

Online MDP with Prototypes Information: A Robust Adaptive Approach

Shuo Sun¹, Meng Qi², Zuo-Jun Max Shen^{1,3}

¹Department of Industrial Engineering and Operations Research, UC Berkeley, Berkeley, California 94720, USA

²SC Johnson College of Business, Cornell University, Ithaca, New York 14853, USA

³Faculty of Engineering and Faculty of Business and Economics, The University of Hong Kong, Hong Kong, China
shuo_sun@berkeley.edu, mq56@cornell.edu, maxshen@berkeley.edu

Abstract

In this work, we consider an online robust Markov Decision Process (MDP) where we have the information of finitely many prototypes of the underlying transition kernel. We consider an adaptively updated ambiguity set of the prototypes and propose an algorithm that efficiently identifies the true underlying transition kernel while guaranteeing the performance of the corresponding robust policy. To be more specific, we provide a sublinear regret of the subsequent optimal robust policy. We also provide an early stopping mechanism and a worst-case performance bound of the value function. In numerical experiments, we demonstrate that our method outperforms existing approaches, particularly in the early stage with limited data. This work contributes to robust MDPs by considering possible prior information about the underlying transition probability and online learning, offering both theoretical insights and practical algorithms for improved decision-making under uncertainty.

Extended version — <https://arxiv.org/abs/2412.14075>

1 Introduction

Markov Decision Processes (MDPs) have become a fundamental framework for sequential decision-making under uncertainty, with applications spanning diverse fields such as control, healthcare and supply chain management. Despite their widespread use, MDPs often face challenges when the true transition dynamics are unknown, potentially leading to suboptimal decisions.

In many real-world scenarios, decision-makers may rely on external datasets to parameterize the MDP model, but have access to multiple plausible model estimates, each potentially leading to different optimal policies. This setting is commonly seen in many applications, for example, the healthcare system (Steimle, Kaufman, and Denton 2021). Consider the context of optimizing its breast cancer screening protocol. Decision-makers might have access to local hospital data, a national cancer research institute’s model, and an international meta-analysis. Each source could suggest a different optimal screening frequency and age range for mammograms. This situation exemplifies the challenge

of determining which model to trust or how to integrate insights from multiple sources to create a robust and effective policy when faced with various plausible model estimates. Similar challenges with multiple transition models arise in recommendation systems, supply chain management, and other domains where early performance and worst-case guarantees are crucial (Chatterjee et al. 2020). Moreover, the concept of multiple parameter models is analogous to the scenario-based stochastic programming literature, where each scenario represents a different possibility of the uncertain parameters.

In this work, we focus on this multi-model setting where there are multiple models (prototypes) of the transition probabilities of the underlying Markov chain and the goal is to identify the true model and therefore solve for the optimal policy. Moreover, we address the problem in an online setting that we need to make real-time decisions with streaming data while knowing the prototypes. These prototypes could be estimated from offline dataset. The key challenge in such settings is two-fold: First, we need to efficiently identify the true underlying transition model while making decisions in real-time. Second, and perhaps more critically, we must ensure good performance during the learning phase when data is limited and model uncertainty is high. Classical online MDP algorithms focus primarily on achieving sublinear regret but may perform poorly in early stages and lack worst-case performance guarantees.

To address these challenges, we propose a novel robust learning algorithm that efficiently identifies the true transition kernel while guaranteeing model performance during the exploration stage. Our approach gradually updates the discrete prototype set and calculates the optimal robust policy, which achieves sublinear regret and provides a lower bound for the algorithm performance at each episode. As data accumulates, we propose a termination mechanism that efficiently identifies the true transition kernel.

Our work differentiates itself from existing approaches in several key aspects. First, we consider an online MDP with structural information of prototypes, which has not been studied before. Moreover, most work in robust MDP considers an offline setting or assumes access to a generator but we consider an online setting. Typically, robust MDP approaches assume a fixed ambiguity set size to calculate the optimal policy in the worst-case scenario. In contrast,

we aim to optimize performance under the true model and gradually shrink the ambiguity set as data accumulates. This fundamental difference in goals sets us apart from existing methods that consider exogenous robustness, where the environment may be perturbed and the goal is to optimize for the worst-case scenario. In those works, the size of the uncertainty set is known, but the nominal transition probability is unknown. We, however, assume the existence of a true nominal system and design an adaptive robust algorithm that remains robust when data points are limited – what we term *endogenous* robustness. Our ambiguity set shrinks as we collect more data. Our work is closest to the online robust MDP work by (Dong et al. 2022). However, our work has an essential difference: they consider exogenous robustness and fix the size of the ambiguity set, whereas we aim to optimize the model performance under the true kernel. It is important to emphasize that there are no existing sublinear regret results for online robust MDP problems, and achieving such results is notoriously difficult in general. In this work, by leveraging known prototypes of the underlying transition probability, we are able to provide sublinear regret bounds. This demonstrates the significant benefit of incorporating useful prior information about the underlying MDP model. Our approach could offer valuable insights for future work on model-based MDPs, particularly in scenarios where structural information is available or can be inferred. The main contributions of our work are as follows:

1. We propose a novel algorithm for learning robust policies in MDPs with multiple transition dynamic prototypes in an online setting (RPO-AAS). We show that our algorithm achieves sublinear regret with respect to the optimal policy for the true model and introduce an early stopping mechanism that allows our algorithm to converge to the true model more quickly with sufficient evidence.
2. We also propose a non-robust algorithm (NRPO-NPC) and analyze the technical performance guarantees. This algorithm does not calculate the robust optimal policy, but selects the prototype that is closest to the empirical distribution and runs the optimal policy corresponding to this prototype. Interestingly, we show that introducing robustness in the algorithm does not sacrifice efficiency.
3. Through numerical experiments, we demonstrate the effectiveness of our approach compared to existing methods, showing improved performance particularly in the early stage with limited data.

2 Related Work

Recent research has explored MDPs with parameter ambiguity using multiple models. (Steimle, Kaufman, and Denton 2021) and (Buchholz and Scheftelowitsch 2019) consider finding a policy that maximizes a weighted performance across multiple models of MDPs. They proved NP-hardness of the problem and developed exact and approximate solution methods. (Ahmed et al. 2017) explore sampling rewards and transition probabilities to generate a finite set of MDPs and find a policy to minimize the maximum regret over the set of MDPs. Our work differs from these approaches in two key aspects. Firstly, we consider an online

setting, whereas previous works focused on offline setting. Secondly, our goal is to identify the true model and optimize its performance during exploration while guaranteeing robustness, rather than optimizing weighted performance for given weights or worst-case regret across all models.

The problem of regret minimization in MDPs with a fixed reward function has been studied extensively since (Burnetas and Katehakis 1997) and (Auer and Ortner 2006). Provably efficient learning algorithms fall into two main categories: The first applies *optimism in the face of uncertainty* principle (Kearns and Singh 2002; Brafman and Tenenholz 2002; Azar, Osband, and Munos 2017) while the second utilizes posterior sampling reinforcement learning (Osband, Russo, and Van Roy 2013; Osband and Van Roy 2017). (Agrawal and Jia 2017) combine these approaches, leveraging both the optimistic principle and posterior sampling to achieve a regret bound for weakly communicating MDPs. Currently the best regret bound for finite MDP is $\tilde{O}(\sqrt{H|\mathcal{S}||\mathcal{A}|T} + H^2|\mathcal{S}|^2|\mathcal{A}| + H\sqrt{T})$ from the UCBVI algorithm, where \mathcal{S} is the finite space of states, \mathcal{A} is the set of finite actions and H is the number of horizons (Azar, Osband, and Munos 2017). Despite these advancements, to our knowledge no existing work considers robust algorithms in MDPs with sublinear regret compared to the optimal reward.

Robust MDPs consider the transition kernels that take values from an uncertainty set and learn an optimal robust policy that maximizes the worst-case value function. Most work in Robust MDP assumes that the the uncertainty set is known (Iyengar 2005; Nilim and El Ghaoui 2005; Xu and Mannor 2010). Recently some work consider the robust optimal policy when the uncertainty set is not exactly known, or say reinforcement learning. Some work assumes that there is a generative model (Panaganti and Kalathil 2022; Yang, Zhang, and Zhang 2022) or assumes an offline dataset is present (Zhou et al. 2021; Qi and Liao 2020; Kallus et al. 2022; Ma et al. 2022). To our knowledge, only (Dong et al. 2022) considers the robust policy learning in online setting. They propose algorithms that achieve a regret of $\tilde{O}(|\mathcal{S}||\mathcal{A}|^2H^2)$ under s -rectangular uncertainty set. However, these work have a different goal from our work. As discussed before, they consider the *exogeneous* robustness, while we consider *endogeneous* robustness. They consider an ambiguity set with fixed size while the radius of our algorithm is shrinking we when collect more data.

Another line of research characterizes the uncertainty through adversarial MDP formulations, where the environment parameters can be adversarially chosen. Most studies focus on the setting where only the reward function can be corrupted, while transition dynamics of the MDP remain fixed but potentially unknown (Neu et al. 2010; Cai et al. 2020; Jin et al. 2020; Rosenberg and Mansour 2019; Jin and Luo 2020; Cai et al. 2020). (Neu et al. 2010) first proposes the online loop-free setting and show a regret of $\tilde{O}(L^2\sqrt{T}|\mathcal{A}|/\alpha)$ under some assumptions, where L is the length of the longest path in the graph, T is the number of episodes, and α is a probability parameter in the assumption. Some work investigates settings where adversaries can corrupt transition metrics. (Lykouris et al. 2021) consider the

setting that the transition is only allowed to be adversarially chosen for C out of the T total episodes and establish a regret of $\tilde{O}(C^2 + \sqrt{T})$. Our prototype elimination approach shares similarities with arm-elimination methods in multi-armed bandit problems (Even-Dar et al. 2006; Audibert and Bubeck 2010), but handles the additional complexity of state transitions rather than simple rewards.

3 Problem Formulation and Preliminaries

3.1 Problem Formulation

We consider a Markov Decision Process defined by a tuple $(\mathcal{S}, \mathcal{A}, P_0, r)$, where \mathcal{S} is the finite state space and \mathcal{A} is the finite action space, $P_0 : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ is the transition kernel, $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the reward function. More specifically, we use $P_0(s, a)$ and $r(s, a)$ to denote the probability distribution of the next state and immediate reward when taking action a at state s . Let $P_0(s'|s, a)$ denote the probability of arriving at state s' when choosing action a at state s . Moreover, we assume the reward $r(s, a)$ is deterministic, and without loss of generality, $r(s, a)$ belongs to $[0, 1]$. Extending the algorithms to the setting with unknown reward does not add significant difficulty.

Loop-Free MDP In this work, we consider an episodic MDP with finite horizons. We assume the MDP has a *loop-free* structure: The state space can be decomposed into $L+1$ non-intersecting layers $\mathcal{S}_0, \dots, \mathcal{S}_L$ such that $\mathcal{S} = \cup_{l=0}^L \mathcal{S}_l$, $\mathcal{S}_i \cap \mathcal{S}_j = \emptyset$ for $i \neq j$. Moreover, the first and the last layers are singletons, i.e., $\mathcal{S}_0 = \{s_0\}$, $\mathcal{S}_L = \{s_L\}$. Let $\mathcal{L}(s)$ denote the layer of state s . The loop-free structure means the transitions are only possible between consecutive layers. These assumptions are not necessary but are commonly adopted in literature, intended to simplify notation and analysis, and can be modified for a more general setup (Rosenberg and Mansour 2019; Jin et al. 2020).

Transition Prototypes In this work, we aim to illustrate the benefit of utilizing prior information about the transition probabilities. Specifically, we consider prototypes that are known to the decision-maker, each of which may correspond to an underlying model or mechanism that is driving the transition of the states. We assume that for each layer l , there are K_l prototypes of the transition kernel in the candidate set, denoted as $\{1, 2, \dots, K_l\}$ and collectively referred to as \mathcal{K}_l . For any layer l , the transition probability at state s and action a defined by prototype $k \in \mathcal{K}_l$ is $P^k(s, a)$. The true transition kernel of each layer l , denoted as k_l^* , must be one of the prototypes, meaning that $\otimes_{s \in \mathcal{S}_l, a \in \mathcal{A}} P_0(s, a) = \otimes_{s \in \mathcal{S}_l, a \in \mathcal{A}} P^{k_l^*}(s, a)$.

In the algorithm, we update the candidate set of prototypes gradually, and we let $\mathcal{K}_{l,t}$ denote the set of prototypes in episode t . We remove the prototypes that are unlikely to be true as we collect more data. For the prototypes, we make the following structural assumption, which essentially states that if the gap between some kernels at a particular state s in the layer and action a is small, then the difference at other states in this layer cannot be too large.

Assumption 1. For any layer $l = 0, \dots, L$, any state $s \in \mathcal{S}_l$, action $a \in \mathcal{A}$, and any prototype $k \in \mathcal{K}_l$, if for some

constant $u \in \mathbb{R}$, the l_1 -norm $\|P^k(s, a) - P^0(s, a)\|_1 \leq u$, then there exists a constant $\gamma \in \mathbb{R}$ such that $\|P^k(s', a') - P^0(s', a')\|_1 \leq \gamma u$ for any other $s' \in \mathcal{S}_l, a' \in \mathcal{A}$.

Assumption 1 reflects that states within the same layer often share similar transition patterns, which is common in practice. The constant γ quantifies the variability of transition probability differences across state-action pairs, while u represents the magnitude of these differences for a reference state-action pair. Importantly, our theoretical results depend solely on γ , not on the absolute differences captured by u . This formulation provides flexibility in accommodating various MDP structures while maintaining analytical tractability. While this assumption helps establish theoretical guarantees, our numerical experiments in Section 7.2 show that the algorithm maintains good performance even with random prototypes where this assumption may not hold.

In this paper, we use $\|\cdot\|_1$ to denote the l_1 -norm between two transition probability vectors. For any two transition kernels at state s and action a , $P_0(s, a)$ and $P_1(s, a)$, we define $\|P_0(s, a), P_1(s, a)\|_1 = \sum_{s' \in \mathcal{S}} |P_0(s'|s, a) - P_1(s'|s, a)|$.

In each episode t , let π_t denote the policy, which is a mapping from the state space \mathcal{S} to action space \mathcal{A} . Given the transition kernel P_0 and policy π_t , the expected reward in episode t is:

$$\mathbb{E}\left[\sum_{l=0}^{L-1} r(s_l, \pi_t(s_l)) | P_0, \pi_t\right],$$

where s_l is the state visited in layer l and episode t and $\pi_t(s_l)$ is the corresponding action. Then, the total expected reward of the learner over T episodes is:

$$R((\pi_t)_{t \in [T]}, P_0) = \sum_{t=1}^T \mathbb{E}\left[\sum_{l=0}^{L-1} r(s_l, \pi_t(s_l)) | P_0, \pi_t\right].$$

For a stationary policy π , with a slight abuse of notation, the total expected reward is given by

$$R(\pi, P_0) = \sum_{t=1}^T \mathbb{E}\left[\sum_{l=0}^{L-1} r(s_l, \pi(s_l)) | P_0, \pi\right].$$

Therefore, the regret can be defined as

$$Reg = R(\pi^*, P_0) - R((\pi_t)_{t \in [T]}, P_0),$$

where $\pi^* \in \arg \max_{\pi} \mathbb{E}[\sum_{l=0}^{L-1} r(s_l, \pi(s_l))]$.

Our regret definition diverges from that in the robust MDP literature (Dong et al. 2022; Zhou et al. 2021) which optimizes worst-case reward over an ambiguity set, with regret measured as the gap between worst-case rewards of the algorithm's policy and the optimal worst-case robust policy. In contrast, we optimize reward under the true transition kernel, aligning with the online MDP framework (Neu et al. 2010).

3.2 Preliminaries

Occupancy measures. We now reformulate the learner's problem using the concept of occupancy measures. We introduce occupancy measures for the purpose of analysis, which has been widely used in the analysis for loop-free

MDP (Jin et al. 2020; Rosenberg and Mansour 2019). Given a policy π and transition kernel P , for any state $s \in \mathcal{S}_l$, $s' \in \mathcal{S}_{l+1}$, the occupancy measure $q^{P,\pi}$ is defined as:

$$q^{P,\pi}(s, a, s') = \mathbb{P}[s_{\mathcal{L}(s)} = s, \pi(s) = a, s_{\mathcal{L}(s)+1} = s' | P, \pi].$$

An occupancy measure satisfies the following two properties and these two properties suffice to define any function $q : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ to be an occupancy measure. (1) The learner traverses every layer in each episode due to the loop-free structure, i.e., for every $l = 0, \dots, L - 1$,

$$\sum_{s \in \mathcal{S}_l} \sum_{a \in \mathcal{A}} \sum_{s' \in \mathcal{S}_{l+1}} q(s, a, s') = 1.$$

(2) The probability of entering a state from the previous layer equals the probability of leaving it. Thus, for every $l = 1, \dots, L - 1$ and $s \in \mathcal{S}_l$,

$$\sum_{s' \in \mathcal{S}_{l+1}} \sum_{a \in \mathcal{A}} q(s, a, s') = \sum_{s' \in \mathcal{S}_{l-1}} \sum_{a \in \mathcal{A}} q(s', a, s).$$

Given an occupancy measure q , the transition function P^q and the policy π^q can be induced as follows:

$$P^q(s' | s, a) = \frac{q(s, a, s')}{\sum_{y \in \mathcal{S}_{\mathcal{L}(s)+1}} q(s, a, y)},$$

$$\pi^q(a | s) = \frac{\sum_{s' \in \mathcal{S}_{\mathcal{L}(s)+1}} q(s, a, s')}{\sum_{b \in \mathcal{A}} \sum_{s' \in \mathcal{S}_{\mathcal{L}(s)+1}} q(s, b, s')}.$$

Then the problem of policy learning can be transformed to learning an occupancy measure $q_t \in \Delta(P_0)$ in each episode t , where $\Delta(P_0)$ is the set of all occupancy measures of an MDP with transition kernel P_0 . With the definition of the occupancy measure, we redefine the expected reward and regret. The total expected reward of the learner is

$$\begin{aligned} R((\pi_t)_{t \in [T]}, P_0) &= \sum_{t=1}^T \mathbb{E} \left[\sum_{l=0}^{L-1} r(s_l, \pi_t(s_l)) | P_0, \pi_t \right] \\ &= \sum_{t=1}^T \langle q^{P_0, \pi_t}, r \rangle \end{aligned}$$

Let $q^* \in \operatorname{argmax}_{q \in \Delta(P_0)} \sum_{t=1}^T \langle q^{P_0, \pi}, r \rangle = q^{P_0, \pi^*}$ denote the occupancy measure corresponding to the optimal policy π^* under P_0 , the regret can be defined as

$$\text{Reg} = \max_{\pi} R(\pi, P_0) - R((\pi_t)_{t \in [T]}, P_0) \quad (1)$$

$$= \sum_{t=1}^T \langle q^* - q^{P_0, \pi_t}, r \rangle. \quad (2)$$

4 The RPO-AAS Algorithm

In this section, we introduce how we update the ambiguity set and calculate the robust optimal policy with respect to the ambiguity set in each episode. The algorithm initializes the policy π to an arbitrary policy π_0 (e.g., a uniform policy) and sets the number of samples $N_1(s, a)$ to zero for each state-action pair (s, a) . In each episode, the following steps

Algorithm 1: Robust Policy Optimization with Adaptive Ambiguity Set (RPO-AAS)

- 1: Initialize: $\pi \leftarrow \pi_0$, number of samples $N_1(s, a) = 0$ for each $s \in \mathcal{S}, a \in \mathcal{A}$
 - 2: **for** $t = 1, \dots, T$ **do**
 - 3: **for** $l = 1, \dots, L$ **do**
 - 4: $s_{tl}, a_{tl} = \operatorname{arg max}_{s \in \mathcal{S}_l, a \in \mathcal{A}} N_t(s, a)$
 - 5: Update the set of candidate prototypes:
 - 6: $\mathcal{K}_{l,t} = \{k \in \mathcal{K}_{l,t-1} : \|P^k(s_{tl}, a_{tl}) - \hat{P}_t(s_{tl}, a_{tl})\|_1 \leq \sqrt{\frac{4|\mathcal{S}_{l+1}| \ln \frac{3LT}{\delta}}{N_t(s_{tl}, a_{tl})}}\}$
 - 7: **end for**
 - 8: Update ambiguity set:
 $\mathcal{U}_t = \bigotimes_{s \in \mathcal{S}, a \in \mathcal{A}} \bigotimes_{k \in \mathcal{K}_{\mathcal{L}(s), t}} P^k(s, a)$
 - 9: Calculate optimal robust policy:
 $\pi_t = \operatorname{arg max}_{\pi} \min_{P \in \mathcal{U}_t} R(\pi, P)$
 - 10: Execute policy π_t for L steps and obtain trajectory s_l, a_l for $l = 1, \dots, L - 1$
 - 11: Update $N_t(s, a)$ and $\hat{P}_t(s, a)$ for all s, a
 - 12: **end for**
-

are performed: First, for each layer $l = 1, \dots, L$, we identify the state-action pair (s_{tl}, a_{tl}) with the maximum number of samples in that layer. Next, we update the set of prototypes $\mathcal{K}_{l,t}$ by eliminating prototypes whose transition probabilities significantly deviate from the empirical transition distribution $\hat{P}_t(s, a)$ for the state-action pair (s_{tl}, a_{tl}) . This update is crucial, as it relies on the state-action pair with the most occurrences, ensuring faster convergence of the empirical distribution to the true distribution. Subsequently, we update the ambiguity set \mathcal{U}_t as the Cartesian product of the ambiguity sets for each state-action pair, where each set comprises the transition probabilities of the remaining prototypes in the corresponding layer. We then calculate the robust optimal policy π_t by maximizing the worst-case value function over the ambiguity set \mathcal{U}_t . Since our ambiguity set satisfies the (s,a)-rectangular property, the optimal policy can be calculated using backward induction. The backward induction and ambiguity set update step takes $O(|\mathcal{S}||\mathcal{A}| + \sum_{l=1}^L |\mathcal{K}_l|)$ time, which is efficient. Moreover, the key advantage of this ambiguity set construction is its high probability of including the true transition kernel as in the following lemma.

Lemma 1. *For the ambiguity set updated as in Algorithm 1, the true transition kernel lies in the ambiguity set \mathcal{U}_t , i.e., $P_0 \in \mathcal{U}_t$ for all $t \in [T]$ with probability at least $1 - \delta$.*

We would like to point out that, this robust setting by considering the ambiguity set and solving for the worst-case value function over it allows one to have a worst-case performance bound, as stated in Proposition 1. To be more specific, with the high-probability ambiguity set, we have that in each episode t , policy π_t has the best worst-case performance and the performance of policy π_t is lower bounded by the optimal objective value of the robust MDP. As we will see later, the non-robust algorithm lacks this robustness and could have poor performance, especially when we don't have enough data at the beginning.

Proposition 1. In episode t , $\min_{P \in \mathcal{U}_t} R(\pi_t, P) \geq \min_{P \in \mathcal{U}_t} R(\pi, P)$ for all policy π . Moreover, with probability at least $1 - \delta$, $\min_{P \in \mathcal{U}_t} R(\pi_t, P)$ provides a lower bound for $R(\pi_t, P_0)$ with probability at least $1 - \delta$.

The proof uses Hoeffding’s inequality to bound the difference between the true and empirical transition probabilities. Due to space limitations, proofs for all results in this paper are provided in the appendix. This proposition implies that in each episode t , π_t has the best worst-case performance, and its actual performance is lower-bounded by the optimal objective value of the robust MDP. In contrast, a non-robust algorithm lacks this guarantee and may perform poorly, especially with limited data at the beginning. While the robust policy has its own advantages, the question remains whether this robust policy has a good performance under the true transition kernel P_0 . In the following section, we prove the theoretical guarantee of the RPO-AAS algorithm under P_0 .

5 Theoretical Results

In this section, we first establish the regret bound, and then show the finite sample guarantee and the convergence result.

5.1 Analysis of Regret

To bound the regret, we begin by decomposing (1) as follows:

$$\text{Reg} = \sum_{t=1}^T \langle q^* - q_t, r \rangle = \sum_{t=1}^T \langle q^* - \hat{q}_t, r \rangle + \langle \hat{q}_t - q_t, r \rangle,$$

where $q_t = q^{P_0, \pi_t}$, $\hat{q}_t = q^{P_t, \pi_t}$ and π_t, P_t is the optimal solution of the robust optimization problem $\max_{\pi} \min_{P \in \mathcal{U}_t} R(\pi, P)$. The high-level idea of our proof of regret has three main steps. First, we upper bound the regret by the total reward difference between the true transition kernel P_0 and the kernel given by P_t under the optimal policy π^* (Lemma 2). We then bound this reward difference in two steps. We first establish a bound on the one-norm difference between P_0 and P_t (Lemma 3), and then bound the difference of total reward (Lemma 5). We begin with Lemma 2, which provides an upper bound on the regret in terms of the total reward difference between the true transition kernel P_0 and the kernel given by robust optimization P_t under the optimal policy π^* .

Lemma 2. With probability at least $1 - \delta$, $\sum_{t=1}^T \langle q^* - \hat{q}_t, r \rangle + \langle \hat{q}_t - q_t, r \rangle \leq \sum_{t=1}^T \|q_t^{P_0, \pi^*} - q_t^{P_t, \pi^*}\|_1$.

Here, $\|q^{P_t, \pi^*} - q^{P_0, \pi^*}\|_1 = \sum_{s, a, s'} |q^{P_t, \pi^*}(s, a, s') - q^{P_0, \pi^*}(s, a, s')|$. So it remains to bound $\sum_{t=1}^T \|q_t^{P_0, \pi^*} - q_t^{P_t, \pi^*}\|_1$. Based on the result from (Rosenberg and Mansour 2019), we bound $\sum_{t=1}^T \|q^{P_0, \pi^*} - q^{P_t, \pi^*}\|_1$ as follows.

Lemma 3. For any policy π and any $P_t \in \mathcal{U}_t$, with probability at least $1 - \delta$, the following holds:

$$\begin{aligned} & \sum_{t=1}^T \|q_t^{P_0, \pi} - q_t^{P_t, \pi}\|_1 \\ & \leq 2 \sum_{t=1}^T \sum_{l=1}^L \sum_{m=0}^{l-1} \sum_{s_m \in \mathcal{S}_m} \sum_{a_m \in \mathcal{A}} q^{P_0, \pi}(s_m, a_m) \xi_t(s_m, a_m), \end{aligned}$$

where $\xi_t(s, a) = \|P_t(\cdot | s, a), P_0(\cdot | s, a)\|_1$.

Thus, to bound the right-hand side in the lemma above, the key is to bound $\xi_t(s, a)$.

Lemma 4. Suppose $P_0 \in \mathcal{U}_t$. Then for any $s \in \mathcal{S}$, $a \in \mathcal{A}$, $t \in [T]$, and for all $k \in K_{t, \mathcal{L}(s)}$, we have:

$$\|P_0(s, a), P^k(s, a)\|_1 \leq \sqrt{\frac{4|\mathcal{S}_{\mathcal{L}(s)+1}||\mathcal{A}| \ln \frac{3LT}{\delta}}{t}} \quad (3)$$

With the established bound for ξ_t , we prove the following bound for the right-hand side of Lemma 3.

Lemma 5. With probability at least $1 - \delta$, the following holds:

$$\begin{aligned} & \sum_{t=1}^T \sum_{l=1}^L \sum_{m=0}^{l-1} \sum_{s_m \in \mathcal{S}_m} \sum_{a_m \in \mathcal{A}} q^{P_0, \pi}(s_m, a_m) \xi_t(s_m, a_m) \\ & \leq L^2 \gamma \sqrt{4T|\mathcal{S}||\mathcal{A}| \ln \frac{3LT}{\delta}}. \end{aligned}$$

By combining Lemma 2, Lemma 3 and 5, we have the following regret bound:

Theorem 1. With probability at least $1 - \delta$, the RPO-AAS algorithm has the following regret bound:

$$\text{Reg} \leq L^2 \gamma \sqrt{4T|\mathcal{S}||\mathcal{A}| \ln \frac{3LT}{\delta}}.$$

It’s worth noting that the state-of-the-art algorithm for general online MDPs achieves a regret bound of $\tilde{O}(\sqrt{H|\mathcal{S}||\mathcal{A}|T + H^2 S^2 |\mathcal{A}| + H\sqrt{T}})$, where H is the number of horizons (Azar, Osband, and Munos 2017). Our regret bound maintains the same dependence on $|\mathcal{S}|$, $|\mathcal{A}|$, and T . This demonstrates that, given structural information, our robust algorithm matches the efficiency of non-robust state-of-the-art approaches. However, it’s important to note that designing efficient robust RL algorithms without structural information remains an open problem in the field.

5.2 Finite-Sample Guarantee and Convergence

In addition to the regret bound, we establish that the policy obtained by the proposed algorithm has a finite-sample performance guarantee and converges to the optimal policy.

Theorem 2 (Finite-sample guarantee). Let $v^\pi(s_0)$ denote the value function at state s_0 under policy π under the true transition kernel. For any $\epsilon > 0$, when $t \geq \frac{4L^4 \gamma^2 |\mathcal{S}||\mathcal{A}| \ln \frac{3LT}{\delta}}{\epsilon^2}$, with probability at least $1 - \delta$, $v^{\pi^*}(s_0) - v^{\pi_t}(s_0) \leq \epsilon$.

This theorem states that after a sufficient number of episodes t , the value function of our algorithm’s policy π_t at the initial state s_0 is within ϵ of the optimal policy π^* ’s value function, with high probability. The required number of episodes is inversely proportional to ϵ^2 . This dependency on ϵ^2 is typical in many MDP problems (Panaganti and Kalathil 2022). We next show that our algorithm can actually identify the true prototype after a finite number of episodes, leading to the optimal policy.

Theorem 3 (Prototype Ambiguity Set Convergence). *Let $h = \min_{s \in \mathcal{S}, a \in \mathcal{A}, k \in [K]} \|P^k(\cdot|s, a), P_0(\cdot|s, a)\|_1$, then when $t \geq \frac{8|\mathcal{S}|^2|\mathcal{A}|\ln \frac{3LT}{\delta}}{h}$, the candidate set of prototypes only include the true prototypes, i.e., $K_{tl} = \{k_l^*\}$, thus $\pi_t = \pi^*$.*

This theorem establishes a finite-time guarantee for our algorithm’s convergence to the true prototype and, consequently, the optimal policy. The result provides a principled stopping criterion, potentially improving the algorithm’s practical efficiency.

6 Extend to Non-robust Algorithm: Selecting the Best Candidate

We propose another algorithm that selects the transition kernel that is nearest to the empirical distribution in each episode, referred to as non-robust policy optimization with nearest prototype-candidate (NRPO-NPC). Then in each episode, we run the optimal policy corresponding for the chosen transition kernel. We demonstrate that this approach provides the same theoretical performance guarantees for regret, convergence, and finite sample guarantees as the robust algorithm. However, it lacks the robustness guarantee. To establish the theoretical results, we first decompose the regret at each episode as follows:

$$\sum_{t=1}^T \langle q^* - q_t, r \rangle = (q^{\pi^*, P_0} - q^{\pi^*, P_t}) + (q^{\pi^*, P_t} - q^{\pi_t, P_t}) + (q^{\pi_t, P_t} - q^{\pi_t, P_0}) \quad (4)$$

The second term, $q^{\pi^*, P_t} - q^{\pi_t, P_t} \leq 0$, since π_t is the optimal policy for transition kernel P_t . Similar to the proof for Theorem 1, we can bound the first term and the third term as long as we can bound the distance between P_0 and P_t , which is shown in the following lemma.

Lemma 6. *For each layer l , let $s_{tl}, a_{tl} = \arg \max_{s \in \mathcal{S}_l, a \in \mathcal{A}} N_t(s, a)$ denote the (s, a) pair with the maximum number of samples in the layer. Let $k_t = \arg \min_{k \in K_{l,t}} \|P^k(s_{tl}, a_{tl}) - \hat{P}_t(s_{tl}, a_{tl})\|_1$. Then for any $s \in \mathcal{S}, a \in \mathcal{A}, t \in [T]$, we have:*

$$\|P_0(\cdot|s, a), P^{k_t}(\cdot|s, a)\|_1 \leq \sqrt{\frac{4|\mathcal{S}_{\mathcal{L}(s)+1}|\mathcal{A}|\ln \frac{3LT}{\delta}}{t}} \quad (5)$$

7 Numerical Experiments

In the numerical experiments, we compare the performance of our proposed robust algorithm with the UCBVI algorithm (Azar, Osband, and Munos 2017), and the two benchmark algorithms we propose that considers the prototype information. We will provide more details later.

We consider a GridWorld experiment of size 5×4 , which is a widely used reinforcement setting from (Sutton and Barto 1998). In each episode, the learner starts from the lower left corner and aims to the upper right corner. Let (x_1, x_2) denote the coordinate, where $x_1 \in \{0, 1, 2, 3, 4\}$ is the coordinate of the horizontal axis and $x_2 \in \{0, 1, 2, 3\}$ is the vertical axis coordinate. The learner collects rewards

at some states, which we call reward states. We set the reward states to be $(2, 2)$, $(1, 1)$ and $(1, 2)$ and the rewards are 3, 5 and 1, respectively. At each state s and a , the learner can either move up ($a = 0$) or right ($a = 1$), with a success probability $z(s, a)$, and the learner goes to the opposite direction with probability $1 - z(s, a)$. $z(s, a)$ is unknown. The learner’s goal is to maximize the total collected rewards. If a learner reaches a boundary, she can only move inward. This problem is an episodic loop-free MDP, where each episode consists of $L = 8$ layers. The number of states is $|\mathcal{S}| = 20$ and the number of actions is $|\mathcal{A}| = 2$.

Prototype configuration. In each instance, we generate K prototypes. We set $K = 4$ and $K = 10$, representing scenarios with few and many prototypes, respectively. For each prototype, we generate random $z_k(s, a)$ from a uniform distribution between 0 and 1. For simplicity, we generate different success probabilities only for different states, meaning $z_k(s, 0)$ remains the same for all states, as does $z_k(s, 1)$. We consider two types of prototype sets: The first set of prototype satisfies our assumption on the structure of transition prototypes (Assumption 1). Specifically, for any s and a , we let $|z_{k_1}(s, a) - z_{k_2}(s, a)|$ be fixed for any kernel k_1 and k_2 . We call this setting *fixed-gap* prototypes. The second set does not satisfy this assumption. In this setting, we generate $z_k(s, 0)$ and $z_k(s, 1)$ for all prototypes randomly.

Algorithms. We compare four algorithms: (1) our robust algorithm (RPO-AAS), (2) UCBVI algorithm, (3) the non-robust nearest prototype-candidate algorithm (NRPO-NPC), and (4) its variant, NRPO-NPC-2. The latter is a heuristic that selects the prototype with the smallest 1-norm distance to the empirical transition probabilities across all states and actions in the layer. Details are provided in the Appendix.

Experiment Environment. We conduct the numerical experiment using *rlberry*, a Python library for reinforcement learning (Domingues et al. 2021). For each setting, we run 100 simulations. In each simulation, we record the average expected rewards in each episode. We then take the average of these simulations. The expected episode reward is the expectation of the total reward under the policy in episode t .

7.1 Structured Prototypes Setting

In the fixed-gap setting with $K = 4$, Figure 1 shows that NRPO-NPC, NRPO-NPC-2, and RO perform significantly faster than the UCBVI algorithm. This indicates that our proposed algorithms can leverage the prototype information effectively, resulting in better performance. NRPO-NPC-2 converges to the optimal policy fastest, although it lacks theoretical guarantees. When $K = 10$, the performance of NRPO-NPC and RO surpasses NRPO-NPC-2 and UCBVI (Figure 2). Notably, in both cases, RO has better performance at the beginning, showcasing the advantage of considering robustness.

7.2 Random Prototypes Setting

We start from $K = 4$ prototypes. Figure 3 shows the performance of the algorithms. In this setting, NRPO-NPC-2 couldn’t converge to the optimal policy. RO yields bet-




Figure 1: Average Expected Episode Rewards of different algorithms with Fixed-gap Prototypes when $K = 4$.




Figure 3: Average Expected Episode Rewards of different algorithms with 4 Random Prototypes.




Figure 2: Average Expected Episode Rewards of different algorithms with Fixed-gap Prototypes when $K = 10$.




Figure 4: Average Expected Episode Rewards of different algorithms with 10 Random Prototypes.

ter policies than UCBVI and NRPO-NPC in the first 2,000 episodes. Moreover, the policy given by RO has lower fluctuations than NRPO-NPC and UCBVI. NRPO-NPC outperforms UCBVI initially but shows greater variance and converges to the optimal solution more slowly than UCBVI. When we increase the number of prototypes to 10, NRPO-NPC, NRPO-NPC-2, and RO continue to outperform UCBVI during the first 400 episodes. RO maintains the lowest variance, indicating that it yields the most stable policy. However, UCBVI converges to the optimal policy more rapidly than RO and NRPO-NPC in many cases, resulting in slightly superior performance after 2,000 episodes.

This observation suggests that as the number of prototypes increases, the benefits of incorporating prototype information diminish. This is logical, as in the limit of infinite prototypes, the algorithm would gain no advantage from prototype information. From Theorem 3, more prototypes potentially reduce h and thus slow convergence, while too few prototypes may fail to include the transition kernel. Therefore, the number of prototypes K presents a practical trade-off. Nevertheless, the RO algorithm maintains its robustness even in this many-prototype setting.

8 Conclusion

In this work, we introduced a novel approach for online MDPs with transition prototypes. Our robust adaptive algorithm efficiently identifies the true transition kernel while guaranteeing performance through robust policies. Theoretical analysis shows the algorithm achieves sublinear regret, provides finite-sample guarantees, and converges to the optimal policy in finite time. Numerical experiments demonstrate its practical advantages, particularly in early learning stages and with structured prototypes. We also extended our analysis to a non-robust algorithm, highlighting the value of prototype information. This work shows the potential of the combination of structural information and robust optimization in reinforcement learning. Future work could explore extensions to more complex MDP settings and investigate robustness-optimality trade-offs in various applications.

Acknowledgments

This research is partially supported by the ITC Mainland-Hong Kong Joint Funding Scheme (No. MHP/192/23).

References

Agrawal, S.; and Jia, R. 2017. Posterior sampling for reinforcement learning: worst-case regret bounds. *arXiv*

preprint [arXiv:1705.07041](https://arxiv.org/abs/1705.07041).

- Ahmed, A.; Varakantham, P.; Lowalekar, M.; Adulyasak, Y.; and Jaillet, P. 2017. Sampling based approaches for minimizing regret in uncertain Markov decision processes (MDPs). *Journal of Artificial Intelligence Research*, 59: 229–264.
- Audibert, J.-Y.; and Bubeck, S. 2010. Best arm identification in multi-armed bandits. In *COLT-23th Conference on learning theory-2010*, 13–p.
- Auer, P.; and Ortner, R. 2006. Logarithmic online regret bounds for undiscounted reinforcement learning. *Advances in neural information processing systems*, 19.
- Azar, M. G.; Osband, I.; and Munos, R. 2017. Minimax regret bounds for reinforcement learning. In *International conference on machine learning*, 263–272. PMLR.
- Brafman, R. I.; and Tennenholtz, M. 2002. R-max—a general polynomial time algorithm for near-optimal reinforcement learning. *Journal of Machine Learning Research*, 3(Oct): 213–231.
- Buchholz, P.; and Scheftelowitsch, D. 2019. Computation of weighted sums of rewards for concurrent MDPs. *Mathematical Methods of Operations Research*, 89: 1–42.
- Burnetas, A. N.; and Katehakis, M. N. 1997. Optimal adaptive policies for Markov decision processes. *Mathematics of Operations Research*, 22(1): 222–255.
- Cai, Q.; Yang, Z.; Jin, C.; and Wang, Z. 2020. Provably efficient exploration in policy optimization. In *International Conference on Machine Learning*, 1283–1294. PMLR.
- Chatterjee, K.; Chmélík, M.; Karkhanis, D.; Novotný, P.; and Royer, A. 2020. Multiple-environment markov decision processes: Efficient analysis and applications. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 30, 48–56.
- Domingues, O. D.; Flet-Berliac, Y.; Leurent, E.; Ménard, P.; Shang, X.; and Valko, M. 2021. `rlberry` - A Reinforcement Learning Library for Research and Education.
- Dong, J.; Li, J.; Wang, B.; and Zhang, J. 2022. Online policy optimization for robust MDP. *arXiv preprint arXiv:2209.13841*.
- Even-Dar, E.; Mannor, S.; Mansour, Y.; and Mahadevan, S. 2006. Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *Journal of machine learning research*, 7(6).
- Iyengar, G. N. 2005. Robust Dynamic Programming. *Mathematics of Operations Research*, 30(2): 257–280.
- Jin, C.; Jin, T.; Luo, H.; Sra, S.; and Yu, T. 2020. Learning adversarial markov decision processes with bandit feedback and unknown transition. In *International Conference on Machine Learning*, 4860–4869. PMLR.
- Jin, T.; and Luo, H. 2020. Simultaneously learning stochastic and adversarial episodic mdps with known transition. *Advances in neural information processing systems*, 33: 16557–16566.
- Kallus, N.; Mao, X.; Wang, K.; and Zhou, Z. 2022. Doubly robust distributionally robust off-policy evaluation and learning. In *International Conference on Machine Learning*, 10598–10632. PMLR.
- Kearns, M.; and Singh, S. 2002. Near-optimal reinforcement learning in polynomial time. *Machine learning*, 49: 209–232.
- Lykouris, T.; Simchowitz, M.; Slivkins, A.; and Sun, W. 2021. Corruption-robust exploration in episodic reinforcement learning. In *Conference on Learning Theory*, 3242–3245. PMLR.
- Ma, X.; Liang, Z.; Blanchet, J.; Liu, M.; Xia, L.; Zhang, J.; Zhao, Q.; and Zhou, Z. 2022. Distributionally robust offline reinforcement learning with linear function approximation. *arXiv preprint arXiv:2209.06620*.
- Neu, G.; György, A.; Szepesvári, C.; et al. 2010. The Online Loop-free Stochastic Shortest-Path Problem. In *COLT*, volume 2010, 231–243. Citeseer.
- Nilim, A.; and El Ghaoui, L. 2005. Robust Control of Markov Decision Processes with Uncertain Transition Matrices. *Operations Research*, 53(5): 780–798.
- Osband, I.; Russo, D.; and Van Roy, B. 2013. (More) efficient reinforcement learning via posterior sampling. *Advances in Neural Information Processing Systems*, 26.
- Osband, I.; and Van Roy, B. 2017. Why is posterior sampling better than optimism for reinforcement learning? In *International conference on machine learning*, 2701–2710. PMLR.
- Panaganti, K.; and Kalathil, D. 2022. Sample complexity of robust reinforcement learning with a generative model. In *International Conference on Artificial Intelligence and Statistics*, 9582–9602. PMLR.
- Qi, Z.; and Liao, P. 2020. Robust batch policy learning in markov decision processes. *arXiv preprint arXiv:2011.04185*.
- Rosenberg, A.; and Mansour, Y. 2019. Online convex optimization in adversarial markov decision processes. In *International Conference on Machine Learning*, 5478–5486. PMLR.
- Steinle, L. N.; Kaufman, D. L.; and Denton, B. T. 2021. Multi-model Markov decision processes. *IIEE Transactions*, 1–16.
- Sutton, R. S.; and Barto, A. G. 1998. Reinforcement learning: an introduction MIT Press. *Cambridge, MA*, 22447: 10.
- Xu, H.; and Mannor, S. 2010. Distributionally Robust Markov Decision Processes. In Lafferty, J. D.; Williams, C. K. I.; Shawe-Taylor, J.; Zemel, R. S.; and Culotta, A., eds., *Advances in Neural Information Processing Systems* 23, 2505–2513. Curran Associates, Inc.
- Yang, W.; Zhang, L.; and Zhang, Z. 2022. Toward theoretical understandings of robust markov decision processes: Sample complexity and asymptotics. *The Annals of Statistics*, 50(6): 3223–3248.
- Zhou, Z.; Zhou, Z.; Bai, Q.; Qiu, L.; Blanchet, J.; and Glynn, P. 2021. Finite-sample regret bound for distributionally robust offline tabular reinforcement learning. In *International Conference on Artificial Intelligence and Statistics*, 3331–3339. PMLR.