

Online Markov Decision Processes Configuration with Continuous Decision Space

Davide Maran*, Pierriccardo Olivieri*, Francesco Emanuele Stradi*,
Giuseppe Urso, Nicola Gatti, Marcello Restelli

Politecnico di Milano,

{davide.maran, pierriccardo.olivieri, francescoemanuele.stradi, nicola.gatti, marcello.restelli}@polimi.it,
giuseppe.urso@mail.polimi.it

Abstract

In this paper, we investigate the optimal online configuration of episodic Markov decision processes when the space of the possible configurations is continuous. Specifically, we study the interaction between a learner (referred to as the configurator) and an agent with a fixed, unknown policy, when the learner aims to minimize her losses by choosing transition functions in online fashion. The losses may be unrelated to the agent’s rewards. This problem applies to many real-world scenarios where the learner seeks to manipulate the Markov decision process to her advantage. We study both deterministic and stochastic settings, where the losses are either fixed or sampled from an unknown probability distribution. We design two algorithms whose peculiarity is to rely on occupancy measures to explore with optimism the continuous space of transition functions, achieving constant regret in deterministic settings and sublinear regret in stochastic settings, respectively. Moreover, we prove that the regret bound is tight with respect to any constant factor in deterministic settings. Finally, we compare the empiric performance of our algorithms with a baseline in synthetic experiments.

Introduction

Reinforcement Learning (RL) investigates the sequential interaction between a learner and an environment, aiming at continually improving the learner’s strategy (Sutton and Barto 2018). In this context, the environment is customarily represented as a Markov Decision Process (MDP) with a fixed but unknown transition function. We study a general scenario where the interaction occurs in episodes, each with a predetermined length. Differently from the standard RL setting, we consider the learner not to be the agent playing the MDP, but the *configurator*. Precisely, at each episode, the learner picks the transition functions for the entire MDP (*i.e.*, a configuration) from a fixed continuous set. Next, she observes the loss suffered and the path traversed by the agent, which depend both on the agent’s fixed policy and the transition chosen for the specific episode. The aim of the configurator is to minimize her regret between her total loss and that provided by an optimal fixed configuration.

Our model represents various real-world situations where the learner aims to manipulate the stochastic nature of the

MDP to her advantage. For example, consider the sale of hotel rooms, where the MDP states are characterized by the number of rooms booked in different categories each day, while the transitions depend on the hotel’s pricing and user behavior. Customarily, the hotels use a fixed pricing strategy that is trained offline and implemented online (technically, it is a post-training scenario). Given that users compare prices across hotels before booking rooms, a *competing* hotel (acting as the MDP configurator) can strategically adjust its pricing to influence user behavior and consequently alter the MDP transitions. Specifically, the competitor seeks to reduce the number of room reservations obtained by the agent to maximize her own.

Although this example illustrates an adversarial setting, our model applies to general scenarios that do not require a relationship between the configurator’s loss and the agent’s reward.

Related Work

Online learning in MDPs Several works initially introduced for on online learning (Cesa-Bianchi and Lugosi 2006; Hazan 2019) have been subsequently extended to MDPs (Auer, Jaksch, and Ortner 2008; Even-Dar, Kakade, and Mansour 2009; Neu et al. 2010). In particular, (Azar, Osband, and Munos 2017) study the problem of optimal exploration in episodic MDPs with unknown transitions and stochastic losses when the feedback is bandit. (Rosenberg and Mansour 2019a) study the online learning problem in episodic MDPs with adversarial losses and unknown transitions when the feedback is full information, presenting an online algorithm which provides a regret upper bound of $\tilde{O}(\sqrt{T})$, where T is the number of episodes. (Rosenberg and Mansour 2019b) study the same setting when the feedback is bandit, providing a regret upper bound of $\tilde{O}(T^{3/4})$, which was subsequently improved to $\tilde{O}(\sqrt{T})$ by (Jin et al. 2019).

Configurable MDPs In MDPs, the transition function is customarily assumed to be fixed, see, *e.g.*, (Sutton and Barto 2018). However, various subsequent works represent environments with non-fixed transition probabilities, as provided in the works by (Satia and Lave 1973), (White and Eldeib 1994), and (Bueno et al. 2017). Recently, the concept of Configurable Markov Decision Processes (Conf-MDPs) was formalized by (Metelli, Mutti, and Restelli 2018). In particu-

*These authors contributed equally.

lar, the authors propose an algorithm capable of optimizing, at the same time, the environment configuration, namely, the transition function and the policy of the learning agent. This line of research has been further expanded upon by (Metelli, Ghelfi, and Restelli 2019) and (Metelli, Manneschi, and Restelli 2022). Moreover, (Ramponi et al. 2021) extend the Conf-MDP setting to an online learning framework. This scenario involves a configurator who chooses online a transition function from a discrete set and aims to maximize her own reward, which is independent from the agent’s one.

Adversarial Attacks Several works deal with adversarial attacks in MDPs, see, *e.g.*, (Ilahi et al. 2021). In the *bounded state attacks* framework, the adversary can manipulate the current state of an MDP in order to force the learning agent to make suboptimal decisions, see, *e.g.*, (Pattanaik et al. 2017), (Korkmaz 2021), and (Wu et al. 2022). Instead, in the *action attacks* setting, the adversary is capable of modifying the agent’s actions, see, *e.g.*, (Lee et al. 2019), (Lee et al. 2021) and (Tan et al. 2020). Finally, in the *model attacks* framework, the attack consists in a (bounded) perturbation of the transition function of the MDP performed by an adversary, see, *e.g.*, (Rakhsha et al. 2020).

Original Contribution

We investigate the problem of *online configuration with continuous decision space* in MDPs, where the rewards may be both *deterministic* or *stochastic*. Precisely, we study the problem of an online configurator which chooses at any round a transition function from a continuous decision space and receives a loss which depends on both the configuration chosen and the fixed policy of the agent she is interacting with. First, we show that our setting can be seen as an instance of the well-known *Lipschitz bandit* framework, as well as a generalization of many post-training *adversarial attacks* models. Then, we propose two algorithms, namely, O-DOSC (Online Deterministic Optimistic Configuration Search) for deterministic settings and O-SOSC (Online Stochastic Optimistic Configuration Search) for the stochastic ones. We prove that O-DOSC achieves constant regret, matching the lower bound that we provide for the deterministic setting. Then, we show that O-SOSC achieves a $\tilde{O}(\sqrt{T})$ regret bound in stochastic settings. Finally, we empirically validate our results with synthetic simulations.

Problem Formulation

Online MDPs

We introduce *online episodic loop-free* MDPs $\mathcal{M} = (X, A, P, \mathcal{R})$ defined as follows.

- T is the number of episodes, with $t \in [T]$ denoting a specific episode.
- X and A are the finite state and action spaces, respectively. By the loop-free property, X is partitioned into H layers X_0, \dots, X_H such that the first and the last layers are singletons, *i.e.*, $X_0 = \{x_0\}$ and $X_H = \{x_H\}$. We will refer to H as the horizon. Moreover, we denote as $h(x)$ the layer of a specific state x .

- $P : X \times A \rightarrow \Delta(X)$ is the transition function, where, for ease of notation, we denote by $P(x'|x, a)$ the probability of going from state $x \in X$ to $x' \in X$ by taking action $a \in A$. By the loop-free property, it holds that $P(x'|x, a) > 0$ only if $x' \in X_{h+1}$ and $x \in X_h$ for some $h \in [0 .. H - 1]$.
- \mathcal{R} is the reward function, which can be *deterministic*, that is, $\mathcal{R} : X \times A \rightarrow [0, 1]$, or *stochastic*, namely a distribution over $[0, 1]$ for every (x, a) . We refer to the reward of a specific state-action pair $x \in X, a \in A$ for a specific episode $t \in [T]$ as $r_t(x, a)$.

Remark 1. Any episodic MDP with horizon H that is not loop-free can be cast into a loop-free one by suitably duplicating the state space H times, *i.e.*, a state x is mapped to a set of new states (x, h) , where $h \in [0 .. H]$.

A policy $\pi : X \rightarrow \Delta(A)$ defines a probability distribution over actions at each state. For ease of notation, we denote by $\pi(\cdot|x)$ the probability distribution for a state $x \in X$, with $\pi(a|x)$ denoting the probability of action $a \in A$.

Continuous Configurable-MDPs

The framework we propose, called Continuous Configurable-MDPs, is characterized by:

- an *agent*, which is characterized by the (optimal) policy π^* of a fixed MDP $\mathcal{M}(X, A, \bar{P}, \mathcal{R})$. We assume, without loss of generality, that π^* is deterministic, since it is well known that MDPs always admit an optimal deterministic policy;
- a *configurator*, which knows X, A, \bar{P}, H, T and at every episode $t \in [T]$ can choose a configuration (*i.e.*, a transition function) P_t from a bounded set \mathcal{I} , in order to minimize her loss. Similarly to the reward function, the loss function can be *deterministic*, that is, $\ell : X \times A \rightarrow [0, 1]$, or *stochastic*, namely a distribution over (x, a) , still bounded in $[0, 1]$. We refer to the loss of a specific state-action pair $x \in X, a \in A$ for a specific episode $t \in [T]$ as $\ell_t(x, a)$. In the *stochastic setting*, $\ell_t(x, a)$ is a sample from $\mathcal{L}(x, a)$, drawn independently from the past. The average of loss distribution for (x, a) , which does not depend on t , is denoted as $\ell(x, a) = \mathbb{E}[\mathcal{L}(x, a)]$.

Customarily in the literature, it is assumed that the configurator’s loss is directly tied to the agent’s reward, namely $\ell_t(x, a) = r_t(x, a)$ for every state-action pair and for every episode. Instead, in our setting, the two functions can be independent.

In Algorithm 1, we report the interaction between the agent and the configurator in the online MDP.

Precisely, at the beginning of each episode t , the loss function is either *deterministically* chosen (although this term may be slightly abused in this context) or *stochastically* chosen (refer to Line 2). Subsequently, the configurator chooses a transition function P_t (as in Line 3), and the MDP is initialized in the state x_0 (as per Line 4). During the episode, the agent traverses all the layers based on her policy π^* (as described in Line 6) and the transition P_t (as per Line 7). Upon completion of the episode, the configurator observes the complete trajectory and losses (as stated in Line 9).

Algorithm 1: Agent-Configurator Interaction

```

1: for  $t \in [T]$  do
2:   For every state-action pair,  $\ell_t(x, a) = \ell(x, a)$  in the
   deterministic setting or  $\ell_t(x, a) \sim \mathcal{L}(x, a)$  in the
   stochastic setting
3:   configurator chooses  $P_t \in \mathcal{I}$ 
4:   state is initialized to  $x_0$ 
5:   for  $h = 0, \dots, H - 1$  do
6:     agent plays  $a_h \sim \pi^*(\cdot|x_h)$ 
7:     environment evolves to  $x_{h+1} \sim P_t(\cdot|x_h, a_h)$ 
8:   end for
9:   configurator observes  $\{x_h, a_h\}_{h=0}^{H-1}$  and suffers
    $\{\ell_t(x_h, a_h)\}_{h=0}^{H-1}$ 
10: end for

```

Occupancy Measures

We introduce the notion of *occupancy measure*, see (Rosenberg and Mansour 2019b). Given a transition function P and a policy π , the occupancy measure $d^{P,\pi} \in [0, 1]^{|X \times A \times X|}$ induced by P and π is such that, for every $x \in X_h$, $a \in A$, and $x' \in X_{h+1}$ with $h \in [0 .. H - 1]$:

$$d^{P,\pi}(x, a, x') = \mathbb{P}[x_h = x, a_h = a, x_{h+1} = x' | P, \pi]. \quad (1)$$

Moreover, we also define:

$$d^{P,\pi}(x, a) = \sum_{x' \in X_{h+1}} d^{P,\pi}(x, a, x'), \quad (2)$$

$$d^{P,\pi}(x) = \sum_{a \in A} d^{P,\pi}(x, a). \quad (3)$$

Then, we can introduce the following lemma, which characterizes *valid* occupancy measures.

Lemma 1. (Rosenberg and Mansour 2019a) *For every $d \in [0, 1]^{|X \times A \times X|}$, it holds that d is a valid occupancy measure of an episodic loop-free MDP if and only if, for every $h \in [0 .. H - 1]$, the following three conditions hold:*

1. $\sum_{x \in X_h} \sum_{a \in A} \sum_{x' \in X_{h+1}} d(x, a, x') = 1$
2. $\sum_{a \in A} \sum_{x' \in X_{h+1}} d(x, a, x') = \sum_{x' \in X_{h-1}} \sum_{a \in A} d(x', a, x)$,
 $\forall x \in X_h$
3. $P^d = P$

where P is the transition function of the MDP and P^d is the one induced by d (see Equation (4)).

Notice that any occupancy measure d induces a transition function P^d and a policy π^d as:

$$P^d(x'|x, a) = \frac{d(x, a, x')}{d(x, a)}, \quad \pi^d(a|x) = \frac{d(x, a)}{d(x)}. \quad (4)$$

Performance Metric

In order to have a proper performance metric for our algorithms, we introduce the notion of objective function of an MDP (in terms of loss).

Definition 1 (Expected Loss). *The expected loss suffered by the configurator at episode t is defined as the expected value of the sum of the losses given the configuration chosen. Namely,*

$$J_t^\pi(P) := \mathbb{E} \left[\sum_{h=1}^H \ell_t(x_h, a_h) \middle| \pi, P \right].$$

Thus, we define the cumulative regret as follows.

Definition 2 (Cumulative Regret). *The cumulative regret is defined as*

$$R_T := \sum_{t=1}^T J_t^\pi(P_t) - J_t^\pi(P^*),$$

where $P^* := \arg \min_{P \in \mathcal{I}} J_t^\pi(P)$.

Following the formulation based on the occupancy measure, the cumulative regret can be written as $R_T := \sum_{t=1}^T \ell^\top d^{P_t, \pi^*} - \min_{P \in \mathcal{I}} \sum_{t=1}^T \ell^\top d^{P, \pi^*}$, or equivalently, $R_T := \sum_{t=1}^T \ell^\top d^{P_t, \pi^*} - \min_{d \in \Delta(\mathcal{I}, \pi^*)} \sum_{t=1}^T \ell^\top d$, where $d^{P,\pi}$ is the occupancy measure vector defined on the tuple (x, a) given a transition function P and a policy π , $\Delta(\mathcal{I}, \pi^*)$ is the space of occupancy measures built given the fixed policy π^* and the transition function space \mathcal{I} , and ℓ is defined as:

- in the *deterministic* setting, ℓ is the loss vector composed by the loss values associated to each tuple (x, a) , namely $\ell(x, a)$,
- in the *stochastic* setting, ℓ , is the vector composed by the expected values of the loss distribution for every (x, a) , namely, $\mathbb{E}[\mathcal{L}(x, a)]$.

Given the definition of this setting, we aim that the regret is sublinear in T , namely $R_T = o(T)$.

The optimization problem described above is linear in the space of the occupancy measures, suggesting the potential adoption of online convex programming tools such as, e.g., Bandit Linear Optimization (BLO) algorithms proposed by (Abernethy, Hazan, and Rakhlin 2008). However, these methods cannot be adopted to our case. Indeed, without the knowledge of the agent's policy, the configurator cannot compute the exact occupancy measure corresponding to her transition and the agent's policy, thus precluding the design of online bandit linear optimization algorithms working on the occupancy measure space. In particular, the configurator can only choose a transition function P_t and the objective function is highly nonlinear in the space of the transition functions.

Generality of the Setting and Interpretation

Our model captures various settings. In the following, we provide two different interpretations. The first focuses on MDPs with adversarial attacks, while the second focuses on Lipschitz bandits.

Interpreting Our Model as an MDP with Adversarial Attacks

Selecting the proper configurator’s decision space \mathcal{I} , several forms of *adversarial attacks* in MDPs can be described by our model. Since the agent’s policy is assumed to be fixed, our model captures post-training attacks.

In the following, the *adversarial attacks* modeled by our setting are presented with the associated configurator’s continuous decision space.

- *Bounded state attacks.* The adversary can modify the agent’s state, substituting it with another state that is similar to the original one. This can be modeled by setting:

$$\mathcal{I} = \{P : \forall x \in X, a \in A, \exists x' \in B(x), \\ P(\cdot|x, a) = \bar{P}(\cdot|x', a)\},$$

where $B(x) = \{x' : d(x, x') < \varepsilon\}$ for some distance function $d(\cdot)$ and $\varepsilon > 0$.

- *Action attacks.* Differently from the state attack scenarios, the adversary can perturb the action of the agent. This kind of attacks can be modeled by setting:

$$\mathcal{I} = \{P : \forall x \in X, a \in A, \exists a' \in B(a), \\ P(\cdot|x, a) = \bar{P}(\cdot|x, a')\},$$

where the set $B(a)$ is defined as in the case of bounded state attacks.

- *Model attacks.* The adversary can change the transition probabilities and the amount of the change is upper bounded according to some metrics. In particular, we adopt the *total variation metric*, denoted with TV. Therefore, \mathcal{I} can be defined as

$$\mathcal{I} = \{P : \text{TV}(P, \bar{P}) < \varepsilon\},$$

for some $\varepsilon > 0$, where $\text{TV}(P, P') := \sum_{x \in X, a \in A} \|P(\cdot|x, a) - P'(\cdot|x, a)\|_1$.

Interpreting Our Model as a Lipschitz Bandit

We can show that the optimization problem faced by the configurator can be seen as a *Lipschitz bandit*, namely, the objective function in the optimization problem is *Lipschitz* continuous. This result is crucial in order to have a proper baseline to compare our theoretical guarantees to. Indeed, standard multi-armed bandits techniques cannot be applied to our setting, since the decision space is continuous.

The Lipschitz continuity of the objective is already well-known for discounted MDPs (Munos and Szepesvári 2008); nevertheless, we report here the result for the case of finite horizon problems.

Theorem 2. *Let P, P' be two transition functions, and π an arbitrary Markovian policy. Then,*

$$|J^\pi(P) - J^\pi(P')| \leq \frac{H^2}{2} \text{TV}(P, P'),$$

where $\text{TV}(P, P') := \sum_{x \in X, a \in A} \|P(\cdot|x, a) - P'(\cdot|x, a)\|_1$.

Theorem 2 suggests that algorithms for Lipschitz bandits can be used to solve our problem. In the specific case of the Zooming algorithm by (Kleinberg, Slivkins, and Upfal 2008)—one of the state-of-the-art algorithms for Lipschitz bandits—, we can derive the following upper regret bound.

Corollary 3. *The Zooming algorithm in our setting achieves a regret of $R_T \leq T^{\frac{1+\mathcal{D}(\mathcal{I})}{2+\mathcal{D}(\mathcal{I})}}$, where $\mathcal{D}(\mathcal{I})$ is the Zooming dimension of the space \mathcal{I} .*

When the decision space \mathcal{I} depends on a family of p continuous parameters, its Zooming dimension is exactly p , so that the regret becomes $T^{\frac{1+p}{2+p}}$. As we show in the following, this regret bound can be dramatically improved and therefore the Zooming algorithm is suboptimal for our problem.

Deterministic Settings

We focus on deterministic settings, and we present our algorithm and its theoretical guarantees. More precisely, we assume there is a fixed function $\ell : X \times A \rightarrow [0, 1]$, such that the configurator will always achieve the same loss whenever the agent chooses a particular action in a given state.

Algorithm

Algorithm 2 provides the pseudo-code of *Online Deterministic Optimistic Configuration Search* (O-DOSC), which tackles deterministic losses. As is customary in the online learning, the configurator needs to face an exploration-exploitation trade-off when searching for the optimal configuration. Specifically, the choice of P_t needs to balance the exploration of unobserved states with the minimization of the configurator’s losses.

As stated above, we assume that the optimal policy π^* in the MDP is deterministic. Thus, our algorithm can safely keep track of the actions played and losses obtained. For this purpose, the set Π is initialized to contain all possible deterministic policies, while the function $\hat{\ell}$ is initialized to return a loss value of 0 for every tuple (x, a) (Lines 1–2). Such an initialization for the function $\hat{\ell}$ is chosen to guarantee optimism vs. uncertainty with respect to the actual loss function.

In order to determine the transition function P_t for each episode, an optimistic approach is adopted. In particular, we minimize the objective over the space of the occupancy measures, which is based on an estimate of the agent’s policy (as reported in Line 4). This approach is optimistic with respect to both the policy and the loss function, which is set to be 0 when non-visited. Additionally, it is possible to simplify the optimization over \mathcal{I} and Π by reducing to the optimization over the space $\Delta(\mathcal{I}, \Pi)$, where $d^{P, \pi} \in \Delta(\mathcal{I}, \Pi)$. For a detailed study of the computational complexity of the minimization update, please refer to the Appendix.

Then, once the agent’s trajectory and losses suffered throughout the path have been observed (Line 5), the following updates are performed. For $\hat{\ell}$, the 0 values associated with the tuples (x, a) visited during the episode are substituted with the observed losses (Line 6). Instead, for the set Π , the actions of the state traversed but not executed by the agent are discarded from the set (Line 7).

Algorithm 2: O-DOSC Algorithm**Require:** X, A, H, \mathcal{I}

- 1: $\Pi \leftarrow$ set of all deterministic policies
- 2: $\hat{\ell}(x, a) \leftarrow 0 \forall (x, a) \in X \times A$
- 3: **for** $t \in [T]$ **do**
- 4: Choose $P_t =$

$$\arg \min_{P \in \mathcal{I}, \pi \in \Pi} \sum_{x, a} \left(\sum_{x' \in X_{h(x)+1}} d^{P, \pi}(x, a, x') \right) \hat{\ell}(x, a)$$

- 5: Observe $\{x_h, a_h, \ell(x_h, a_h)\}_{h=0}^{H-1}$
- 6: $\hat{\ell}(x_h, a_h) \leftarrow \ell(x_h, a_h) \quad \forall h \in [0..H-1]$
- 7: $\Pi \leftarrow \Pi \setminus \{x_h, a\} \quad \forall a \neq a_h, \forall h \in [0..H-1]$
- 8: **end for**

Upper and Lower Regret Bounds

In this section, we present the theoretical guarantees of our O-DOSC algorithm in deterministic settings. Initially, we state the regret bound achieved by our algorithm, and, subsequently, we show that the regret bound matches the lower bound for our specific setting.

In deterministic settings, we show that Algorithm 2 achieves a constant regret bound.

Theorem 4. *In deterministic settings, Algorithm 2 guarantees a regret upper bound*

$$R_T \leq (H+1)|X|.$$

Proof Sketch. First, notice that, due to the deterministic nature of the loss, visiting each state one time is sufficient to compute the optimal policy.

Still, it is necessary to deal with states for which the associated occupancy measure is very low under every transition $P \in \mathcal{I}$ available to the configurator. Since these states are unlikely to be visited in the whole episode, they may prevent the configurator from learning the optimal configuration.

The key observation is that the aforementioned states cannot contribute significantly to the regret: defining the estimation error as follows, $\varepsilon_t := J^\pi(P_t) - LB_t(P_t)$, where $LB_t(P_t) := \min_{P \in \mathcal{I}, \pi \in \Pi} \sum_x d^{P, \pi}(x) \hat{\ell}(x)$, we can show that,

1. The regret is bounded by $R_T \leq \mathbb{E} \left[\sum_{t=1}^T \varepsilon_t \right]$
2. At each step, the undiscovered states \bar{X}_t satisfy $\sum_{x \in \bar{X}_t} d^{P_t, \pi}(x) \geq \frac{\varepsilon_t}{H+1}$

Given the previous steps, it can be derived that,

$$R_T \leq (H+1) \sum_{x \in \bar{X}_t} d^{P_t, \pi}(x),$$

which implies that the higher regret at one step, the higher probability to discover more states, thus lower regret for the future. \square

The previous result is rather intuitive. Indeed, since both the optimal policy and reward function are deterministic,

once the configurator visited the entire MDP, the optimal configuration has been found.

The reader may wonder if the the regret bound shown in Theorem 4 is tight for the setting. In the following, we show that our result is the best *any* algorithm can achieve. Therefore, Algorithm 2 matches the lower bound of the deterministic setting. Indeed, we can show that,

Theorem 5. *In deterministic settings, any algorithm achieves a regret of order $\Omega(H|X|)$.*

Proof Sketch. We properly build a Markov decision process as follows: the second layer is composed by approximately $|X|$ states while the remaining layers by two states only. The agent policy chooses a non-trivial action only at the second step, after which the loss is uniquely determined. Thus, the following $H-3$ steps have the effect of increasing the loss $H-3$ times. In order to get to the optimal state (in the second layer), the configurator has to pull in expectation approximately $|X|/2$ suboptimal configurations, thus augmenting the regret of a $H-3$ factor. This in turns implies a regret of approximately $(H-3)|X|/2$. \square

Stochastic Settings

We focus on stochastic settings, and we present our algorithm and its theoretical guarantees. Precisely, we assume that there is a fixed probability distribution for every state-action pair, denoted as $\mathcal{L}(x, a)$, which drives the sampling of losses from the interval $[0, 1]$ every time the agent chooses an action in a given state.

Algorithm

Algorithm 3 provides the pseudo-code of *Online Stochastic Optimistic Configuration Search* (O-SOSC) for stochastic losses. Similarly to what happens in deterministic settings, the configurator needs to address an exploration-exploitation trade-off when seeking for the optimal configuration. Again, the choice of P_t is required to balance the exploration of non-visited states with the minimization of the configurator's losses. Furthermore, in this case we introduce an additional complexity, given by the way losses are chosen.

By the theory of MDPs, we can safely assume that the optimal policy π^* for the MDP is deterministic. Algorithm 3 keeps track of the action played and the losses obtained by the configurator. For this purpose, the set Π is initialized to containing all possible deterministic policies, while $\hat{\ell}_t$ is initialized to return a loss value of 0 for every tuple (x, a) (Lines 1–2). We choose this initialization for the function $\hat{\ell}_t$ to be optimistic with respect to the actual loss function.

To determine the transition function P_t for each episode, we take an optimistic approach by minimizing the objective over the space of occupancy measures based on an estimate of the agent's policy (as reported in Line 4). It is worth noting that this update is optimistic with respect to both the policy and the loss function, which is set to 0 when non-visited, and is computed with UCB-like lower bound once traversed. Moreover, it is possible to simplify the optimization over \mathcal{I} and Π by reducing it to the optimization over the space $\Delta(\mathcal{I}, \Pi)$, where $d^{P, \pi} \in \Delta(\mathcal{I}, \Pi)$. For a detailed

Algorithm 3: O-SOSC Algorithm**Require:** $X, A, H, \mathcal{I}, \delta, T$

- 1: $\Pi \leftarrow$ set of all deterministic policies
- 2: $\widehat{\ell}_1(x, a) \leftarrow 0 \quad \forall (x, a) \in X \times A$
- 3: **for** $t \in [T]$ **do**
- 4: Choose $P_t =$

$$\arg \min_{P \in \mathcal{I}, \pi \in \Pi} \sum_{x, a} \left(\sum_{x' \in X_{h(x)+1}} d^{P, \pi}(x, a, x') \right) \widehat{\ell}_t(x, a)$$

- 5: Observe $\{x_h, a_h, \ell_t(x_h, a_h)\}_{h=0}^{H-1}$
- 6: **Bonus** $= \sqrt{\frac{-\log(\delta) + \log(N_t(x_h)(N_t(x_h)+1))}{2N_t(x_h)}}$
- 7: $\widehat{\ell}_{t+1}(x_h, a_h) \leftarrow \max(0, \bar{\ell}_t(x_h, a_h) - \text{Bonus})$
 $\quad \quad \quad \forall h \in [0..H-1]$
- 8: $\Pi \leftarrow \Pi \setminus \{x_h, a\} \quad \forall a \neq a_h, \forall h \in [0..H-1]$
- 9: **end for**

study of the computational complexity of the minimization update, please refer to the Appendix.

Once the agent's trajectory and losses suffered throughout the path have been observed (Line 5), the following updates are performed. For $\widehat{\ell}_t$, the values associated with the tuples (x, a) visited during the episode are updated with a UCB-like term that depends on the number of visits of a specific state $N_t(x)$ (Line 6), which is subtracted to the empirical mean $\bar{\ell}_t(x, a)$ of the losses observed. For the set Π , the actions of the state traversed but not executed by the agent are discarded from the set (Line 7).

Theoretical Guarantees

In this section, we present the theoretical result for Algorithm 3. First of all, we can derive a simple lower bound of the regret in the stochastic case.

Theorem 6. *In stochastic settings, any algorithm achieves a regret of order $\Omega(\sqrt{|X|T})$.*

Proof. In order to find a proper lower bound, we restrict to the (simpler) discrete decision space case. In such a setting, our model can be seen as a generalization of the multi-armed bandit setting. Specifically, given any multi-armed bandit problem, we can build an equivalent instance of our problem as follows. For every arm of the bandit problem, we have a transition function in \mathcal{I} bringing deterministically from a common initial state to a different state. This implies that the number of transition functions in the MDP equals the number of arms of the bandit problem ($|\mathcal{I}| = |X|$). Therefore, the standard instance independent lower bound for multi-armed bandits with $|\mathcal{I}| = |X|$ number of arms leads to a regret of $R_T = \Omega(\sqrt{|X|T})$ which represents a lower bound for our problem. \square

Finally, we show that Algorithm 3 achieves a sublinear regret bound that matches the aforementioned lower bound in the number of episodes T .

Theorem 7. *In the stochastic setting, for the choice $\delta = T^{-1/2}$, Algorithm 3 achieves a regret upper bounded as follows,*

$$R_T = \widetilde{O}\left(|X|\sqrt{T} + H|X|\right).$$

Proof Sketch. In the stochastic setting, the configurator has to visit each state more than one time, due to the stochastic nature of the loss.

Analogously to stochastic multi-armed bandits, the aforementioned issue is faced defining a *good event* under which the configurator is able to learn a good policy by only visiting all the states for a fixed and not too large amount of times.

We refer to the *good events* $s E^c$, where E corresponds to

$$E := \left\{ \exists x \in X, t \in [T] : |\bar{\ell}_t(x) - \ell(x)| > \sqrt{\frac{-\log(\delta) + \log(N_t(x)(N_t(x)+1))}{2N_t(x)}} \right\},$$

and where $N_t(x)$ is the number of times state x has been visited by round t , and $\delta = T^{-1/2}$. With a standard concentration inequality and a union bound, we can prove that $\mathbb{P}(E) \leq 2|X|\delta$. Next, it is necessary to focus on the *good event* E^c .

Under E^c , the proof follows the one for the deterministic case with a fixed number of visits for every state, showing that the expected regret conditioned to the event is of order $\widetilde{O}\left(|X|\sqrt{T} + H|X|\right)$. \square

Empirical Evaluation

In this section, we experimentally evaluate the performance of Algorithms 2 and 3 in terms of empiric regret. We describe the results obtained in the deterministic and stochastic settings separately. In each case, we conduct experiments with both discrete and continuous decision spaces \mathcal{I} .

As a baseline, we opt for UCB1 (Auer, Cesa-Bianchi, and Fischer 2002) since, in the case of discrete decision spaces, UCB1 is a standard baseline, while, in the case of continuous decision spaces, UCB1 can be preferred to Zooming (Kleinberg, Slivkins, and Upfal 2008) for two reasons. The first reason is that the design of a suitable covering oracle for Zooming raises several conceptual and computational issues due to the high number of dimensions whose solutions is open. The second reason is that, in our experimental settings, the optimal solution is one of the arms, and, in these cases, UCB1 is a more severe baseline than Zooming as it guarantees a much better regret bound.

In the following experiments, we consider a Markov decision process structured as follows. The MDP consists of four layers. As is standard in the loop-free model, the first and the last layers are singletons, while the second and third layers each comprise two states. Additionally, every state is associated with two actions. For reasons of space, the description of the experimental settings and additional details on the experimental results can be found in the Appendix.




Figure 1: Average cumulative regret with a 95% confidence interval over 10 experiments in deterministic settings with discrete (a, b) and continuous (c) decision spaces.




Figure 2: Average cumulative regret with a 95% confidence interval over 10 experiments in stochastic settings with discrete (a) and continuous (b) decision spaces.

Deterministic Settings We report in Figure 1 the experimental results obtained with deterministic settings where the cumulative regret is averaged over 10 runs.

In particular, Figure 1(a) shows the results with discrete settings, while Figure 1(c) shows the results with continuous settings. In both cases, O-DOSC dramatically outperforms UCB1. Figure 1(b) clearly shows that O-DOSC effectively computes the optimal transition function during the very initial rounds, and subsequently it ceases to explore. Indeed, once O-DOSC visited all the states, it can numerically compute the optimal transition. Instead, UCB1 keeps exploring for a long time.

Stochastic Settings We report in Figure 2 the experimental results obtained with deterministic settings where the cumulative regret is averaged over 10 runs.

Precisely, Figure 2(a) shows the results with discrete settings, while Figure 2(b) shows the results with continuous settings. In both cases, O-SOSC outperforms UCB1. Differently from what happens in deterministic settings, O-SOSC does not find the optimal solution in the initial rounds, and additional exploration is required. However, the performance exhibited by O-SOSC in this setting is remarkably impressive.

Conclusions

In this paper, we propose the problem of *online configuration* of Markov decision processes with *continuous decision spaces*. We study the problem both when the losses are deterministic and stochastic. We propose O-DOSC algorithm, which achieves constant regret in deterministic settings, and we show that this result is tight for any constant with respect to the lower bound. Then, we propose O-SOSC which achieves a sublinear regret bound when the losses are stochastic. Finally, we empirically validate our theoretical results with synthetic simulations.

Future Works In future work, we are interested in studying the problem when losses are *adversarial*, namely no statistical assumption are made. Furthermore, we aim to study the problem of *online configurations* against a learning agent, namely, when the policy of the agent is allowed to be dynamic. In such a setting, we are interested in understanding how the results change when the agent is employing a no-regret optimizer and when she is omniscient, namely, she can observe the transitions chosen by the configurator and then commit to a policy as in a Stackelberg game.

References

- Abernethy, J. D.; Hazan, E.; and Rakhlin, A. 2008. Competing in the Dark: An Efficient Algorithm for Bandit Linear Optimization. In *Annual Conference Computational Learning Theory*.
- Auer, P.; Cesa-Bianchi, N.; and Fischer, P. 2002. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47(2): 235–256.
- Auer, P.; Jaksch, T.; and Ortner, R. 2008. Near-optimal Regret Bounds for Reinforcement Learning. In Koller, D.; Schuurmans, D.; Bengio, Y.; and Bottou, L., eds., *Advances in Neural Information Processing Systems*, volume 21. Curran Associates, Inc.
- Azar, M. G.; Osband, I.; and Munos, R. 2017. Minimax Regret Bounds for Reinforcement Learning.
- Bueno, T. P.; Mauá, D. D.; Barros, L. N.; and Cozman, F. G. 2017. Modeling Markov Decision Processes with Imprecise Probabilities Using Probabilistic Logic Programming. In Antonucci, A.; Corani, G.; Couso, I.; and Destercke, S., eds., *Proceedings of the Tenth International Symposium on Imprecise Probability: Theories and Applications*, volume 62 of *Proceedings of Machine Learning Research*, 49–60. PMLR.
- Cesa-Bianchi, N.; and Lugosi, G. 2006. *Prediction, learning, and games*. Cambridge university press.
- Even-Dar, E.; Kakade, S. M.; and Mansour, Y. 2009. Online Markov decision processes. *Mathematics of Operations Research*, 34(3): 726–736.
- Hazan, E. 2019. Introduction to Online Convex Optimization. *CoRR*, abs/1909.05207.
- Ilahi, I.; Usama, M.; Qadir, J.; Janjua, M. U.; Al-Fuqaha, A.; Hoang, D. T.; and Niyato, D. 2021. Challenges and Countermeasures for Adversarial Attacks on Deep Reinforcement Learning. arXiv:2001.09684.
- Jin, C.; Jin, T.; Luo, H.; Sra, S.; and Yu, T. 2019. Learning Adversarial MDPs with Bandit Feedback and Unknown Transition.
- Kleinberg, R.; Slivkins, A.; and Upfal, E. 2008. Multi-Armed Bandits in Metric Spaces. arXiv:0809.4882.
- Korkmaz, E. 2021. Investigating Vulnerabilities of Deep Neural Policies. arXiv:2108.13093.
- Lee, X. Y.; Esfandiari, Y.; Tan, K. L.; and Sarkar, S. 2021. Query-based Targeted Action-Space Adversarial Policies on Deep Reinforcement Learning Agents. arXiv:2011.07114.
- Lee, X. Y.; Ghadai, S.; Tan, K. L.; Hegde, C.; and Sarkar, S. 2019. Spatiotemporally Constrained Action Space Attacks on Deep Reinforcement Learning Agents. arXiv:1909.02583.
- Metelli, A. M.; Ghelfi, E.; and Restelli, M. 2019. Reinforcement Learning in Configurable Continuous Environments. In Chaudhuri, K.; and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, 4546–4555. PMLR.
- Metelli, A. M.; Manneschi, G.; and Restelli, M. 2022. Policy space identification in configurable environments. *Machine Learning*, 111(6): 2093–2145.
- Metelli, A. M.; Mutti, M.; and Restelli, M. 2018. Configurable Markov Decision Processes. arXiv:1806.05415.
- Munos, R.; and Szepesvári, C. 2008. Finite-Time Bounds for Fitted Value Iteration. *Journal of Machine Learning Research*, 9(5).
- Neu, G.; Antos, A.; György, A.; and Szepesvári, C. 2010. Online Markov decision processes under bandit feedback. *Advances in Neural Information Processing Systems*, 23.
- Pattanaik, A.; Tang, Z.; Liu, S.; Bommannan, G.; and Chowdhary, G. 2017. Robust Deep Reinforcement Learning with Adversarial Attacks. arXiv:1712.03632.
- Rakhsha, A.; Radanovic, G.; Devidze, R.; Zhu, X.; and Singla, A. 2020. Policy Teaching via Environment Poisoning: Training-time Adversarial Attacks against Reinforcement Learning. arXiv:2003.12909.
- Ramponi, G.; Metelli, A. M.; Concetti, A.; and Restelli, M. 2021. Learning in Non-Cooperative Configurable Markov Decision Processes. In Ranzato, M.; Beygelzimer, A.; Dauphin, Y.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems*, volume 34, 22808–22821. Curran Associates, Inc.
- Rosenberg, A.; and Mansour, Y. 2019a. Online Convex Optimization in Adversarial Markov Decision Processes. In Chaudhuri, K.; and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, 5478–5486. PMLR.
- Rosenberg, A.; and Mansour, Y. 2019b. Online Stochastic Shortest Path with Bandit Feedback and Unknown Transition Function. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; d Alché-Buc, F.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Satia, J. K.; and Lave, R. E. 1973. Markovian Decision Processes with Uncertain Transition Probabilities. *Operations Research*, 21(3): 728–740.
- Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement Learning: An Introduction*. Cambridge, MA, USA: A Bradford Book. ISBN 0262039249.
- Tan, K. L.; Esfandiari, Y.; Lee, X. Y.; Aakanksha; and Sarkar, S. 2020. Robustifying Reinforcement Learning Agents via Action Space Adversarial Training. arXiv:2007.07176.
- White, C. C.; and Eldeib, H. K. 1994. Markov Decision Processes with Imprecise Transition Probabilities. *Operations Research*, 42(4): 739–749.
- Wu, F.; Li, L.; Huang, Z.; Vorobeychik, Y.; Zhao, D.; and Li, B. 2022. CROP: Certifying Robust Policies for Reinforcement Learning through Functional Smoothing. arXiv:2106.09292.