathbb{E}}^{\bar{\pi},\pi} D_{\infty}^{\bar{\pi},\pi} + \beta D_{\mathbb{E}}^{\bar{\pi},\pi} D_{\infty}^{\bar{P},P} + \beta D_{\infty}^{\bar{\pi},\pi} D_{\mathbb{E}}^{\bar{P},P}) \frac{\Delta Q^{P,\pi}}{2}, \\ \frac{\partial B}{\partial \beta} &= \frac{\mathbb{A}_{P,\pi,\mu}^{\bar{P},\pi}}{1-\gamma} - \frac{\gamma}{(1-\gamma)^2} (\alpha D_{\mathbb{E}}^{\bar{\pi},\pi} D_{\infty}^{\bar{P},P} + \alpha D_{\infty}^{\bar{\pi},\pi} D_{\mathbb{E}}^{\bar{P},P} + 2\gamma \beta D_{\infty}^{\bar{P},P} D_{\mathbb{E}}^{\bar{P},P}) \frac{\Delta Q^{P,\pi}}{2}.\end{aligned}$$

When the target policy is different from the current one and, symmetrically, the target model is different from the current model the linear system of the derivatives admits a unique solution. We compute the second order derivative to discover the nature of such point:

$$\begin{aligned}\frac{\partial B^2}{\partial^2 \alpha} &= -\frac{\gamma}{(1-\gamma)^2} D_{\mathbb{E}}^{\bar{\pi},\pi} D_{\infty}^{\bar{\pi},\pi} \Delta Q^{P,\pi}, \\ \frac{\partial B^2}{\partial \alpha \partial \beta} &= \frac{\partial B^2}{\partial \beta \partial \alpha} = -\frac{\gamma}{(1-\gamma)^2} (D_{\mathbb{E}}^{\bar{\pi},\pi} D_{\infty}^{\bar{P},P} + D_{\infty}^{\bar{\pi},\pi} D_{\mathbb{E}}^{\bar{P},P}) \frac{\Delta Q^{P,\pi}}{2}, \\ \frac{\partial B^2}{\partial^2 \beta} &= -\frac{\gamma^2}{(1-\gamma)^2} D_{\infty}^{\bar{P},P} D_{\mathbb{E}}^{\bar{P},P} \Delta Q^{P,\pi},\end{aligned}$$

from the second order derivatives we can write the Hessian matrix:

$$\mathcal{H}(\alpha, \beta) = -\frac{\gamma \Delta Q^{P,\pi}}{(1-\gamma)^2} \begin{pmatrix} 2D_{\mathbb{E}}^{\bar{\pi},\pi} D_{\infty}^{\bar{\pi},\pi} & D_{\mathbb{E}}^{\bar{\pi},\pi} D_{\infty}^{\bar{P},P} + D_{\infty}^{\bar{\pi},\pi} D_{\mathbb{E}}^{\bar{P},P} \\ D_{\mathbb{E}}^{\bar{\pi},\pi} D_{\infty}^{\bar{P},P} + D_{\infty}^{\bar{\pi},\pi} D_{\mathbb{E}}^{\bar{P},P} & 2\gamma D_{\infty}^{\bar{P},P} D_{\mathbb{E}}^{\bar{P},P} \end{pmatrix},$$

having trace and determinant:

$$\begin{aligned}\text{tr}(\mathcal{H}(\alpha, \beta)) &= -\frac{\gamma \Delta Q^{P,\pi}}{(1-\gamma)^2} (2D_{\mathbb{E}}^{\bar{\pi},\pi} D_{\infty}^{\bar{\pi},\pi} + 2\gamma D_{\infty}^{\bar{P},P} D_{\mathbb{E}}^{\bar{P},P}) \leq 0, \\ \det(\mathcal{H}(\alpha, \beta)) &= \frac{\gamma^2 \Delta Q_{P,\pi}^{P',\pi^2}}{(1-\gamma)^4} \left((4\gamma - 2) D_{\mathbb{E}}^{\bar{\pi},\pi} D_{\infty}^{\bar{\pi},\pi} D_{\infty}^{\bar{P},P} D_{\mathbb{E}}^{\bar{P},P} - (D_{\mathbb{E}}^{\bar{\pi},\pi} D_{\infty}^{\bar{P},P})^2 - (D_{\infty}^{\bar{\pi},\pi} D_{\mathbb{E}}^{\bar{P},P})^2 \right) \\ &\leq \frac{\gamma^2 \Delta Q_{P,\pi}^{P',\pi^2}}{(1-\gamma)^4} \left(2D_{\mathbb{E}}^{\bar{\pi},\pi} D_{\infty}^{\bar{\pi},\pi} D_{\infty}^{\bar{P},P} D_{\mathbb{E}}^{\bar{P},P} - (D_{\mathbb{E}}^{\bar{\pi},\pi} D_{\infty}^{\bar{P},P})^2 - (D_{\infty}^{\bar{\pi},\pi} D_{\mathbb{E}}^{\bar{P},P})^2 \right) \\ &\leq -\frac{\gamma^2 \Delta Q_{P,\pi}^{P',\pi^2}}{(1-\gamma)^4} \left(D_{\mathbb{E}}^{\bar{\pi},\pi} D_{\infty}^{\bar{P},P} - D_{\infty}^{\bar{\pi},\pi} D_{\mathbb{E}}^{\bar{P},P} \right)^2 \leq 0.\end{aligned}$$

When $\bar{P} \neq P$ and $\bar{\pi} \neq \pi$ we observe that the Hessian matrix is indefinite since both the trace and the determinant are negative. This means that the unique stationary point is a saddle point which is uninteresting for optimization purposes. By the way, $B(\alpha, \beta)$ is a quadratic function, therefore it is continuous on the compact set $[0, 1]^2$ and therefore, from Weierstrass theorem, it admits a global maximum (and minimum). Since such point is not a stationary point it must lie on the boundary of $[0, 1]^2$.

Then, by setting to zero the equations $\frac{\partial B}{\partial \alpha}|_{\beta=0}$, $\frac{\partial B}{\partial \alpha}|_{\beta=1}$, $\frac{\partial B}{\partial \beta}|_{\alpha=0}$, $\frac{\partial B}{\partial \beta}|_{\alpha=1}$ we can obtain the following optimal values (which are clipped to lie in the interval $[0, 1]$):

$$\alpha_0^* = \frac{(1-\gamma)}{\gamma \Delta Q^{P,\pi}} \frac{\mathbb{A}_{P,\pi,\mu}^{P,\bar{\pi}}}{D_{\infty}^{\bar{\pi},\pi} D_{\mathbb{E}}^{\bar{\pi},\pi}},$$

$$\begin{aligned}\alpha_1^* &= \frac{(1-\gamma)}{\gamma\Delta Q^{P,\pi}} \frac{\mathbb{A}_{P,\pi,\mu}^{P,\bar{\pi}}}{D_\infty^{\bar{\pi},\pi} D_\mathbb{E}^{\bar{\pi},\pi}} - \frac{1}{2} \left(\frac{D_\mathbb{E}^{\bar{P},P}}{D_\infty^{\bar{\pi},\pi}} + \frac{D_\infty^{\bar{P},P}}{D_\infty^{\bar{\pi},\pi}} \right), \\ \beta_0^* &= \frac{(1-\gamma)}{\gamma^2\Delta Q^{P,\pi}} \frac{\mathbb{A}_{P,\pi,\mu}^{\bar{P},\pi}}{D_\infty^{\bar{P},P} D_\mathbb{E}^{\bar{P},P}}, \\ \beta_1^* &= \frac{(1-\gamma)}{\gamma^2\Delta Q^{P,\pi}} \frac{\mathbb{A}_{P,\pi,\mu}^{\bar{P},\pi}}{D_\infty^{\bar{P},P} D_\mathbb{E}^{\bar{P},P}} - \frac{1}{2\gamma} \left(\frac{D_\mathbb{E}^{\bar{\pi},\pi}}{D_\infty^{\bar{P},P}} + \frac{D_\infty^{\bar{\pi},\pi}}{D_\infty^{\bar{P},P}} \right).\end{aligned}$$

Instead, for $\gamma \in (0, 1)$, the Hessian is singular when either the target policy or the target model are equal to the current one. Those cases can be treated separately and clearly yield maxima points. When $\bar{P} = P$ then we have $\alpha^* = \alpha_0^*$, when $\bar{\pi} = \pi$ we have $\beta^* = \beta_0^*$.

We report the values of the decoupled bound $B(\alpha, \beta)$ (Theorem 3) in correspondence of the optimal coefficients:

$$\begin{aligned}B(\alpha_0^*, 0) &= \frac{(1-\gamma)\mathbb{A}_{P,\pi,\mu}^{P,\bar{\pi}}{}^2}{2\gamma\Delta Q^{P,\pi} D_\infty^{\bar{\pi},\pi} D_\mathbb{E}^{\bar{\pi},\pi}}, \\ B(0, \beta_0^*) &= \frac{(1-\gamma)\mathbb{A}_{P,\pi,\mu}^{\bar{P},\pi}{}^2}{2\gamma^2\Delta Q^{P,\pi} D_\infty^{\bar{P},P} D_\mathbb{E}^{\bar{P},P}}, \\ B(1, \beta_1^*) &= \mathbb{A}_{P,\pi,\mu}^{P,\bar{\pi}} + \frac{(1-\gamma)\mathbb{A}_{P,\pi,\mu}^{\bar{P},\pi}{}^2}{2\gamma^2\Delta Q^{P,\pi} D_\infty^{\bar{P},P} D_\mathbb{E}^{\bar{P},P}} - \frac{\mathbb{A}_{P,\pi,\mu}^{\bar{P},\pi}}{2\gamma} \left(\frac{D_\mathbb{E}^{\bar{\pi},\pi}}{D_\infty^{\bar{P},P}} + \frac{D_\infty^{\bar{\pi},\pi}}{D_\infty^{\bar{P},P}} \right) \\ &\quad + \frac{\Delta Q^{P,\pi}}{2(1-\gamma)} \left(\frac{1}{2} D_\mathbb{E}^{\bar{P},P} D_\mathbb{E}^{\bar{\pi},\pi} D_1 + \frac{1}{2} D_\infty^{\bar{P},P} D_\mathbb{E}^{\bar{\pi},\pi} D_1 - \frac{1}{4} D_\infty^{\bar{P},P} D_\mathbb{E}^{\bar{P},P} D_1^2 - \gamma D_\mathbb{E}^{\bar{\pi},\pi} D_\infty^{\bar{\pi},\pi} \right), \\ B(\alpha_1^*, 1) &= \mathbb{A}_{P,\pi,\mu}^{\bar{P},\pi} + \frac{(1-\gamma)\mathbb{A}_{P,\pi,\mu}^{P,\bar{\pi}}{}^2}{2\gamma\Delta Q^{P,\pi} D_\infty^{\bar{\pi},\pi} D_\mathbb{E}^{\bar{\pi},\pi}} - \frac{\mathbb{A}_{P,\pi,\mu}^{P,\bar{\pi}}}{2} \left(\frac{D_\mathbb{E}^{\bar{P},P}}{D_\infty^{\bar{\pi},\pi}} + \frac{D_\infty^{\bar{P},P}}{D_\infty^{\bar{\pi},\pi}} \right) \\ &\quad + \frac{\gamma\Delta Q^{P,\pi}}{2(1-\gamma)} \left(\frac{1}{2} D_\mathbb{E}^{\bar{\pi},\pi} D_\mathbb{E}^{\bar{P},P} D_2 + \frac{1}{2} D_\infty^{\bar{\pi},\pi} D_\mathbb{E}^{\bar{P},P} D_2 - \frac{1}{4} D_\infty^{\bar{\pi},\pi} D_\mathbb{E}^{\bar{\pi},\pi} D_2^2 - \gamma D_\mathbb{E}^{\bar{P},P} D_\infty^{\bar{P},P} \right),\end{aligned}$$

where

$$D_1 = \frac{D_\mathbb{E}^{\bar{\pi},\pi}}{D_\infty^{\bar{P},P}} + \frac{D_\infty^{\bar{\pi},\pi}}{D_\infty^{\bar{P},P}}, \quad D_2 = \frac{D_\mathbb{E}^{\bar{P},P}}{D_\infty^{\bar{\pi},\pi}} + \frac{D_\infty^{\bar{P},P}}{D_\infty^{\bar{\pi},\pi}}.$$

■




Figure 3: An example of some additional plots. Conf-MDP with local maxima.

Appendix C. Examples of Conf-MDPs

In this section we report two examples of Conf-MDPs having some interesting behaviors. In all figures, the transition probabilities are reported on the edges and the reward is written below the state name.

C.1 An example of Conf-MDP with local optima

Let us consider the Conf-MDP represented in Figure 3 where $\omega \in [0, 1]$ is the parameter, $p \in [0, 1]$ is a small fixed probability (say 0.1) and M is a large positive number. In each state there is only one action available (i.e., all policies are optimal). The vertex models are obtained for $\omega \in \{0, 1\}$. For both target models there is a small probability to get the punishment $-M$ since for $\omega = 0$ the probability to reach state C from B is p and for $\omega = 1$ state B is reachable from A with probability p . We expect that by mixing the two target models we can only worsen the performance. It is simple to realize that the expected return is a cubic function of ω . We report the expression for $p = 0.1$ and $\gamma = 1$:

$$J_{\mu}^{P_{\omega}} = \frac{1}{2} \left(0.512\omega^3 + (0.64M - 1.088)\omega^2 - (0.64M + 0.296)\omega + 1.981 - 0.09M \right).$$

We can find the stationary points by looking at the derivative:

$$\frac{\partial J_{\mu}^{P_{\omega}}}{\partial \omega} = 0.768\omega^2 + (0.64M - 1.088)\omega - 0.32M - 0.148.$$

For M sufficiently large the derivative has one sign variation thus it has two solutions of opposite sign, having expression:

$$\omega_{1,2} = \frac{1}{24} \left(17 - 10M \pm 10\sqrt{M^2 - M - 4} \right).$$

Clearly, we are interested only in the solutions within $[0, 1]$ thus we discard the negative one. It is simple to see that the positive solution is approximately $\frac{1}{2}$ for M sufficiently large, as:

$$\lim_{M \rightarrow +\infty} \frac{1}{24} \left(17 - 10M + 10\sqrt{M^2 - M - 4} \right) = \frac{1}{2}.$$

However, having a look at the second derivative we realize that this is a point of minimum, since

$$\frac{\partial^2 J_{\mu}^{P_{\omega}}}{\partial \omega^2} = 1.536\omega + 0.64M - 1.088 \Big|_{\omega=\frac{1}{2}} > 0.$$




Figure 4: An example of Conf-MDP with mixed optimal model.

Notice that in the unfortunate case in which SPMI is initialized at this value of ω the expected relative advantage (which is the same as the gradient) is zero for both the vertex models and therefore there would be no update. Therefore, the maximum must lie on the border, specifically either for $\omega = 0$ or $\omega = 1$. It is simple to see that $J_\mu^{P_1} > J_\mu^{P_0}$. Moreover, if we compute the value of the gradient for $\omega = 0$ and $\omega = 1$ we realize that in both cases the value is negative. Having a negative advantage, SPMI would never make any step even when the model is initialized at the lower performance vertex $\omega = 0$.

C.2 An example of Conf-MDP with mixed optimal model

We consider the Conf-MDP as represented in Figure 4. As in the previous case, the parameter is $\omega \in [0, 1]$ and $p \in [0, 1]$ is a fixed probability. We want to show that there exists no value of ω such that P_ω maximizes the value function in all states, while there exist one value of ω maximizing the expected return. It is simple to compute the value function in each state:

$$\begin{aligned} V^{P_\omega}(A) &= \gamma^2(\omega p + (1 - \omega)(1 - p))(\omega(1 - p) + (1 - \omega)p), \\ V^{P_\omega}(B) &= \gamma(\omega(1 - p) + (1 - \omega)p), \\ V^{P_\omega}(C) &= 1, \\ V^{P_\omega}(D) &= 0. \end{aligned}$$

Since the initial state is A we have that $J_\mu^{P_\omega} = V^{P_\omega}(A)$ which is maximized for $\omega = 1/2$. However, there is no value of ω for which the value function of each state is maximized. As shown in Figure 5, while $V^{P_\omega}(A)$ is maximal in $\omega = 1/2$, $V^{P_\omega}(B)$ is maximal for $\omega = 1$. All values of $\omega \in [1/2, 1]$ are indeed Pareto optimal (Figure 6).

With some boring calculation we can determine the expression of the expected relative advantage functions:

$$\begin{aligned} \mathbb{A}_{P_{\omega,\mu}}^{P_1} &= \gamma^2(1 - \omega)(1 - 2\omega)(1 - 2p)^2 \\ \mathbb{A}_{P_{\omega,\mu}}^{P_2} &= -\gamma^2\omega(1 - 2\omega)(1 - 2p)^2. \end{aligned}$$

We clearly see that they both vanish for $\omega = 1/2$.

Appendix D. Environment description

In this appendix we provide a more detailed description of the environments used in the experimental evaluation.




Figure 5: The state value function of the Conf-MDP in Figure 4 as a function of the parameter.




Figure 6: The state value function of states A and B (the only ones varying with the parameter) of the Conf-MDP in Figure 4. The green continuous line is the Pareto frontier.

D.1 Student-Teacher domain

The teaching/learning process involves two entities: the teacher and the student (learner) that interact. We assume both entities share the same goal, i.e., maximizing the learning. The teaching model, however, should be suited for the specific learning policy of the student. For instance, not all students have the same skills and are able to capture the information provided by the teacher with the same speed and effectiveness. Thus, the teaching model should be tailored in order to meet the student’s needs. Given the goal of maximizing learning, a teaching model induces an optimal learning policy (within the space of the policies that a certain student can play). Symmetrically, a learning policy determines an optimal teaching model (within the space of models available to the teacher). The question we want to answer in this experiment is: can we dynamically adapt the teaching model to the learning policy and the learning policy to the teaching model, so to maximize the learning?

We take inspiration from (Rafferty et al., 2011) and we formalize the teaching/learning process as an MDP in which the student is the agent and the teacher is the environment. To fit our framework to this context, we can think to the teacher as an online learning platform that can be configured by the student in order to improve the learning experience. As in (Rafferty et al., 2011) we test the model on the “alphabet arithmetic” a concept-learning task in which literals are mapped to numbers.

We consider n literals L_1, \dots, L_n , to which the student can assign the values $\{0, \dots, m\}$. The teacher, at each time step, provides an “example”, i.e., an equation where a number of distinct literals (from 2 to $p \leq n$) sum to a numerical answer. The set of all possible examples is given by:

$$\mathcal{E} = \left\{ \sum_{i \in I} L_i = l : I \subseteq \{1, \dots, n\}, 2 \leq |I| \leq p, l \in \{0, \dots, |I|m\} \right\}.$$

The student reacts to an example by performing an action, i.e., an assignment of literals. The set of all assignments, i.e., actions, is given by:

$$\mathcal{A} = \left\{ L_1 = l_1, L_2 = l_2, \dots, L_n = l_n : l_i \in \{0, \dots, m\}, i = 1, \dots, n \right\},$$

thus $|\mathcal{A}| = (m + 1)^n$. In order to model the student policy space we assume that a student can modify an arbitrary number of literals under the assumption that two consecutive assignments satisfy $\sum_{i=1}^n |l'_i - l_i| \leq k$. This models the learning limitations of the student, in particular how hard is for the student to capture the teacher information. We assume that the teacher can provide any example. The set of states is the combination of an assignment and an example, i.e., $\mathcal{S} = \mathcal{E} \times \mathcal{A}$.

The goal of the student is to perform assignments that are consistent with the teacher's examples (within its limitations on the possible assignments). So, while the student is learning the optimal policy it can configure the teacher to provide more suitable examples. The reward is 1 when the assignment is consistent, 0, when it is not and the horizon H is finite. Notice that we don't have goal state, differently from (Rafferty et al., 2011) so the problem becomes fully observable. We assume that, at the beginning, both policy and model are uniform distribution on the allowed actions/states.

D.2 Racetrack Simulator

The Racetrack Simulator aims to represent a basic abstraction of an autonomous car driving problem. In this context the autonomous driver, the learning agent, has to optimize a driving policy in order to run the vehicle to the track finish line as fast as possible. The vehicle and the track naturally compose the model of the learning process, however there is the possibility to tune a set of vehicle parameters, s.t. aerodynamic profile (to affect the vehicle stability) and engine setting. Therefore, to maximize the performance, the driving policy of the agent and the model configuration has to be jointly considered. It is noteworthy that a specific model parametrization (vehicle setting) induces an optimal driving policy and, in the other hand, a driving policy determines an optimal model parametrization. Moreover a model-policy pair that results to be optimal for a specific track may not be optimal for a (morphologically) different track. Then, the question we aim to answer with this experiment is the following: can we learn the optimal model-policy pair for a given track by dynamically adapt the vehicle parametrization to the driving policy and, conversely, the driving policy to the vehicle parametrization during the learning process?

We formalize the learning process as an MDP in which the driver is the agent and the environment is composed by the track and the vehicle. The track is represented by a grid of positions, each grid point is either of type *roadway*, *wall*, *initial position*, *goal position*. A state in the learning process belongs to the set:

$$\mathcal{S} = \left\{ (x, y, v_x, v_y) : x \in \{0, \dots, x_{\max}\}, y \in \{0, \dots, y_{\max}\}, v_x \in \{v_{\min}, \dots, v_{\max}\}, v_y \in \{v_{\min}, \dots, v_{\max}\} \right\},$$

where (x, y) corresponds to a grid position and (v_x, v_y) are the speed along the coordinate axes. At each step the agent can increment the speed along a coordinate direction or do nothing, then the action space is represented by the following:

$$\mathcal{A} = \left\{ \textit{keep}, \textit{increment } v_x, \textit{increment } v_y, \textit{decrement } v_x, \textit{decrement } v_y \right\}.$$

The learning process starts at the state corresponding to the initial position with zero velocities; the agent collects reward 1 when it reaches a state corresponding to the goal position within the finite horizon H , he collects 0 reward in any other case.

The transition model induces a success probability to any action, a failed action causes a random action to occur instead of the one selected by the agent. This probability aims to model the stability of the vehicle, the more the vehicle is unstable, the more is hard for the agent to drive it (or select an action). The model also induces a failure probability to every action: a failure represent a break of the vehicle, thus it directly cause the end of the episode. This feature represents the pressure on the vehicle engine, the more performance the driver asks for, the more it may break down. We formalize the transition model as a convex combination between a set of vertex models: these correspond to vehicle configuration pushed towards the limit in terms of the aspects described above. For our purpose we define a model dichotomy related to vehicle stability: $P_highspeed$ (P_hs) trades stability at lower speed to have more stability (or high action success probability) in high speed situations, $P_lowspeed$ (P_ls), instead, provides more stability in low speed situation and poor stability at higher speed. We define also a model dichotomy related to engine boost: P_boost (P_b) guarantees higher engine performance and a lower reliability (or higher failure probability), at the opposite $P_noboost$ (P_nb) provides higher reliability but poor engine performance. In Figure 7 we propose a graphical representation of the features of these extreme models.




Figure 7: Graphical representation of the racetrack extreme models.

Considering any possible combination of stability and engine setting, we define the model set (set of vertex models) $\mathbf{P} = \{P_hs_b, P_hs_nb, P_ls_b, P_ls_nb\}$. Each model in this set is obtained by taking, for each state-action pair, the product of the transition probabilities of the components (e.g., $P_hs_b(s, a) = P_hs(s, a) \times P_b(s, a)$). Then, we derive the model space as the convex hull of the vertices in the model set:

$$\mathcal{P}_\omega = \left\{ P_\omega = \omega_0 P_hs_b + \omega_1 P_hs_nb + \omega_2 P_ls_b + \omega_3 P_ls_nb \right\}.$$

While the agent is learning the optimal driving policy, the model parametrization can be configured (selecting a vector $(\omega_0, \omega_1, \omega_2, \omega_3)$) trying to fit the vehicle settings to the driving policy and simultaneously trying to fit the policy-settings pair to the morphology of the track. At the beginning of the learning process, we assume the policy to be a uniform distribution on the action space and the model to be $(0, 0.5, 0, 0.5)$, that we can consider the most conservative parametrization in our context. We also report in Figure 8 an illustrative representation of the tracks used in the experiments.

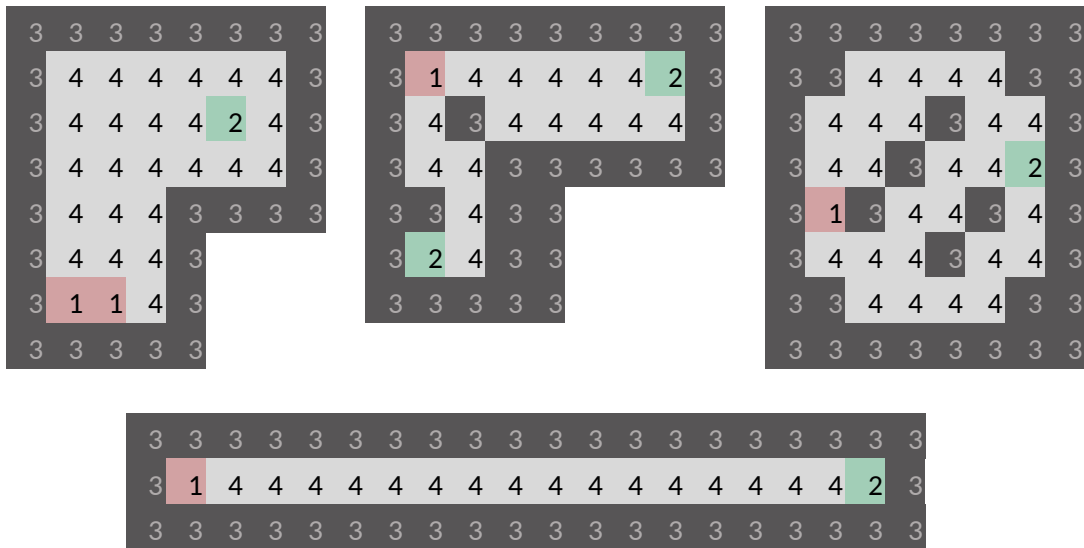


Figure 8: Basic representation of the tracks used in the Racetrack Simulator. From left to right: T1, T3, T4 and T2 just below. Each position have a type label: 1 for initial states, 2 for goal states, 3 for walls, 4 for roadtracks.

Appendix E. Experiments details

In this appendix, we report the hyper-parameters used for each experiment and some additional plots.

E.1 Student-Teacher domain

In Table 2 we report the number of iterations to convergence for the different problem settings we considered. We can see that SPMI is the first or the second to converge in most of the cases. In Table 3 we report the hyper-parameters we used for the runs in the different problem settings. In Figures 9, 10, 11 we provide numerous plots showing different interesting metrics for the problem settings we considered in the experimental evaluation.

Table 2: Number of steps to convergence for the update strategies in different problem settings of the Student-Teacher domain. In **bold** the best algorithm and underlined the second best. The run has been stopped after 50000 iterations.

Problem	SPMI	SPMI-sup	SPMI-alt	SPI+SMI	SPI+SMI
2-1-1-2	<u>16234</u>	18054	30923	22130	7705
2-1-2-2	2839	3194	5678	2839	12973
2-2-1-2	20345	<u>18287</u>	>50000	39722	10904
2-2-2-2	12025	<u>14315</u>	>50000	>50000	15257
2-3-1-2	14187	13391	11772	>50000	<u>12183</u>
3-1-1-2	<u>15410</u>	17929	22707	31122	14257
3-1-2-2	3313	3313	8434	3313	22846
3-1-3-2	2945	3435	5891	2945	18090

Table 3: Additional information and hyper-parameters used for the different problem settings of the Student-Teacher domain.

Problem	$ \mathcal{S} $	$ \mathcal{A} $	horizon	$\Delta Q^{P,\pi}$	γ	ϵ
2-1-1-2	12	4	10	$\frac{1-\gamma^{10}}{1-\gamma} \simeq 9.56$	0.99	0
2-1-2-2	12	4	10	9.56	0.99	0
2-2-1-2	45	8	10	9.56	0.99	0
2-2-2-2	45	8	10	9.56	0.99	0
2-3-1-2	112	16	10	9.56	0.99	0
3-1-1-2	72	9	10	9.56	0.99	0
3-1-2-2	72	9	10	9.56	0.99	0
3-1-3-2	72	9	10	9.56	0.99	0

CONFIGURABLE MARKOV DECISION PROCESSES

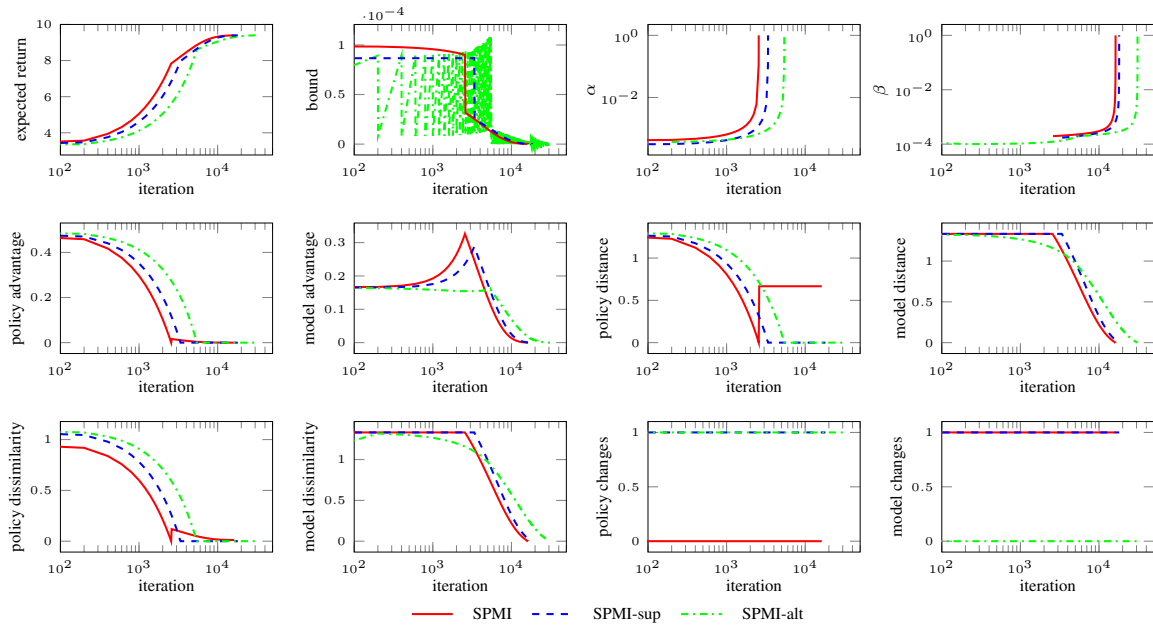


Figure 9: Several statistics of the persistent target choice for different update strategies in the case of Student-Teacher domain 2-1-1-2.

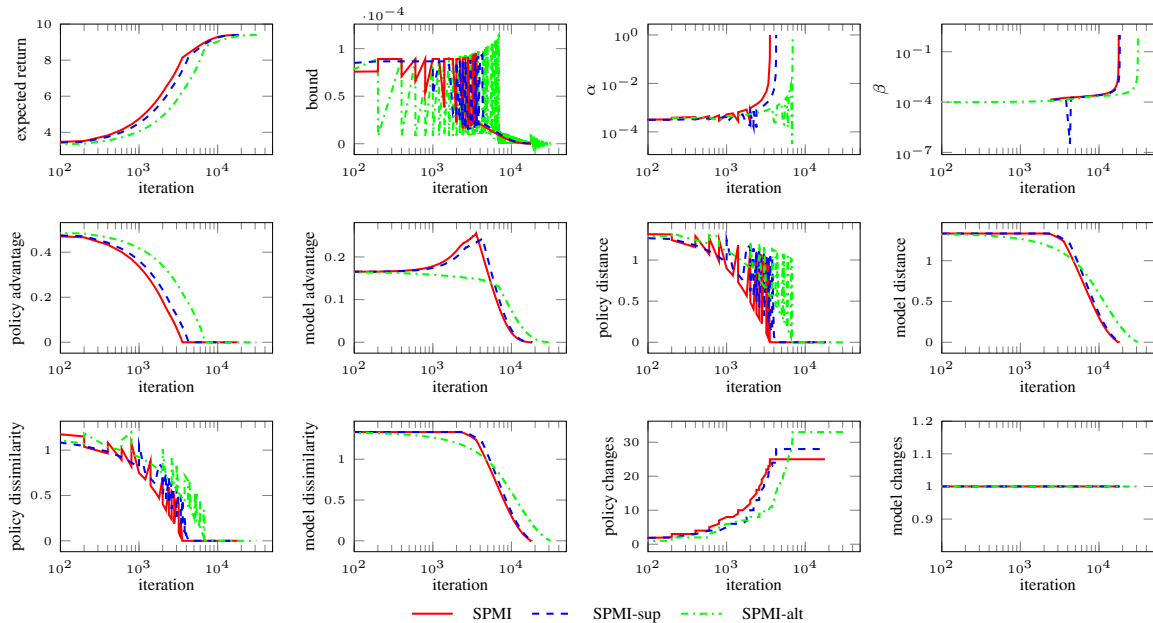


Figure 10: Several statistics of the greedy target choice for different update strategies in the case of Student-Teacher domain 2-1-1-2.

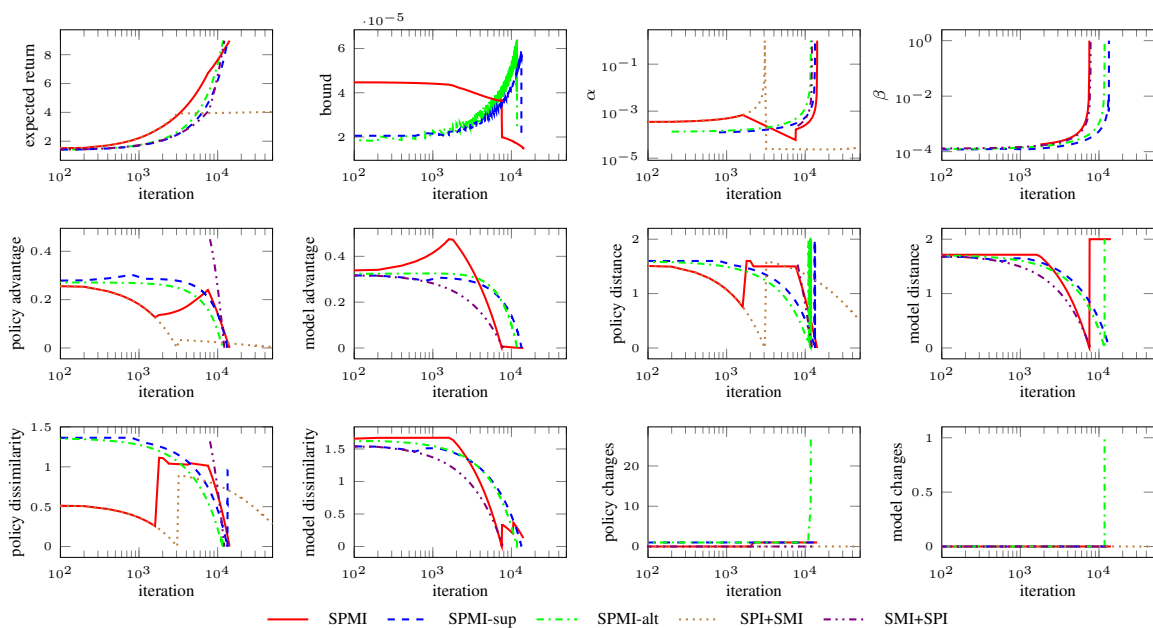


Figure 11: Several statistics of the persistent target choice for different update strategies in the case of Student-Teacher domain 2-3-1-2.

E.2 Racetrack Simulator

In this section, we provide additional information about the experiments in the Racetrack Simulator environment (Table 4) along with a more detailed presentation of the results reported in the Experimental Evaluation section of the paper (Figure 12 and Figure 13). Moreover, we present the performance obtained in some cases not mentioned in the paper (Figure 14).

Table 4: Additional information and hyper-parameters used for the different problem settings of the Racetrack Simulator. In the cases with two vertices only stability vehicle configurations are considered.

Track	vertices	$ \mathcal{S} $	$ \mathcal{A} $	horizon	$\Delta Q^{P,\pi}$	γ	ϵ
T1	2	675	5	20	1	0.9	0
T3	2	450	5	20	1	0.9	0
T4	2	675	5	20	1	0.9	0
T2	2	450	5	20	1	0.9	0
T2	4	451	5	20	1	0.9	0

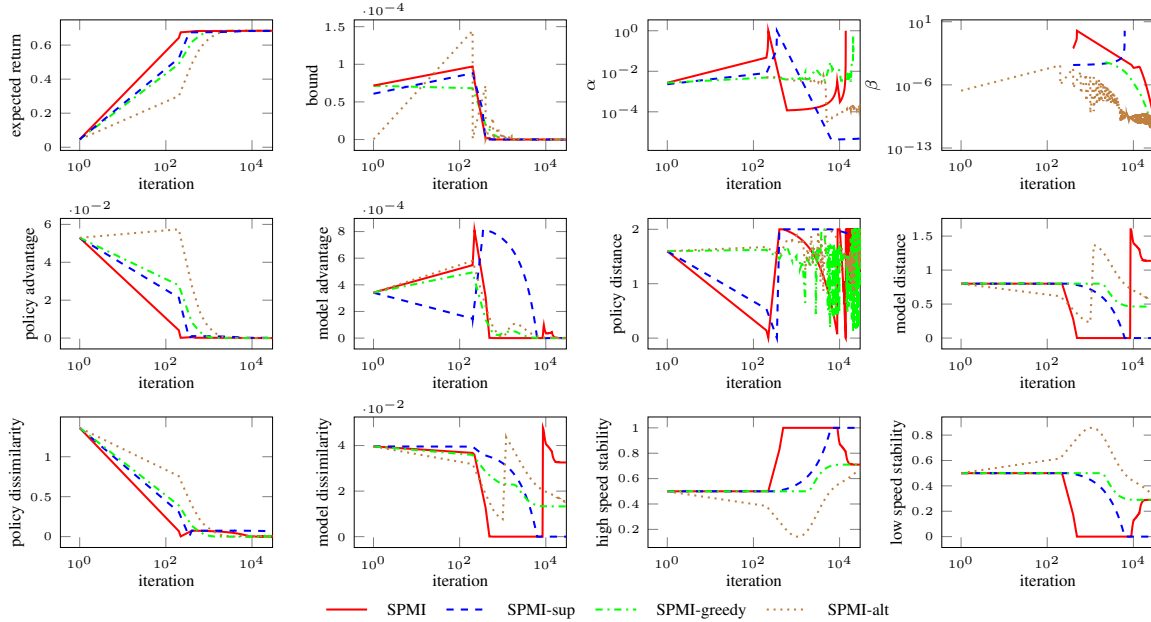


Figure 12: Several statistics of different update strategies and target choices in the case of Racetrack Simulator in the T1 considering vehicle stability only.

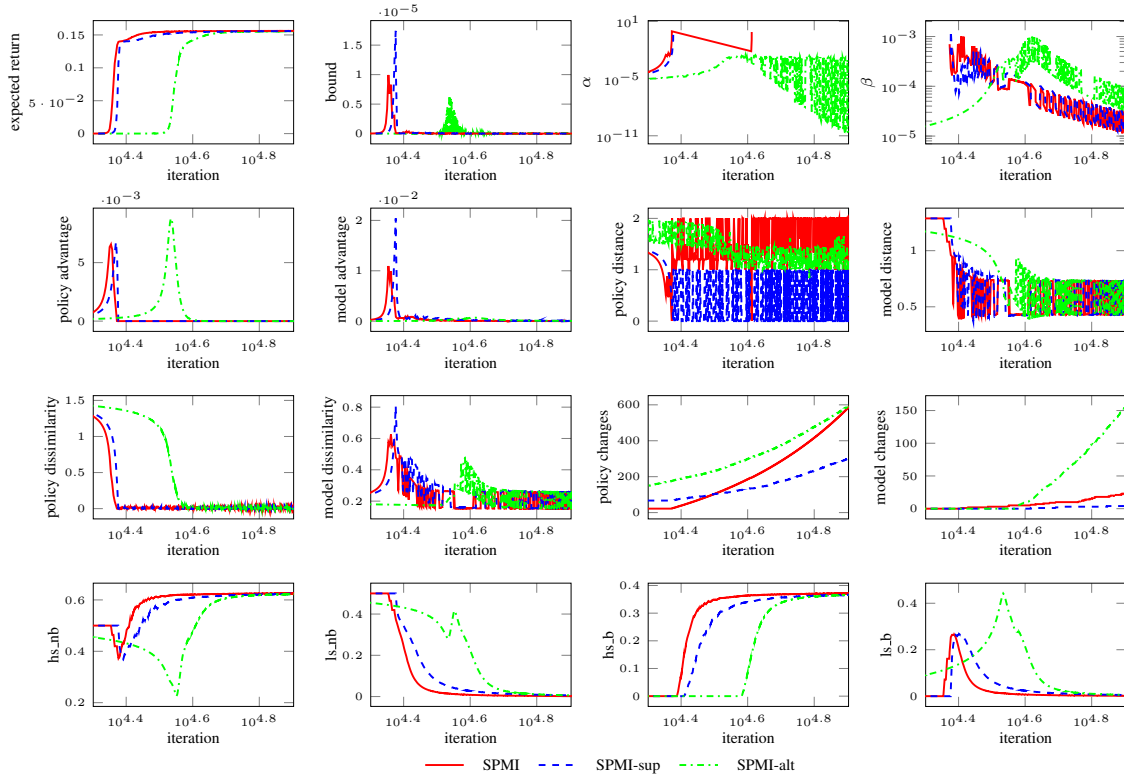


Figure 13: Several statistics of different update strategies in the case of Racetrack Simulator in the T2 considering vehicle stability and engine setting.




Figure 14: Expected return of the Racetrack Simulator in the T2, T3, T4 for different update strategies and considering vehicle stability configuration only.