

# When Is Partially Observable Reinforcement Learning Not Scary?

**Qinghua Liu**

*Princeton University*

QINGHUAL@PRINCETON.EDU

**Alan Chung**

*Princeton University*

ALAN.CHUNG@PRINCETON.EDU

**Csaba Szepesvári**

*DeepMind and University of Alberta*

SZEPESVA@UALBERTA.CA

**Chi Jin**

*Princeton University*

CHIJ@PRINCETON.EDU

**Editors:** Po-Ling Loh and Maxim Raginsky

## Abstract

Applications of Reinforcement Learning (RL), in which agents learn to make a sequence of decisions despite lacking complete information about the latent states of the controlled system, that is, they act under partial observability of the states, are ubiquitous. Partially observable RL can be notoriously difficult—well-known information-theoretic results show that learning partially observable Markov decision processes (POMDPs) requires an exponential number of samples in the worst case. Yet, this does not rule out the existence of large subclasses of POMDPs over which learning is tractable.

In this paper we identify such a subclass, which we call *weakly revealing* POMDPs. This family rules out the pathological instances of POMDPs where observations are uninformative to a degree that makes learning hard. We prove that for weakly revealing POMDPs, a simple algorithm combining *optimism* and *Maximum Likelihood Estimation* (MLE) is sufficient to guarantee polynomial sample complexity. To the best of our knowledge, this is the first provably sample-efficient result for learning from interactions in *overcomplete POMDPs*, where the number of latent states can be larger than the number of observations.

## 1. Introduction

A wide range of modern artificial intelligence challenges can be cast as Reinforcement Learning (RL) problems under *partial observability*, in which agents learn to make a sequence of decisions despite lacking complete information about the underlying state of system. For example, in robotics the agent has to cope with noisy sensors, occlusions, and unknown dynamics (Akkaya et al., 2019), while in imperfect information games the player makes only local observations (Vinyals et al., 2019; Brown and Sandholm, 2019). Further applications of partially observable RL include autonomous driving (Levinson et al., 2011), resource allocation (Bower and Gilbert, 2005), medical diagnostic systems (Hauskrecht and Fraser, 2000), recommendation (Li et al., 2010), business management (De Brito and Van Der Laan, 2009), etc. As such, learning and acting under partial observability has been an important topic in operation research, control, and machine learning.

Because of the non-Markovian nature of the observations, learning and planning in partially observable environments requires an agent to maintain *memory* and possibly reason about *beliefs* over the states, all while exploring to collect information about the environment. As such, partial observability can significantly complicate learning and planning under uncertainty. While practical

RL systems have succeeded in a set of partially observable problems including Poker (Brown and Sandholm, 2019), Starcraft (Vinyals et al., 2019) and certain robotic tasks (Cassandra et al., 1996), the theoretical understanding of learning to act in partially observable systems remains limited. Most existing results in RL theory focus on fully observable systems or, more generally, learning when the features of states are accessible and can faithfully represent value functions. As such, algorithms developed for this case need not to reason about what the latent state may be and in particular do not need to resort to using the observation histories. Thus, the resulting algorithms can be fundamentally limited and may not work beyond the narrow settings that they are designed for. Owing to the ubiquity of partially observable problems, addressing the theoretical challenges of partial observability is vital to closing the gap between the typical applications and the scope of available theoretical works.

This paper considers Partially Observable Markov Decision Process (POMDPs)—the standard model in reinforcement learning that captures the partial-information structure. Despite the existence of many efficient algorithms for learning MDPs in the fully observable settings, learning POMDPs is notoriously difficult in theory—well-known complexity-theoretic results show that learning and planning in partially observable environments is indeed statistically and computationally *intractable* in general (Papadimitriou and Tsitsiklis, 1987; Mundhenk et al., 2000; Vlassis et al., 2012; Mossel and Roch, 2005), even if in the favorable setting with a small number of states, actions, and observations. However, these complexity barriers are of a worst case nature, and they do not preclude efficient algorithms for learning rich sub-classes of POMDPs which could potentially cover interesting practical applications. This leaves an important question:

**Can we identify a rich sub-class of POMDPs that empowers sample-efficient RL?**

Prior efforts on sample-efficient learning of POMDPs focus either on special cases of POMDPs such as latent MDPs (Kwon et al., 2021b,a), or on general POMDPs but with restrictive assumptions. In particular, Azizzadenesheli et al. (2016); Guo et al. (2016) do not address strategic exploration—a core challenge in RL; Jin et al. (2020a) considers the exploration setting but only addresses undercomplete POMDPs, where the number of states must be no larger than the number of observations.

This paper answers the highlighted question above affirmatively. We identify a rich family of tractable POMDPs—*weakly revealing* POMDPs (see Section 3), which rule out the pathological instances whose observations contain no information to distinguish latent states. *Weakly revealing* POMDPs are very rich—it contains a majority of existing POMDPs classes which are known to be tractable (Azizzadenesheli et al., 2016; Guo et al., 2016; Jin et al., 2020a); it also handles overcomplete POMDPs where the number of latent states can be larger than the number of observations.

We further propose a new simple algorithm for learning POMDPs—*Optimistic Maximum Likelihood Estimation* (OMLE). As its name suggests, OMLE (read, Oh-Em-El-Eeh) combines optimism with classical maximum likelihood estimation (Wilks, 1938). In contrast to the algorithm of Jin et al. (2020a) which heavily exploit the undercomplete structure, our algorithm is generic, does not explicitly rely on any special structure, and can be used for any POMDPs. We prove that OMLE learns a near-optimal policy for any weakly revealing POMDP within a polynomial number of samples (Theorem 4 and 7). To the best of our knowledge, this is the first provably sample-efficient result for learning overcomplete POMDPs in settings where exploration is necessary. Our result also reasserts that optimism is a powerful tool to address exploration needs, regardless of whether states are observable. We complement our positive results with lower bounds showing that certain polynomial dependency on the problem parameters in our sample complexity is necessary.

Finally, we remark that our algorithm, as well as all existing algorithms for learning large classes of POMDPs, remains *computationally* inefficient. This is due to the inherent computational hardness of learning POMDPs: planning (i.e., computing the optimal policy *given* model parameters) alone is already PSPACE-complete (Papadimitriou and Tsitsiklis, 1987), not mentioning the additional computation required for model estimation and exploration. We leave the challenge of computationally efficient learning for future work.

### 1.1. Overview of techniques

The major technical challenge of this paper is to establish the sample efficiency guarantee of OMLE for learning any weakly revealing POMDPs despite the simplicity of the algorithm. Our results rely on the following three key ideas. To the best of our knowledge, the second and the third ideas are novel in the context of learning POMDPs, while the first technique was used by Jin et al. (2020a).

- **Observable Operator Model (OOM) (Jaeger, 2000):** OOM provides an alternative parameterization of the POMDP model, by representing the probability of a trajectory over observations and actions as the product of a series of linear operators, which is known as *observable operators*. For more details, see Section 5.1. Such linear structure facilitates us to use existing tools from matrix analysis to analyze POMDPs. Although OMLE algorithm does not explicitly utilize the OOM representation, the observable operators serve as important intermediate quantities in our analysis. They help us to bound the suboptimality of the learned policy as a function of the size of our confidence set.
- **MLE-based Confidence Set:** In contrast to the current mainstream approaches of learning POMDPs which use *spectral methods* to directly estimate either the model parameters or the observable operators (see, e.g., Azizzadenesheli et al., 2016; Guo et al., 2016; Jin et al., 2020a), we use the *maximum likelihood estimation* (MLE) approach, which provides implicit guarantees on learning observable operators. We achieve this by adapting the classic techniques for analyzing MLE (e.g., Geer et al., 2000). An appealing feature of the MLE approach is its generality and that the confidence set construction does not need to rely on the specific structure of the problem. The strength of this unified approach is that it allows OMLE to be used with almost no changes in both undercomplete and overcomplete POMDPs. In contrast, spectral-based algorithms require more careful designs that are adjusted to specific problems (such as undercomplete vs. overcomplete settings). These adjustments, if not done optimally, easily lead to requirement of unnecessary, artificial assumptions.
- **$\ell_1$ -norm eluder Dimension:** To prove the sample efficiency of optimistic algorithms, one needs to argue that, after a sufficient number of iterations, the size of the maintained confidence set is small enough to guarantee near-optimality of the learned policy. In the tabular setting, this is typically achieved by resorting to the pigeon-hole principle (e.g., Azar et al., 2017; Jin et al., 2018), while in the linear setting, one typically uses the so-called elliptical potential lemma (e.g., Lattimore and Szepesvári, 2020). To generalize these argument, Russo and Van Roy (2013) introduced the notion of eluder dimension for sets of real-valued functions with a common domain. The use of MLE-based confidence set requires us to develop a new result which is stronger than the standard elliptical potential arguments: While we have linear structures, the  $\ell_2$ -norms typically used are not suitable for our purposes. As such, the standard eluder dimension (which is tied to the  $\ell_2$ -norm) is also unsuitable. To address these

challenges, we introduce a variation of the eluder dimension, which is called  $\ell_1$ -norm eluder dimension, and which might be of independent interest.

## 1.2. Related Works

Due to space limitation, we postpone a more detailed discussion of related works to Appendix A.

## 2. Preliminaries

For a positive integer  $n$ , we let  $[n] = \{1, \dots, n\}$ . We consider episodic, tabular, partially observable Markov decision processes (POMDP). These processes generalize the standard Markov decision processes by making agents observe a “noisy function” of the state of a controlled Markov process. We consider time inhomogeneous, fixed horizon version of POMDPs. Formally, such a POMDP is specified by a tuple  $(\mathcal{S}, \mathcal{A}, \mathcal{O}; H, \mu_1, \mathbb{T}, \mathbb{O}; r)$ . Here  $\mathcal{S}$ ,  $\mathcal{A}$  and  $\mathcal{O}$  denote the space of state, action and observation respectively, with respective cardinalities  $|\mathcal{S}| = S$ ,  $|\mathcal{A}| = A$  and  $|\mathcal{O}| = O$ ;  $H$  denotes the length of each episode;  $\mu_1 \in \Delta_S$  denotes the distribution of the initial state where  $\Delta_S$  is the  $(S - 1)$ -dimensional probability simplex which we identify with the set of distributions over the states  $\mathcal{S}$ ;  $\mathbb{T} = \{\mathbb{T}_{h,a}\}_{(h,a) \in [H-1] \times \mathcal{A}}$  denotes the collection of transition matrices where  $\mathbb{T}_{h,a}$  is the  $S \times S$  transition matrix of action  $a$  at step  $h$  such that  $\mathbb{T}_{h,a}(\cdot | s)$  gives the distribution of the next state if the agent takes action  $a$  at state  $s$  and step  $h$ ;  $\mathbb{O} = \{\mathbb{O}_h\}_{h \in [H]}$  denotes the collection of emission matrices of size  $O \times S$  so that  $\mathbb{O}_h(\cdot | s)$  gives the distribution over observations at step  $h$  conditioned on the current hidden state being  $s$ ; and  $r = \{r_h\}_{h \in [H]}$  are the known reward functions from  $\mathcal{O}$  to  $[0, 1]$  such that the agent will receive reward  $r_h(o)$  when she observes  $o \in \mathcal{O}$  at step  $h$ .<sup>1</sup>

In a POMDP, the states are generally hidden from the agent: in every step a controlling agent can only see the observations and her own actions. At the beginning of each episode, the environment samples an initial state  $s_1$  from  $\mu_1$ . At each step  $h \in [H]$ , the agent first observes  $o_h$  that is sampled from  $\mathbb{O}_h(\cdot | s_h)$ , the observation distribution of hidden state  $s_h$  at step  $h$ . Then the agent receives reward  $r_h(o_h)$  that is computed from  $o_h$ , and takes action  $a_h$ . After this, the environment transitions to  $s_{h+1}$ , whose distribution follows  $\mathbb{T}_{h,a_h}(\cdot | s_h)$ . The current episode terminates immediately after  $a_H$  is taken. We use  $\tau_h = (o_1, a_1, \dots, o_h, a_h)$  to denote a trajectory from step 1 to step  $h$ .

A policy  $\pi = \{\pi_h : \mathcal{T}_h \rightarrow \Delta_A\}_{h=1}^H$  is a collection of  $H$  functions where  $\mathcal{T}_h = (\mathcal{O} \times \mathcal{A})^{h-1} \times \mathcal{O}$  denotes the set of all length- $h$  histories. Given a policy  $\pi$ , we use  $V^\pi$  to denote its value, which is defined as the expected total reward received under policy  $\pi$ :

$$V^\pi := \mathbb{E}_\pi \left[ \sum_{h=1}^H r_h(o_h) \right], \quad (1)$$

where the expectation is with respect to the randomness of the transitions, observations and the policy. Since the state, action, observation spaces and the horizon are all finite, there always exists an optimal policy  $\pi^*$  that achieves the optimal value  $V^* := \sup_\pi V^\pi$ . Different from MDPs, the optimal policies in POMDPs are in general history-dependent instead of only depending on the

1. This is equivalent to assuming that reward information is contained in the observation. We consider this setup to avoid the leakage of information about the latent states through rewards beyond observations. We remark that all results in this paper immediately extend to the more general setting where reward  $r(\tau_H)$  can be a function of the entire observation-action trajectory  $\tau_H$ , and is only received at the end of each episode.

current observation, which makes not only learning, but already computing a near-optimal policy in known POMDPs more challenging than doing the same in MDPs.

**Learning objective.** Our goal is to learn an  $\varepsilon$ -optimal policy  $\pi$  in the sense that  $V^\pi \geq V^* - \varepsilon$ , using a number of samples polynomial in all relevant parameters. We also consider the problem of learning with low regret. Suppose the agent interacts with POMDPs for  $K$  episodes, and plays a policy  $\pi_k$  in the  $k^{\text{th}}$  iteration for any  $k \in [K]$ . The total (expected) regret is then defined as:

$$\text{Regret}(K) = \sum_{k=1}^K [V^* - V^{\pi_k}].$$

The question then is whether a learner can keep the regret small.

**Notation.** We use bold upper-case letters  $\mathbf{B}$  to denote matrices and bold lower-case letters  $\mathbf{b}$  to denote vectors. Given a matrix  $\mathbf{B} \in \mathbb{R}^{m \times n}$ , we use  $\mathbf{B}_{ij}$  to denote its  $(i, j)^{\text{th}}$  entry,  $\sigma_k(\mathbf{B})$  to denote its  $k^{\text{th}}$  largest singular value, and  $\mathbf{B}^\dagger$  to denote its Moore-Penrose inverse. For a vector  $\mathbf{b} \in \mathbb{R}^m$ , we use  $\text{diag}(\mathbf{b})$  to denote a diagonal matrix with  $[\text{diag}(\mathbf{b})]_{ii} = \mathbf{b}_i$ .

### 3. Weakly Revealing POMDPs

The purpose of this section is to define the class of weakly revealing POMDPs. We first motivate our definition by revisiting the pathological instances which prevent sample-efficient learning of POMDPs in general. We then introduce the formal definition of weakly revealing POMDPs in the *undercomplete* setting when  $S \leq O$  and finally extend it to the *overcomplete* setting when  $S > O$ . All the proofs for this section are deferred to Appendix I.

#### 3.1. Hard instances of POMDPs

Here we revisit the hardness results and the pathological instances constructed by [Krishnamurthy et al. \(2016\)](#) and [Jin et al. \(2020a\)](#). As it turns out, learning POMDPs is statistically hard in the worst-case due to the existence of POMDPs with uninformative observations.

**Proposition 1 ([Krishnamurthy et al. \(2016\)](#); [Jin et al. \(2020a\)](#))** *There exists a class of 2-states  $H$ -horizon POMDPs whose observations reveal no information about the underlying states up to the end, such that any algorithm requires at least  $A^{\Omega(H)}$  samples to learn an  $\mathcal{O}(1)$ -optimal policy with a probability of  $1/2$  or higher.*

The hard instance is a combinatorial lock with unobserved states. Consider POMDPs with states  $s_{h,\text{good}}$  and  $s_{h,\text{bad}}$ ,  $h = 1, \dots, H$ . The emission probability  $\mathbb{O}_h(\cdot | s_{h,\text{good}})$  is precisely the same as  $\mathbb{O}_h(\cdot | s_{h,\text{bad}})$  for the all steps except the last one, so that the agent has absolutely no information about the latent state during the first  $H - 1$  steps. Let the initial state be  $s_{1,\text{good}}$ . Consider a special action sequence  $\{a_h^*\}_{h=1}^{H-1}$ , and construct the transition dynamics such that at each step  $1 \leq h \leq H - 1$ , the next state is  $s_{h+1,\text{good}}$  only if the previous state is  $s_{h,\text{good}}$  and the action taken is  $a_h^*$ . In all other cases, the environment transitions to  $s_{h+1,\text{bad}}$ . Finally, the agent will receive a reward of one only if she is in  $s_{H,\text{good}}$  at step  $H$ ; the agent receives zero reward otherwise.

It is not hard to see the optimal policy will take action  $a_h^*$  at step  $h$ , which will give a total reward of 1. However, since the agent effectively has no observation in the first  $H - 1$  steps, she has no option but to try out all possible action sequences, which requires  $A^{\Omega(H)}$  episodes to find the correct action sequence with constant probability.

### 3.2. Weakly revealing condition in the undercomplete setting

Based on the hard instances constructed above, we conclude that if the observations do not contain information to distinguish two different latent states, then learning these POMDPs is statistically hard. For POMDPs with more than two states, the hardness result above can be easily extended to the case where there exist two mixtures of latent states with disjoint support such that the observations do not contain any information to distinguish these two mixtures. Concretely, by a mild abuse of language, a mixture of states is identified by a probability vector  $\nu \in \Delta_S$ ;  $\nu_1$  and  $\nu_2$  are said to have disjoint support if  $\text{supp}(\nu_1) \cap \text{supp}(\nu_2) = \emptyset$ .

A direct approach to rule out the above-described pathological instances is to just assume that any two latent state mixtures  $\nu_1, \nu_2$  that have disjoint support induce distinct distributions over observations, that is,  $\mathbb{O}_h \nu_1 \neq \mathbb{O}_h \nu_2$  for all  $h \in [H]$  where  $\mathbb{O}_h$  is the  $O \times S$  emission matrix at step  $h$ . A linear algebraic argument then shows that this condition is equivalent to that the rank of the emission matrix  $\mathbb{O}_h$  is  $S$ .

**Proposition 2** *The emission matrix  $\mathbb{O}_h$  is rank  $S$  if and only if the induced distributions over observations are distinct for any two mixtures of latent states with disjoint support.*

The weakly revealing condition is simply a robust version of the condition that the rank of the emission matrices is  $S$ —it assumes the  $S^{\text{th}}$  singular value of emission matrix  $\mathbb{O}_h$  is lower bounded. This condition in the undercomplete setting was first identified by Jin et al. (2020a) as a technical condition to ensure the sample efficiency of their algorithms.

**Assumption 1 ( $\alpha$ -weakly revealing condition)** *There exists  $\alpha > 0$ , such that  $\min_h \sigma_S(\mathbb{O}_h) \geq \alpha$ .*

This condition ensures that the observations contain enough information to distinguish any two mixtures of states given a sufficiently large number of samples.

We call Assumption 1 the “*weakly*” revealing condition to distinguish it from the setup known as rich observation or block MDP in the literature (Jiang et al., 2017; Du et al., 2019; Misra et al., 2020). The latter setup considers the problem where the latent state can be directly recovered from any single observation and the stage  $h$  in the episode. That is, the latent state is completely revealed by the observation. Therefore, technically speaking, block MDPs are fully observable, which is in a way “diagonally opposite” to the setting we consider.

Finally, we note that since  $\mathbb{O}_h$  is a matrix of size  $O \times S$ , Assumption 1 implicitly requires  $S \leq O$ . That is, it only holds in the undercomplete setting.

### 3.3. Weakly revealing condition in the overcomplete setting

In the overcomplete setting, we have  $S > O$ . It is information-theoretically impossible to distinguish any two mixtures of latent states by inspecting observations only in a single step. The key observation here is that we should instead inspect the distribution of observations for  $m$  consecutive steps. We note that the number of all possible observable sequence  $(o_1, a_1, \dots, a_{m-1}, o_m)$  of length  $m$  is  $O^m A^{m-1}$ , which is larger than  $S$  when  $m \geq \Omega(\log S)$ .

To state our assumption, we define the  $m$ -step emission-action matrices

$$\{\mathbb{M}_h \in \mathbb{R}^{(A^{m-1}O^m) \times S}\}_{h \in [H-m+1]}$$

as follows: For an observation sequence  $\mathbf{o}$  of length  $m$ , initial state  $s$  and action sequence  $\mathbf{a}$  of length  $m - 1$ , we let  $[\mathbb{M}_h]_{(\mathbf{a}, \mathbf{o}), s}$  be the probability of receiving  $\mathbf{o}$  provided that the action sequence  $\mathbf{a}$  is used from state  $s$  and step  $h$ :

$$[\mathbb{M}_h]_{(\mathbf{a}, \mathbf{o}), s} = \mathbb{P}(o_{h:h+m-1} = \mathbf{o} \mid s_h = s, a_{h:h+m-2} = \mathbf{a}) \quad \text{for all } (\mathbf{a}, \mathbf{o}) \in \mathcal{A}^{m-1} \times \mathcal{O}^m \text{ and } s \in \mathcal{S}. \quad (2)$$

Similar to the undercomplete case, the weakly revealing condition in the overcomplete setting assumes that the  $S^{\text{th}}$  singular value of the  $m$ -step emission matrix  $\mathbb{M}_h$  is lower bounded.

**Assumption 2 ( $m$ -step  $\alpha$ -weakly revealing condition)** *There exists  $m \in \mathbb{N}$ ,  $\alpha > 0$  such that  $\min_{h \in [H-m+1]} \sigma_S(\mathbb{M}_h) \geq \alpha$  where  $\mathbb{M}_h$  is the  $m$ -step emission matrix defined in (2).*

Assumption 2 ensures that the observable sequence in the next  $m$  consecutive steps contain enough information to distinguish any two mixtures of states given a sufficiently large number of observations. Assumption 1 is a special case of Assumption 2 with  $m = 1$ .

Finally, we remark that in case that  $O^m \geq S$ , a sufficient condition to make Assumption 2 hold is that: for any stage  $h$ , there exists a  $(m - 1)$ -step action sequence such that the  $m$ -step observation sequences under this action sequence is  $\alpha$ -weakly revealing the hidden state. Formally, for any  $h \in [H]$  and  $\mathbf{a} \in \mathcal{A}^{m-1}$  let  $\mathbb{M}_{h, \mathbf{a}}$  stands for the  $O^m \times S$  matrix obtained from  $\mathbb{M}_h$  by selecting the rows of  $\mathbb{M}_h$  where the row-index corresponds to  $\mathbf{a}$ . That is,  $(\mathbb{M}_{h, \mathbf{a}})_{\mathbf{o}, s} = [\mathbb{M}_h]_{(\mathbf{a}, \mathbf{o}), s}$ .

**Proposition 3** *Assume that  $O^m \geq S$ , then Assumption 2 holds if  $\max_{\mathbf{a} \in \mathcal{A}^{m-1}} \sigma_S(\mathbb{M}_{h, \mathbf{a}}) \geq \alpha$  for all  $h \in [H - m + 1]$ .*

## 4. Main Results

In this section, we present our algorithm—*Optimistic Maximum Likelihood Estimation* (OMLE) and its theoretical guarantees for learning weakly revealing POMDPs in both the undercomplete and the overcomplete settings.

### 4.1. Undercomplete setting

For clarity, we first present the algorithm and results for learning undercomplete POMDPs under Assumption 1. As we will see in the later section, with a minor modification this algorithm also generalizes to learning overcomplete POMDPs under Assumption 2.

**Algorithm description** To condense notations, we use  $\theta = (\mathbb{T}, \mathbb{O}, \mu_1)$  to denote the model parameters of a POMDP and use  $\Theta$  to denote the collections of all possible model parameters  $\theta$  that correspond to POMDPs with  $S$  states,  $A$  actions, and  $O$  observations. To make the dependence on  $\theta$  explicit, we will use  $V^\pi(\theta)$  to denote the value of a policy  $\pi$ , while we use  $\mathbb{P}_\theta^\pi(\tau)$  to denote the probability of observing a trajectory  $\tau$  under policy  $\pi$ , when the underlying POMDP is given by  $\theta$ . We also use  $\mathbb{O}_h(\theta)$  ( $\mathbb{M}_h(\theta)$ ) to denote the emission matrix of  $\theta$  (respectively, the multistep emission matrix of  $\theta$ ).

Algorithm 1 gives the pseudocode of OMLE. As can be seen from this pseudocode, in each episode  $k$  there are two main steps:

---

**Algorithm 1** OPTIMISTIC MAXIMUM LIKELIHOOD ESTIMATION (OMLE)
 

---

- 1: **Initialize:**  $\mathcal{B}^1 = \{\hat{\theta} \in \Theta : \min_h \sigma_S(\mathbb{O}_h(\hat{\theta})) \geq \alpha\}$ ,  $\mathcal{D} = \{\}$
- 2: **for**  $k = 1, \dots, K$  **do**
- 3:   compute  $(\theta^k, \pi^k) = \operatorname{argmax}_{\hat{\theta} \in \mathcal{B}^k, \pi} V^\pi(\hat{\theta})$
- 4:   execute policy  $\pi^k$  to collect a trajectory  $\tau^k := (o_1^k, a_1^k, \dots, o_h^k, a_h^k)$
- 5:   add  $(\pi^k, \tau^k)$  into  $\mathcal{D}$  and update

$$\mathcal{B}^{k+1} = \left\{ \hat{\theta} \in \Theta : \sum_{(\pi, \tau) \in \mathcal{D}} \log \mathbb{P}_{\hat{\theta}}^\pi(\tau) \geq \max_{\theta' \in \Theta} \sum_{(\pi, \tau) \in \mathcal{D}} \log \mathbb{P}_{\theta'}^\pi(\tau) - \beta \right\} \cap \mathcal{B}^1 \quad (3)$$


---

- Optimistic planning (Lines 3-4): find the POMDP model  $\theta^k$  with the highest optimal value in the confidence set  $\mathcal{B}^k$  and follow the associated optimal policy  $\pi^k$  in the episode to collect a trajectory  $\tau^k$ .<sup>2</sup>
- Confidence set update (Line 5): add the newly collected policy-trajectory pair into the dataset, and then update the confidence set to include those models that assign a total log-likelihood to the data that is “close” to the maximum possible such total log-likelihood. In particular, the form of the confidence set is

$$\left\{ \hat{\theta} \in \Theta : \sum_{(\pi, \tau) \in \mathcal{D}} \log \mathbb{P}_{\hat{\theta}}^\pi(\tau) \geq \max_{\theta' \in \Theta} \sum_{(\pi, \tau) \in \mathcal{D}} \log \mathbb{P}_{\theta'}^\pi(\tau) - \beta \right\} \cap \mathcal{B}^1,$$

where  $\mathcal{B}^1$  is the initial confidence set that contains all  $\alpha$ -weakly revealing models of a given size.

Compared to the standard maximum likelihood estimation (MLE) approach, all  $\alpha$ -weakly revealing models with a sufficiently high likelihood are allowed and the size of this set is controlled by  $\beta \geq 0$ . In particular, if  $\beta = 0$ , the confidence set collapses to the solutions of MLE.

In our algorithm, the choice of  $\beta$  is governed by the magnitude of the “statistical noise” introduced by various random events. By analyzing this noise, one can choose the value of  $\beta$  to guarantee that the true POMDP model is always contained in the resulting confidence set with a prescribed probability (see Proposition 13 for a rigorous statement).

We emphasize that the algorithm design of MLE is considerably simpler than that of prior provably sample-efficient algorithms for learning POMDPs (see, e.g., Azizzadenesheli et al., 2016; Guo et al., 2016; Jin et al., 2020a), which rely on spectral methods.

**Theoretical guarantees** Our main result, which shows that OMLE will achieve small regret in any *weakly revealing POMDPs* (Assumption 1), is as follows:

**Theorem 4 (Regret of OMLE)** *There exists an absolute constant  $c > 0$  such that for any  $\delta \in (0, 1]$  and  $S, A, O, H, K \in \mathbb{N}$ , if we choose  $\beta = c (H(S^2 A + SO) \log(SAOHK) + \log(K/\delta))$  in*

---

2. Our algorithm, as well as all existing algorithms for learning large classes of POMDPs, is computationally inefficient. In particular, a naive implementation of optimistic planning (Line 3) is to enumerate all POMDP models in an  $\varepsilon$ -cover of the confidence set and compute their optimal policies, which requires  $e^{\Omega(HS^2 A + HSO)}$  time in the worst case.



Algorithm 1, then, for any POMDP with  $S$  states,  $A$  actions,  $O$  observations and horizon  $H$  and satisfying Assumption 1, with probability at least  $1 - \delta$ ,

$$\text{Regret}(k) \leq \text{poly}(S, A, O, H, \alpha^{-1}, \log(\delta^{-1}K)) \cdot \sqrt{k} \quad \text{for all } k \in [K].$$

The proof, as well as the specific polynomial dependency, is presented in Appendix F. Note that the growth rate of regret as a function  $k$  is optimal (Auer et al., 1995).

Moreover, by the standard online-to-batch conversion (Cesa-Bianchi et al., 2004), the regret bound immediately implies the following sample complexity result:

**Corollary 5 (Sample Complexity of OMLE)** *Under the same setting as Theorem 4, when  $K \geq \text{poly}(S, A, O, H, \alpha^{-1}, \log(\varepsilon^{-1}\delta^{-1})) \cdot \varepsilon^{-2}$ , with probability at least  $1 - \delta$ , the uniform mixture of the policies produced by OMLE is  $\varepsilon$ -optimal. I.e.,  $(1/K) \cdot \sum_{k=1}^K V^{\pi^k} \geq V^* - \varepsilon$ .*

Here, the  $\tilde{O}(\varepsilon^{-2})$  dependence is also optimal up to log factors. Previous work by Jin et al. (2020a) also provides polynomial sample-complexity guarantee for learning  $\alpha$ -weakly revealing POMDPs under Assumption 1. The present result improves over the results of Jin et al. (2020a) in the following aspects:

- While the OOM-UCB algorithm of Jin et al. (2020a) heavily exploited the special structure of undercomplete POMDPs, OMLE appears in a much simpler form and the algorithm design arguably does not use this special structure. As a result, OMLE can be easily extended to learning *multi-step weakly revealing* POMDPs, while to the best of our knowledge OOM-UCB cannot.
- In terms of theoretical guarantees, OMLE enjoys a near-optimal  $\sqrt{k}$ -regret while OOM-UCB was only shown to achieve a regret of size  $O(k^{2/3})$ . The higher regret of OOM-UCB is due to the limitation of its exploration mechanism.<sup>3</sup>

Finally, observe that the upper bound in Theorem 4 depends polynomially on the inverse of  $\alpha$ —an upper bound on the  $\ell_1$ -norm of the pseudoinverse of the emission matrices in Assumption 1. This polynomial dependence turns out to be unavoidable as is shown by the following lower bound.

**Theorem 6 (Necessity of  $\text{poly}(\alpha^{-1})$  dependency)** *For any  $\alpha \in (0, 1/2)$  and  $H, A \in \mathbb{N}^+$ , there exists an undercomplete  $\alpha$ -weakly revealing POMDP with  $S, O = \mathcal{O}(1)$  so that any algorithm requires at least  $\Omega(\min\{\frac{1}{\alpha H}, A^{H-1}\})$  samples to learn a  $(1/2)$ -optimal policy with probability  $1/6$  or higher.*

Theorem 6 implies that a polynomial dependence on  $1/\alpha$  is in general unavoidable in the sense that any algorithm either needs to suffer a regret exponential in the horizon  $H$ , or its regret needs to be polynomially dependent on  $1/\alpha$ . The proof of Theorem 6 is provided in Appendix H.

## 4.2. Overcomplete setting

We now turn to the more challenging setting of learning in overcomplete POMDPs, where the number of hidden states can be larger than the number of observations. We show that a simple variant of OMLE is able to learn weakly-revealing overcomplete POMDPs in a polynomial number of samples. As we shall see, we pay a nontrivial price for the increased generality: while we can still achieve rate-optimal PAC-results, we compromise on the regret of the algorithm.

3. OOM-UCB itself is not a no-regret algorithm. However, combining its  $\tilde{O}(\text{poly}(\cdot)/\varepsilon^2)$  sample complexity guarantee with the explore-then-commit strategy implies a  $\tilde{O}(\text{poly}(\cdot) \times k^{2/3})$ -regret.

---

**Algorithm 2** MULTI-STEP OPTIMISTIC MAXIMUM LIKELIHOOD ESTIMATION
 

---

- 1: **Initialize:**  $\mathcal{B}^1 = \{\hat{\theta} \in \Theta : \min_h \sigma_S(\mathbb{M}_h(\hat{\theta})) \geq \alpha\}$ ,  $\mathcal{D} = \{\}$
- 2: **for**  $k = 1, \dots, K$  **do**
- 3:    $(\theta^k, \pi^k) = \operatorname{argmax}_{\hat{\theta} \in \mathcal{B}^k, \pi} V^\pi(\hat{\theta})$
- 4:   **for**  $h = 0, \dots, H - m$  **do**
- 5:     execute policy  $\pi_{1:h}^k \circ \operatorname{uniform}(\mathcal{A})$  to collect a trajectory  $\tau^{k,h}$   
       then add  $(\pi_{1:h}^k \circ \operatorname{uniform}(\mathcal{A}), \tau^{k,h})$  into  $\mathcal{D}$
- 6:   update

$$\mathcal{B}^{k+1} = \left\{ \hat{\theta} \in \Theta : \sum_{(\pi, \tau) \in \mathcal{D}} \log \mathbb{P}_{\hat{\theta}}^\pi(\tau) \geq \max_{\theta' \in \Theta} \sum_{(\pi, \tau) \in \mathcal{D}} \log \mathbb{P}_{\theta'}^\pi(\tau) - \beta \right\} \cap \mathcal{B}^1 \quad (4)$$


---

**Algorithm description** Algorithm 2 shows the pseudo-code of OMLE suitable for  $m$ -step  $\alpha$ -weakly revealing overcomplete POMDPs. While the basic structure of the method is the same as before, the general OMLE, which we call *multi-step OMLE*, differs from the basic version in two important aspects:

- Instead of merely following the optimistic policy, Algorithm 2 adopts a more active strategy for exploration. Specifically, for each optimistic policy  $\pi^k$ , the learner will one by one experiments with  $(H - m + 1)$  policies that are obtained by picking a within-episode time index  $h \in \{0, \dots, H - m\}$  and then following policy  $\pi^k$  for the first  $h$  steps, and then picking actions uniformly at random in the remaining steps of the episode. We denote the resulting policy by  $\pi_{1:h}^k \circ \operatorname{uniform}(\mathcal{A})$ , which abuses notation, but should improve readability.
- When constructing the confidence set, Algorithm 2 requires the minimum singular value of the  $m$ -step emission-action matrix (defined in equation (2)) to be lower bounded by  $\alpha$ , which enforces the multi-step  $\alpha$ -weakly revealing condition in Assumption 2.

By trying random action sequences after executing  $\pi^k$  for the initial  $h$  steps, the learner can gather more information about the hidden states reachable by  $\pi^k$  at step  $h$  and therefore can better learn the system dynamics under  $\pi^k$ . The price of trying random actions is that the algorithm as described here will in general have linear regret. Nevertheless, with an online-to-batch conversion, Algorithm 2 serves as a suitable approach to learning a good policy with low sample complexity.

**Theoretical guarantees** Our main result in this section bounds the total suboptimality of the policies  $\pi^1, \dots, \pi^k$  chosen by OMLE. Note that since OMLE is not following these policies, the regret of OMLE is different (in general, higher) than the total suboptimality.

**Theorem 7 (Total suboptimality of multi-step OMLE)** *There exists an absolute constant  $c > 0$  such that for any  $\delta \in (0, 1]$  and  $S, A, O, K, H \in \mathbb{N}$ , if we choose parameter  $\beta$  in Algorithm 2 as  $\beta = c (H(S^2 A + SO) \log(SAOH) + \log(KH/\delta))$ , then, for any POMDP with  $S$  states,  $A$  actions,  $O$  observations and horizon  $H$  and satisfying Assumption 2, with probability at least  $1 - \delta$ ,*

$$\sum_{t=1}^k (V^* - V^{\pi^t}) \leq \operatorname{poly}(S, A^m, O, H, \alpha^{-1}, \log(\delta^{-1}K)) \cdot \sqrt{k} \quad \text{for all } k \in [K].$$

The specific polynomial dependency is presented in Appendix G. This form of the result is preferred as it makes a comparison to Theorem 4 more direct and it also reveals a bit of the proof strategy. The significance of this result is that, using the standard online-to-batch conversion (Cesa-Bianchi et al., 2004), we get the following sample complexity results.

**Corollary 8 (Sample Complexity of multi-step OMLE)** *Under the same setting as Theorem 7, when  $K \geq \text{poly}(S, A^m, O, H, \alpha^{-1}, \log(\varepsilon^{-1} \delta^{-1})) \cdot \varepsilon^{-2}$ , with probability at least  $1 - \delta$ , the uniform mixture of the policies produced by multi-step OMLE is  $\varepsilon$ -optimal. I.e.,  $(1/K) \cdot \sum_{k=1}^K V^{\pi^k} \geq V^* - \varepsilon$ .*

Up to polylogarithmic factors, the dependence on  $\varepsilon$  in this result is unimprovable. Using an explore-then-exploit strategy, this latter result gives rise to a method that enjoys  $\tilde{O}(K^{2/3})$  regret, where the constants hidden are still polynomial in the relevant quantities. To our knowledge, for small fixed  $m$ , this is the first sample-efficient result for learning overcomplete POMDPs in the *exploration* setting where the algorithm needs to reason about how to collect information efficiently.

A natural question here is whether the exponential dependence on  $m$  in Theorem 7 is necessary. We answer this question by providing the following lower bound, which rules out the possibility of an upper bound polynomial in  $m$ .

**Theorem 9 (Necessity of  $A^{\Omega(m)}$  dependency)** *For any  $m, A \in \mathbb{N}^+$ , there exists a POMDP with  $S, H, O = O(m)$  and satisfying Assumption 2 with  $\alpha \geq 1$  so that any algorithm requires at least  $\Omega(A^{m-1})$  samples to learn a  $(1/2)$ -optimal policy with probability at least  $1/2$ .*

## 5. Proof Overview

We provide a proof overview of Theorem 4 for learning undercomplete *weakly revealing* POMDPs (Assumption 1). We defer the full proof to Appendix F. The proof for learning overcomplete POMDPs (Theorem 7) follows a similar strategy, which is described in Appendix G.

### 5.1. Observable operator models

To begin with, we introduce the observable operators (Jaeger, 2000) that provide an alternate parameterization of POMDPs. These operators will serve as intermediate quantities in our analysis: They will allow us to bound the suboptimality of the learned policies as a function of the “width” of the MLE confidence set. Given the transition matrices  $\{\mathbb{T}_{h,a}\}_{(h,a) \in [H] \times \mathcal{A}}$ , the observation matrices  $\{\mathbb{O}_h\}_{h \in [H]}$ , and the initial distribution  $\mu_1$ , the observable operators  $\{\mathbf{B}_h(o, a)\}_{(h,o,a) \in [H-1] \times \mathcal{O} \times \mathcal{A}}$  and the initial  $\mathbf{b}_0$  observation distribution are given by

$$\mathbf{B}_h(o, a) = \mathbb{O}_{h+1} \mathbb{T}_{h,a} \text{diag}(\mathbb{O}_h(o | \cdot)) \mathbb{O}_h^\dagger, \quad \mathbf{b}_0 = \mathbb{O}_1 \mu_1,$$

where  $\mathbb{O}_h(o | \cdot) \in \mathbb{R}^S$  denotes the  $o^{\text{th}}$  row of  $\mathbb{O}_h$ . It is known that these operators give an equivalent parameterization of the POMDPs: For any policy, the distribution induced by a POMDP over the possible trajectories of observation-action pairs can be described solely using these operators. In particular, the probability of observing trajectory  $\tau_h = (o_1, a_1, \dots, o_h, a_h)$  under policy  $\pi$  in POMDP model  $\theta$  is given by

$$\mathbb{P}_\theta^\pi(\tau_h) = \pi(\tau_h) \cdot \left( \mathbf{e}_{o_h}^\top \mathbf{B}_{h-1}(o_{h-1}, a_{h-1}; \theta) \cdots \mathbf{B}_1(o_1, a_1; \theta) \mathbf{b}_0(\theta) \right), \quad (5)$$

where  $\pi(\tau_h) := \prod_{h'=1}^h \pi(a_{h'} \mid o_{h'}, \tau_{h-1})$  represents the part of the probability of  $\tau_h$  that can be attributed to the randomness of the policy and we used  $\mathbf{B}_j(\cdot; \theta)$  to denote the observable operators underlying  $\theta$ . One important advantage of adopting this operator representation of POMDPs is that the linear structure facilitates us to use existing tools from matrix analysis to analyze the error of operator estimates.

## 5.2. Step 1: bound the regret by the error of operator estimates

By analyzing the relaxed MLE condition, one can prove that the groundtruth POMDP model  $\theta^*$  is contained in confidence set  $\mathcal{B}^k$  for all  $k \in [K]$  with high probability (see Proposition 13 in Appendix B). Therefore, from now on assume that  $\theta^* \in \bigcap_{k \in [K]} \mathcal{B}^k$  holds. Now, recall that we choose the model estimate and the behavior policy optimistically in Algorithm 1, i.e.,  $(\theta^k, \pi^k) = \operatorname{argmax}_{\hat{\theta} \in \mathcal{B}^k, \pi} V_{\hat{\theta}}^{\pi}$ . As a result, we have  $V^* = \max_{\pi} V_{\theta^*}^{\pi} \leq \max_{\hat{\theta} \in \mathcal{B}^k, \pi} V_{\hat{\theta}}^{\pi} = V_{\theta^k}^{\pi^k}$  for all  $k \in [K]$ . From this, we get

$$\sum_{t=1}^k V_{\theta^*}^* - V_{\theta^*}^{\pi^t} \leq \sum_{t=1}^k V_{\theta^t}^{\pi^t} - V_{\theta^*}^{\pi^t} \leq H \sum_{t=1}^k \sum_{\tau_H} |\mathbb{P}_{\theta^t}^{\pi^t}(\tau_H) - \mathbb{P}_{\theta^*}^{\pi^t}(\tau_H)|, \quad (6)$$

where  $\tau_H = (o_1, a_1, \dots, o_H, a_H)$  denotes a whole trajectory and the second inequality uses the fact that the cumulative reward of each trajectory is bounded by  $H$ . Therefore, to prove Theorem 1, it suffices to bound the total cumulated error in estimating the probability of the individual trajectories, cf. the RHS of (6).

By using the OOM representations in (5), it turns out that we can bound the RHS of (6) by the error in estimating each observable operator. To simplify notation, we abbreviate  $\mathbf{B}_h(o, a; \theta^*)$ ,  $\mathbf{b}_0(\theta^*)$  as  $\mathbf{B}_h(o, a)$ ,  $\mathbf{b}_0$ , and denote  $\mathbf{B}_h^t(o, a) := \mathbf{B}_h(o, a; \theta^t)$ ,  $\mathbf{b}_0^t := \mathbf{b}_0(\theta^t)$ . With this notation, we have the following result:

**Lemma 10** *For any  $k \in \mathbb{N}$ , the RHS of (6) is upper bounded by*

$$\frac{H\sqrt{S}}{\alpha} \left( \sum_{t=1}^k \sum_{h=1}^{H-1} \sum_{\tau_h} \left\| (\mathbf{B}_h(o_h, a_h) - \mathbf{B}_h^t(o_h, a_h)) \mathbf{b}(\tau_{h-1}) \right\|_1 \times \pi^t(\tau_h) + \|\mathbf{b}_0 - \mathbf{b}_0^t\|_1 \right), \quad (7)$$

where  $\mathbf{b}(\tau_h) := \left( \prod_{h'=1}^h \mathbf{B}_{h'}(o_{h'}, a_{h'}) \right) \mathbf{b}_0$  is the “belief vector” associated with trajectory  $\tau_h = (o_1, a_1, \dots, o_h, a_h)$ .

In Equation (7) we abused notation in a few ways: In the innermost sum over the observation-action trajectories  $\tau_h$  of length  $h$ ,  $\tau_{h-1}$  refers to the prefix of  $\tau_h$  where the last observation-action is dropped. Also, in this sum,  $o_h, a_h$  refer to the last observation-action pair of  $\tau_h$ .

Lemma 10 is obtained from Lemma 27, which states the same result for an arbitrary sequence of observable operators. As a result, in order to control the regret, it suffices to control the estimation error of each operator. Importantly, here we do not need to recover the operators accurately at all entries, which, in general, is also impossible when there are hard-to-reach latent states. Instead, we only care about the projections of the errors onto the belief vectors, which are further reweighted by the probability of the behavior policies. Therefore, it suffices to learn the operators accurately only in those directions that are adequately covered by the reweighted belief vectors.

### 5.3. Step 2: derive constraints for the operator estimates from OMLE

Now let us make a detour to see what guarantees OMLE can provide for our operator estimates. As a result of the classic MLE analysis (e.g., Geer et al., 2000), we can show under the same choice of  $\beta$  as Theorem 4, with high probability

$$\sum_{t=1}^{k-1} \left\| \mathbb{P}_{\theta^k}^{\pi^t}(\tau_h = \cdot) - \mathbb{P}_{\theta^*}^{\pi^t}(\tau_h = \cdot) \right\|_1^2 = \mathcal{O}(\beta) \quad \text{for all } (k, h) \in [K] \times [H]. \quad (8)$$

(Proposition 14 in Appendix B gives the precise result.) In brief, this means the model estimate in the  $k^{\text{th}}$  iteration, that is  $\theta^k$ , can be used to predict the behavior of the policies followed *before* the  $k^{\text{th}}$  iteration to a certain accuracy. To proceed, we represent the probabilities in equation (8) by products of operators using equation (5) and perform further algebraic transformations, which eventually leads to the following lemma for our operator estimates. The proof of this lemma is given in Appendix F.2.

**Lemma 11** *Suppose the relation in equation (8) holds, then for all  $(k, h) \in [K] \times [H]$*

$$\sum_{t=1}^{k-1} \sum_{\tau_h} \left\| \left( \mathbf{B}_h(o_h, a_h) - \mathbf{B}_h^k(o_h, a_h) \right) \mathbf{b}(\tau_{h-1}) \right\|_1 \times \pi^t(\tau_h) = \mathcal{O} \left( \frac{\sqrt{S\beta k}}{\alpha} \right). \quad (9)$$

Intuitively, the constraints above imply the operator estimates in the  $k^{\text{th}}$  iteration are close to the true operators when being projected onto the belief vectors that are reweighted by the historical policies. However, a careful examination shows that (9) cannot be directly used to control (7) because (7) involves the operator error of  $\theta^t$  reweighted by  $\pi^t$  that is the behavior policy in the *same* iteration. This is very different from (9). We deal with this problem in Step 3.

### 5.4. Step 3: bridge Step 1 and 2 via $\ell_1$ -norm eluder dimension

To prove the sample efficiency of optimistic algorithms, one needs to argue that, after a sufficient number of iterations, the size of the maintained confidence set is small enough to guarantee near-optimality of the learned policy. This is typically achieved by resorting to the pigeon-hole principle in the tabular setting (e.g., Azar et al., 2017; Jin et al., 2018), or to the elliptical potential lemma in the linear setting (e.g., Lattimore and Szepesvári, 2020).

In the context of this paper, by further algebraic transformations, we reduce the problem of bounding (7) by (9) to proving the following algebraic inequality, which plays a similar role as the elliptical potential lemma. The full inequality is more involved (see Proposition 22 in Appendix D); here we present a simplified version for the sake of simplicity.

**Proposition 12** *Suppose sequences  $\{w_{k,j}\}_{(k,j) \in [K] \times [m]}$  and  $\{x_{k,i}\}_{(k,i) \in [K] \times [n]}$  satisfy that  $w_{k,j}, x_{k,i} \in \mathbb{R}^d$  for all  $(k, i, j) \in [K] \times [n] \times [m]$ . Suppose that we further have*

$$\sum_{t=1}^{k-1} \sum_{j=1}^m \sum_{i=1}^n |w_{k,j}^\top x_{t,i}| \leq \sqrt{k}, \quad \sum_{j=1}^m \|w_{k,j}\|_2 \leq 1 \quad \text{and} \quad \sum_{i=1}^n \|x_{k,i}\|_2 \leq 1 \quad \text{for all } k \in [K].$$

*Then we have  $\sum_{t=1}^k \sum_{j=1}^m \sum_{i=1}^n |w_{t,j}^\top x_{t,i}| = \tilde{\mathcal{O}}(\sqrt{\zeta k})$  for all  $k \in [K]$ , where  $\zeta$  is a parameter that depends on  $d$  only.*

At a high level, the precondition and the target in Proposition 12 correspond to equation (9) and (7), respectively (see Appendix F.3 for details). In the special case of  $m = n = 1$ , Proposition 12 reduces to

$$\text{if } \sum_{t=1}^{k-1} |w_k^\top x_t| \leq \sqrt{k} \text{ for all } k \in [K], \quad \text{then } \sum_{t=1}^k |w_t^\top x_t| = \tilde{O}(\sqrt{\zeta k}) \text{ for all } k \in [K]. \quad (10)$$

We compare this with the standard elliptical potential lemma in linear bandit literature (e.g., Lattimore and Szepesvári, 2020), which is typically of the form:

$$\text{if } \sum_{t=1}^{k-1} |w_k^\top x_t|^2 \leq 1 \text{ for all } k \in [K], \quad \text{then } \sum_{t=1}^k |w_t^\top x_t| = \tilde{O}(\sqrt{dk}) \text{ for all } k \in [K]. \quad (11)$$

We remark that the precondition in (11) directly implies the precondition in (10) by the Cauchy-Swartz inequality. That is, Proposition 12 is stronger than the standard elliptical potential lemma, and we need to develop new techniques to prove Proposition 12.

Noting the close relation between the elliptical potential lemma and the framework of eluder dimension (Russo and Van Roy, 2013) (in its original  $\ell_2$ -norm form), we develop a new framework based on the  $\ell_1$ -norm counterpart of eluder dimension, and adapt corresponding techniques to prove that in (10) and Proposition 12 we can allow the choice of  $\zeta = d^2$  which is one  $d$  factor worse than the standard elliptical potential lemma. We defer the details of this framework to Appendix D.

## 6. Conclusion

In this paper, we identified a new rich class of POMDPs, which we call *weakly revealing* POMDPs. *Weakly revealing* POMDPs subsume a majority of existing POMDPs that are known to be sample-efficiently learnable, and include both undercomplete and overcomplete POMDPs. We further propose a new simple algorithm, OMLE, which combines optimism with maximum likelihood estimation. We prove that OMLE can learn a near-optimal policy for any weakly revealing POMDP using polynomial samples. We complement our positive results with two lower bounds to justify the necessity of the appearance of certain problem-dependent quantities in our upper bounds. Finally, while our work shows that sample-efficient learning is possible in large classes of POMDPs, computationally efficient learning of POMDPs remains challenging, which we leave for future work.

## Acknowledgments

We thank Nan Jiang for valuable discussions on the sample complexity of *multi-step OMLE*.

## References

- Alekh Agarwal, Sham Kakade, Akshay Krishnamurthy, and Wen Sun. Flambe: Structural complexity and representation learning of low rank mdps. *NeurIPS*, 2020.
- Ilge Akkaya, Marcin Andrychowicz, Maciek Chociej, Mateusz Litwin, Bob McGrew, Arthur Petron, Alex Paino, Matthias Plappert, Glenn Powell, Raphael Ribas, et al. Solving Rubik’s cube with a robot hand. *arXiv preprint arXiv:1910.07113*, 2019.

- Animashree Anandkumar, Rong Ge, Daniel Hsu, Sham M Kakade, and Matus Telgarsky. Tensor decompositions for learning latent variable models. *Journal of machine learning research*, 15: 2773–2832, 2014.
- Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. Gambling in a rigged casino: The adversarial multi-armed bandit problem. In *Proceedings of IEEE 36th annual foundations of computer science*, pages 322–331. IEEE, 1995.
- Alex Ayoub, Zeyu Jia, Csaba Szepesvari, Mengdi Wang, and Lin Yang. Model-based reinforcement learning with value-targeted regression. In *International Conference on Machine Learning*, pages 463–474. PMLR, 2020.
- Mohammad Gheshlaghi Azar, Ian Osband, and Rémi Munos. Minimax regret bounds for reinforcement learning. In *International Conference on Machine Learning*, pages 263–272. PMLR, 2017.
- Kamyar Azizzadenesheli, Alessandro Lazaric, and Animashree Anandkumar. Reinforcement learning of POMDPs using spectral methods. In *Conference on Learning Theory*, pages 193–256. PMLR, 2016.
- Kamyar Azizzadenesheli, Yisong Yue, and Animashree Anandkumar. Policy gradient in partially observable environments: Approximation and convergence. *arXiv preprint arXiv:1810.07900*, 2018.
- Joseph L Bower and Clark G Gilbert. *From resource allocation to strategy*. Oxford University Press, 2005.
- Noam Brown and Tuomas Sandholm. Superhuman AI for multiplayer poker. *Science*, 365(6456): 885–890, 2019.
- Anthony R Cassandra, Leslie Pack Kaelbling, and James A Kurien. Acting under uncertainty: Discrete Bayesian models for mobile-robot navigation. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems. IROS’96*, volume 2, pages 963–972. IEEE, 1996.
- Nicolo Cesa-Bianchi, Alex Conconi, and Claudio Gentile. On the generalization ability of online learning algorithms. *IEEE Transactions on Information Theory*, 50(9):2050–2057, 2004.
- Christoph Dann, Tor Lattimore, and Emma Brunskill. Unifying PAC and regret: Uniform PAC bounds for episodic reinforcement learning. *Advances in Neural Information Processing Systems*, 30, 2017.
- Marisa P De Brito and Erwin A Van Der Laan. Inventory control with product returns: The impact of imperfect information. *European journal of operational research*, 194(1):85–101, 2009.
- Simon Du, Akshay Krishnamurthy, Nan Jiang, Alekh Agarwal, Miroslav Dudik, and John Langford. Provably efficient RL with rich observations via latent state decoding. In *International Conference on Machine Learning*, pages 1665–1674. PMLR, 2019.
- Simon Du, Sham Kakade, Jason Lee, Shachar Lovett, Gaurav Mahajan, Wen Sun, and Ruosong Wang. Bilinear classes: A structural framework for provable generalization in RL. In *International Conference on Machine Learning*, pages 2826–2836. PMLR, 2021.

- Yonathan Efroni, Chi Jin, Akshay Krishnamurthy, and Sobhan Miryoosefi. Provable reinforcement learning with a short-term memory. *arXiv preprint arXiv:2202.03983*, 2022.
- Eyal Even-Dar, Sham M Kakade, and Yishay Mansour. Reinforcement learning in POMDPs without resets. 2005.
- Dylan J Foster, Sham M Kakade, Jian Qian, and Alexander Rakhlin. The statistical complexity of interactive decision making. *arXiv preprint arXiv:2112.13487*, 2021.
- Sara A Geer, Sara van de Geer, and D Williams. *Empirical Processes in M-estimation*, volume 6. Cambridge University Press, 2000.
- Noah Golowich, Ankur Moitra, and Dhruv Rohatgi. Planning in observable POMDPs in quasipolynomial time. *arXiv preprint arXiv:2201.04735*, 2022.
- Zhaohan Daniel Guo, Shayan Doroudi, and Emma Brunskill. A PAC RL algorithm for episodic POMDPs. In *Artificial Intelligence and Statistics*, pages 510–518. PMLR, 2016.
- Milos Hauskrecht and Hamish Fraser. Planning treatment of ischemic heart disease with partially observable Markov decision processes. *Artificial Intelligence in Medicine*, 18(3):221–244, 2000.
- Daniel Hsu, Sham M Kakade, and Tong Zhang. A spectral algorithm for learning hidden Markov models. *Journal of Computer and System Sciences*, 78(5):1460–1480, 2012.
- Herbert Jaeger. Observable operator models for discrete stochastic time series. *Neural computation*, 12(6):1371–1398, 2000.
- Mehdi Jafarnia-Jahromi, Rahul Jain, and Ashutosh Nayyar. Online learning for unknown partially observable mdps. *arXiv preprint arXiv:2102.12661*, 2021.
- Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, John Langford, and Robert E Schapire. Contextual decision processes with low Bellman rank are PAC-learnable. In *International Conference on Machine Learning*, pages 1704–1713. PMLR, 2017.
- Chi Jin, Zeyuan Allen-Zhu, Sebastien Bubeck, and Michael I Jordan. Is Q-learning provably efficient? *Advances in neural information processing systems*, 31, 2018.
- Chi Jin, Sham M Kakade, Akshay Krishnamurthy, and Qinghua Liu. Sample-efficient reinforcement learning of undercomplete POMDPs. *NeurIPS*, 2020a.
- Chi Jin, Zhuoran Yang, Zhaoran Wang, and Michael I Jordan. Provably efficient reinforcement learning with linear function approximation. In *Conference on Learning Theory*, pages 2137–2143. PMLR, 2020b.
- Chi Jin, Qinghua Liu, and Sobhan Miryoosefi. Bellman eluder dimension: New rich classes of RL problems, and sample-efficient algorithms. *Advances in Neural Information Processing Systems*, 34, 2021.
- Akshay Krishnamurthy, Alekh Agarwal, and John Langford. PAC reinforcement learning with rich observations. *Advances in Neural Information Processing Systems*, 29, 2016.



- P Kumar and A Becker. A new family of optimal adaptive controllers for Markov chains. *IEEE Transactions on Automatic Control*, 27(1):137–146, 1982.
- Jeongyeol Kwon, Yonathan Efroni, Constantine Caramanis, and Shie Mannor. Reinforcement learning in reward-mixing MDPs. *Advances in Neural Information Processing Systems*, 34, 2021a.
- Jeongyeol Kwon, Yonathan Efroni, Constantine Caramanis, and Shie Mannor. RL for latent MDPs: Regret guarantees and a lower bound. *Advances in Neural Information Processing Systems*, 34, 2021b.
- Tze Leung Lai. Adaptive treatment allocation and the multi-armed bandit problem. *The Annals of Statistics*, pages 1091–1114, 1987.
- Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- Jesse Levinson, Jake Askeland, Jan Becker, Jennifer Dolson, David Held, Soeren Kammel, J Zico Kolter, Dirk Langer, Oliver Pink, Vaughan Pratt, et al. Towards fully autonomous driving: Systems and algorithms. In *2011 IEEE intelligent vehicles symposium (IV)*, pages 163–168. IEEE, 2011.
- Lihong Li, Wei Chu, John Langford, and Robert E Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th international conference on World wide web*, pages 661–670, 2010.
- Akshay Mete, Rahul Singh, Xi Liu, and PR Kumar. Reward biased maximum likelihood estimation for reinforcement learning. In *Learning for Dynamics and Control*, pages 815–827. PMLR, 2021.
- Dipendra Misra, Mikael Henaff, Akshay Krishnamurthy, and John Langford. Kinematic state abstraction and provably efficient rich-observation reinforcement learning. In *International conference on machine learning*, pages 6961–6971. PMLR, 2020.
- Elchanan Mossel and Sébastien Roch. Learning nonsingular phylogenies and hidden Markov models. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 366–375, 2005.
- Martin Mundhenk, Judy Goldsmith, Christopher Lusena, and Eric Allender. Complexity of finite-horizon Markov decision process problems. *Journal of the ACM (JACM)*, 47(4):681–720, 2000.
- Christos H Papadimitriou and John N Tsitsiklis. The complexity of Markov decision processes. *Mathematics of operations research*, 12(3):441–450, 1987.
- Pascal Poupart and Nikos Vlassis. Model-based Bayesian reinforcement learning in partially observable domains. In *Proc Int. Symp. on Artificial Intelligence and Mathematics*, pages 1–2, 2008.
- Stephane Ross, Brahim Chaib-draa, and Joelle Pineau. Bayes-adaptive POMDPs. *Advances in neural information processing systems*, 20, 2007.
- Daniel Russo and Benjamin Van Roy. Eluder dimension and the sample complexity of optimistic exploration. In *NIPS*, pages 2256–2264. Citeseer, 2013.

- Masatoshi Uehara, Xuezhou Zhang, and Wen Sun. Representation learning for online and offline RL in low-rank MDPs. In *International Conference on Learning Representations*, 2021.
- Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michael Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.
- Nikos Vlassis, Michael L Littman, and David Barber. On the computational complexity of stochastic controller optimization in POMDPs. *ACM Transactions on Computation Theory (TOCT)*, 4(4): 1–8, 2012.
- Samuel S Wilks. The large-sample distribution of the likelihood ratio for testing composite hypotheses. *The annals of mathematical statistics*, 9(1):60–62, 1938.
- Yi Xiong, Ningyuan Chen, Xuefeng Gao, and Xiang Zhou. Sublinear regret for learning POMDPs. *arXiv preprint arXiv:2107.03635*, 2021.
- Andrea Zanette, Alessandro Lazaric, Mykel Kochenderfer, and Emma Brunskill. Learning near optimal policies with low inherent Bellman error. In *International Conference on Machine Learning*, pages 10978–10989. PMLR, 2020.
- Tong Zhang. From  $\varepsilon$ -entropy to KL-entropy: Analysis of minimum information complexity density estimation. *The Annals of Statistics*, 34(5):2180–2210, 2006.

## Appendix A. Related works

Reinforcement learning has been extensively studied in the fully observable setting (see, e.g., [Azar et al., 2017](#); [Dann et al., 2017](#); [Jin et al., 2018, 2020b](#); [Zanette et al., 2020](#); [Jiang et al., 2017](#), and the references therein). For the purpose of this paper, we focus our attention on reviewing the theoretical results for partially observable reinforcement learning.

**Hardness of learning POMDPs.** There is a line of well-known computational hardness results for planning and learning in POMDPs. Firstly, even when the parameters of a POMDP are known, computing the optimal policy (i.e., planning) is PSPACE-complete ([Papadimitriou and Tsitsiklis, 1987](#)). Moreover, even if one only wants to find the optimal memoryless policy, the problem is still NP-hard ([Vlassis et al., 2012](#)). In addition, the model estimation of POMDPs is also computationally hard—[Mossel and Roch \(2005\)](#) proved an average-case computational result showing that estimating the model parameters for a subclass of Hidden Markov Models (HMMs) is at least as hard as learning parity with noise<sup>4</sup>. Since HMMs can be viewed as special cases of POMDPs without action control, their result directly implies estimating the model parameters of POMDPs is hard.

Learning POMDPs is also known to be statistically hard: [Krishnamurthy et al. \(2016\)](#) proved that finding a near-optimal policy of a POMDP in the worst case requires a number of samples that is exponential in the episode length. The hard instances are those pathological POMDPs where the observations contain no useful information for identifying the system dynamics.

---

4. Learning parity with noise is conjectured to be NP-hard in the theory of computational complexity.

**Positive results for learning POMDPs.** Despite the worst-case hardness results, there is a long history of learning sub-classes of POMDPs. [Even-Dar et al. \(2005\)](#) studied POMDPs without resets, where the proposed algorithm has sample complexity scaling exponentially with a certain horizon time. [Poupart and Vlassis \(2008\)](#); [Ross et al. \(2007\)](#) developed Bayesian methods to learn POMDPs, while [Azizzadenesheli et al. \(2018\)](#) considered learning the optimal memoryless policies with policy gradient methods. PAC or regret bounds are not known for these approaches.

In the category of polynomial sample results, a sequence of recent works (e.g., [Guo et al., 2016](#); [Azizzadenesheli et al., 2016](#); [Jin et al., 2020a](#); [Xiong et al., 2021](#)) applied spectral methods ([Hsu et al., 2012](#); [Anandkumar et al., 2014](#)) to learning POMDPs and obtained polynomial sample complexity results. Among them, [Guo et al. \(2016\)](#); [Azizzadenesheli et al. \(2016\)](#); [Xiong et al. \(2021\)](#) made strong reachability assumptions and did not address the exploration problem. Furthermore, these results assume that both the transition and emission matrices are full rank, which are stronger than the weakly revealing conditions considered in this paper. [Jafarnia-Jahromi et al. \(2021\)](#) proposed a posterior sampling-based algorithm, and provided sample-efficient guarantees *assuming* either sufficient separability between different models, or the success of belief state and transition kernel estimation. These assumptions significantly reduce the difficulty of estimating model dynamics—a core challenge in learning POMDPs, and thus reduce the generality of the results.

The most related work to us is [Jin et al. \(2020a\)](#), which addressed the exploration problem in learning undercomplete POMDPs, where the number of latent states must be no greater than the number of observations. Their algorithm is specially designed to exploit the undercomplete structure of POMDPs. It remains unclear if their techniques can be extended to the overcomplete setting. In contrast, this paper presents a new generic algorithm based on MLE, which enjoys provable sample-efficiency in the exploration settings of both undercomplete and overcomplete POMDPs.

Very recently, [Golowich et al. \(2022\)](#) developed the first quasi-polynomial time planning algorithm for a subclass of POMDPs. Their result holds under the  $\gamma$ -observability condition, which is very similar to the weakly-revealing condition presented in this paper.<sup>5</sup> Compared to this paper, the result in [Golowich et al. \(2022\)](#) purely focuses on the computational efficiency. It is restricted to the undercomplete setting, and addresses only planning but not estimation or exploration, all of which are important components for learning POMDPs.

**Latent MDPs.** Latent MDPs ([Kwon et al., 2021b](#))—where an MDP is randomly drawn from a set of  $M$  possible MDPs at the beginning of the interaction—can be considered as a special class of overcomplete POMDPs. [Kwon et al. \(2021b\)](#) proved that learning latent MDPs remains statistically hard in the worst case. They also provided several positive results for learning latent MDPs with additional assumptions, such as revealing the latent contexts at the end of each episode. [Kwon et al. \(2021a\)](#) provided positive results for latent MDPs without these additional assumptions, but the results only apply to the setting of  $M = 2$  with a shared transition. Latent MDPs and weakly revealing POMDPs do not contain each other.

**Decodable POMDPs.** Block MDPs ([Krishnamurthy et al., 2016](#)) are POMDPs whose current latent state can be uniquely determined by the current observation. By simple algebra, one can verify that block MDPs are special cases of single-step weakly revealing POMDPs that satisfy Assumption 1 with  $\alpha \geq 1/\sqrt{O}$ . The recently proposed  $m$ -step decodable POMDPs ([Efroni et al.,](#)

5.  $\gamma$ -observability condition requires that  $\min_h \|\mathbb{O}_h(b - b')\|_1 \geq \gamma \|b - b'\|_1$  for any  $b, b' \in \Delta_S$ , while  $\alpha$ -weakly-revealing condition (Assumption 1) assumes  $\min_h \|\mathbb{O}_h x\|_2 \geq \alpha \|x\|_2$  for any  $x \in \mathbb{R}^O$ . In Lemma 35, we prove that  $\frac{\alpha}{\sqrt{S}} \leq \gamma \leq 4\sqrt{O}\alpha$ . As a result, these two conditions are “equivalent” up to a factor of at most  $\mathcal{O}(\sqrt{O})$ .

2022) are generalizations of block MDPs, in which the latent state can be uniquely decoded from the most recent history (of observations and actions) of a short length  $m$ . This multistep decodability assumption can also be viewed as a special case of general “weakly revealing”-type of conditions. However, Efroni et al. (2022) assume that  $m$ -step history decodes (weakly reveals) the current state, while this paper assumes that  $m$ -step future weakly reveals the current state. Finally, we remark that most existing results for block MDPs or  $m$ -step decodable POMDPs (see, e.g., Krishnamurthy et al., 2016; Jiang et al., 2017; Du et al., 2019; Misra et al., 2020; Efroni et al., 2022) further involve decoder class or value function approximation, which is beyond the scope of this paper.

**RL with function approximation.** There is a recent line of research (e.g., Jiang et al., 2017; Ayoub et al., 2020; Du et al., 2021; Jin et al., 2021; Foster et al., 2021) on reinforcement learning with general function approximation. This line of results proposed certain complexity measure for sequential decision making problems, and developed generic algorithms which have sample-efficient guarantees as long as the complexity measure of RL problems is small. These frameworks are known to cover a special subclass of POMDPs—reactive POMDPs (Jiang et al., 2017), where the optimal value only depends on the current-step observation-action pair. It remains highly unclear whether weakly revealing POMDPs identified in this paper can be covered by those general frameworks. We remark that investigating this problem requires us to compute those complexity measures for POMDPs, which is highly non-trivial and may require techniques developed in this paper.

**MLE approaches in bandit and RL.** The idea of using the MLE principle in the confidence set construction can be traced back to Lai (1987), which considers the problem of Bernoulli bandits. MLE-based approaches are also used in the framework of reward-biased MLE (Kumar and Becker, 1982; Mete et al., 2021), which balances the reward with the likelihood value, for learning tabular MDPs. Recently, the MLE-based approaches are also used in the setting of representation learning in RL (Agarwal et al., 2020; Uehara et al., 2021).

## Appendix B. Maximum Likelihood Estimation

In this section, we analyze the maximum likelihood estimation (MLE) approach for the following meta-algorithm. Since Algorithm 1 and 2 can be viewed as special cases of Algorithm 3, all the results developed in this section directly apply to their analysis.

---

### Algorithm 3 Meta-algorithm

---

**for**  $t = 1, \dots, T$  **do**  
  choose policy  $\pi^t$  as a deterministic function of  $\{(\pi^i, \tau^i)\}_{i=1}^{t-1}$   
  execute policy  $\pi^t$  and collect a trajectory  $\tau^t$

---

For the reader’s convenience, we recall the definitions of the following notations: (a)  $\theta = (\mathbb{T}, \mathbb{O}, \mu)$  denotes the ensemble of all the parameters of a POMDP model, (b)  $\Theta$  denotes the collections of all such POMDP parameter ensembles, and (c)  $\theta^*$  denotes the parameter ensemble of the groundtruth POMDP that we are interacting with. We will view  $\Theta$  as a subset of a Euclidean space; in particular,  $\Theta \subset \mathbb{R}^{H(S^2A+SO)+S}$ , so  $\theta, \theta^* \in \mathbb{R}^{H(S^2A+SO)+S}$ .

The first proposition shows that, up to certain error, with high probability, the log-likelihood of the groundtruth model computed using the historical data is close to the maximum log-likelihood.

**Proposition 13** *There exists an absolute constant  $c$  such that for any  $\delta \in (0, 1]$ , with probability at least  $1 - \delta$ : the following inequality holds for all  $t \in [T]$  and **all**  $\theta \in \Theta$*

$$\sum_{i=1}^t \log \left( \frac{\mathbb{P}_{\theta}^{\pi^i}(\tau^i)}{\mathbb{P}_{\theta^*}^{\pi^i}(\tau^i)} \right) \leq c (H(S^2 A + SO) \log(TSAOH) + \log(T/\delta)). \quad (12)$$

Our second claim shows that any POMDP model, whose log-likelihood on the historical data is comparable to that of the groundtruth model, will produce similar distributions of trajectories as the groundtruth model under historical polices.

**Proposition 14** *There exists a universal constant  $c$  such that for any  $\delta \in (0, 1]$ , with probability at least  $1 - \delta$  for all  $t \in [T]$  and **all**  $\theta \in \Theta$ , it holds that*

$$\begin{aligned} & \sum_{i=1}^t \left( \sum_{\tau \in (\mathcal{O} \times \mathcal{A})^H} \left| \mathbb{P}_{\theta}^{\pi^i}(\tau) - \mathbb{P}_{\theta^*}^{\pi^i}(\tau) \right| \right)^2 \\ & \leq c \left( \sum_{i=1}^t \log \left( \frac{\mathbb{P}_{\theta^*}^{\pi^i}(\tau^i)}{\mathbb{P}_{\theta}^{\pi^i}(\tau^i)} \right) + H(S^2 A + SO) \log(TSAOH) + \log(T/\delta) \right). \end{aligned} \quad (13)$$

The proofs of Proposition 13 and 14 can be found in Appendix C.

## Appendix C. Proofs for Maximum Likelihood Estimation

In this section, we prove the two propositions stated in Appendix B. For technical purposes, we introduce the concept of optimistic  $\varepsilon$ -discretization. Specifically, we denote  $\bar{\theta}$  as the optimistic  $\varepsilon$ -discretization of  $\theta$  so that  $\bar{\theta}_i = \lceil \theta_i / \varepsilon \rceil \times \varepsilon$  for all coordinate  $i$ . We comment that although  $\bar{\theta}$  is not a legal POMDP parameterization, it can still be used to compute the probability of observing any trajectory  $\tau$  under any policy  $\pi$  by simply replacing  $\theta_i$  with  $\bar{\theta}_i$  in the computation. In particular, for  $\theta = (\mathbb{T}, \mathbb{O}, \mu_1) \in \mathbb{R}^d$  with  $d = H(S^2 A + SO) + S$ , given a trajectory  $\tau = (o_1, a_1, \dots, o_H, a_H)$  and a policy  $\pi$  we define

$$\begin{aligned} \mathbb{P}_{\bar{\theta}}^{\pi}(\tau) &= \sum_{s_1, \dots, s_H \in \mathcal{S}} \mu_1(s_1) \mathbb{O}_1(o_1 | s_1) \pi_1(a_1 | o_1) \mathbb{T}_{1, a_1}(s_2 | s_1) \times \\ & \dots \\ & \times \mathbb{O}_1(o_{H-1} | s_{H-1}) \pi_1(a_{H-1} | o_1, a_1, \dots, o_{H-2}, a_{H-2}, o_{H-1}) \mathbb{T}_{H-1, a_{H-1}}(s_H | s_{H-1}) \\ & \times \mathbb{O}_1(o_H | s_H) \pi_1(a_H | o_1, a_1, \dots, o_{H-1}, a_{H-1}, o_H). \end{aligned}$$

Note that the right-hand side gives a probability measure over the trajectories when  $\theta \in \Theta$ , but not in general. To simplify the language, we will still call  $\mathbb{P}_{\bar{\theta}}^{\pi}(\tau)$  the ‘‘probability’’ of  $\tau$  no matter whether this is indeed a probability. The significance of using optimistic discretization is that for any  $\theta \in \Theta$  and trajectory  $\tau$ , we always have that

$$\mathbb{P}_{\bar{\theta}}^{\pi}(\tau) \geq \mathbb{P}_{\theta}^{\pi}(\tau).$$

We denote by  $\bar{\Theta}$  the collections of all such  $\bar{\theta}$ , i.e.,  $\bar{\Theta} = \{\bar{\theta} : \theta \in \Theta\}$ . Throughout this section, we will choose a fixed  $\varepsilon$  satisfying that for any  $\theta \in \Theta$  and any policy  $\pi$ ,

$$\|\mathbb{P}_{\theta}^{\pi} - \mathbb{P}_{\bar{\theta}}^{\pi}\|_1 \leq 1/T. \quad (14)$$

It is not hard to see that one can choose  $\varepsilon \leq 1/(C(S+O+ A)HT)$  for this condition to be satisfied where  $C$  is some large absolute constant.

Since  $\bar{\theta}$  belongs to  $[0, 1]^d$  with  $d = S + H(S^2A + SO)$ , the log-cardinality of  $\bar{\Theta}$  defined by this particular  $\varepsilon$  is at most  $\mathcal{O}(H(S^2A + SO) \log(TSAOH))$ .

### C.1. Proof of Proposition 13

**Proof** The proof is rather standard (e.g., see Geer et al., 2000) and uses Cramér-Chernoff's method. Pick any  $\bar{\theta} \in \bar{\Theta}$  and  $t \in [T]$ . Denote  $\mathbb{E}_t[\cdot] = \mathbb{E}[\cdot \mid \{(\pi^i, \tau^i)\}_{i=1}^{t-1} \cup \{\pi^t\}]$ . We have

$$\begin{aligned} & \mathbb{E} \left[ \exp \left( \sum_{i=1}^t \log \left( \frac{\mathbb{P}_{\bar{\theta}}^{\pi^i}(\tau^i)}{\mathbb{P}_{\theta^*}^{\pi^i}(\tau^i)} \right) \right) \right] \\ &= \mathbb{E} \left[ \exp \left( \sum_{i=1}^{t-1} \log \left( \frac{\mathbb{P}_{\bar{\theta}}^{\pi^i}(\tau^i)}{\mathbb{P}_{\theta^*}^{\pi^i}(\tau^i)} \right) \right) \cdot \mathbb{E}_t \left[ \exp \left( \log \left( \frac{\mathbb{P}_{\bar{\theta}}^{\pi^t}(\tau^t)}{\mathbb{P}_{\theta^*}^{\pi^t}(\tau^t)} \right) \right) \right] \right] \\ &= \mathbb{E} \left[ \exp \left( \sum_{i=1}^{t-1} \log \left( \frac{\mathbb{P}_{\bar{\theta}}^{\pi^i}(\tau^i)}{\mathbb{P}_{\theta^*}^{\pi^i}(\tau^i)} \right) \right) \cdot \mathbb{E}_t \left[ \frac{\mathbb{P}_{\bar{\theta}}^{\pi^t}(\tau^t)}{\mathbb{P}_{\theta^*}^{\pi^t}(\tau^t)} \right] \right] \\ &= \mathbb{E} \left[ \exp \left( \sum_{i=1}^{t-1} \log \left( \frac{\mathbb{P}_{\bar{\theta}}^{\pi^i}(\tau^i)}{\mathbb{P}_{\theta^*}^{\pi^i}(\tau^i)} \right) \right) \cdot \|\mathbb{P}_{\bar{\theta}}^{\pi^t}(\tau = \cdot)\|_1 \right] \\ &\leq \mathbb{E} \left[ \exp \left( \sum_{i=1}^{t-1} \log \left( \frac{\mathbb{P}_{\bar{\theta}}^{\pi^i}(\tau^i)}{\mathbb{P}_{\theta^*}^{\pi^i}(\tau^i)} \right) \right) \cdot \left( 1 + \frac{1}{T} \right) \right] \leq \dots \leq e, \end{aligned}$$

where the first inequality follows from (14). Therefore, by Markov's inequality, we have

$$\begin{aligned} \mathbb{P} \left( \sum_{i=1}^t \log \left( \frac{\mathbb{P}_{\bar{\theta}}^{\pi^i}(\tau^i)}{\mathbb{P}_{\theta^*}^{\pi^i}(\tau^i)} \right) > \log(1/\delta) \right) &\leq \mathbb{E} \left[ \exp \left( \sum_{i=1}^t \log \left( \frac{\mathbb{P}_{\bar{\theta}}^{\pi^i}(\tau^i)}{\mathbb{P}_{\theta^*}^{\pi^i}(\tau^i)} \right) \right) \right] \cdot \exp[-\log(1/\delta)] \\ &= e\delta. \end{aligned}$$

Taking a union bound for all  $(\bar{\theta}, t) \in \bar{\Theta} \times [T]$  and rescaling  $\delta$ , we obtain

$$\mathbb{P} \left( \max_{(\bar{\theta}, t) \in \bar{\Theta} \times [T]} \sum_{i=1}^t \log \left( \frac{\mathbb{P}_{\bar{\theta}}^{\pi^i}(\tau^i)}{\mathbb{P}_{\theta^*}^{\pi^i}(\tau^i)} \right) > c(H(S^2A + SO) \log(TSAOH) + \log(T/\delta)) \right) \leq \delta,$$

where  $c > 0$  is some absolute constant. Finally, recall  $\bar{\theta}$  is an optimistic discretization of  $\theta$ , which implies  $\mathbb{P}_{\bar{\theta}}^{\pi}(\tau) \leq \mathbb{P}_{\theta}^{\pi}(\tau)$  for all  $\theta, \pi, y$ . As a result, we conclude that

$$\mathbb{P} \left( \max_{(\theta, t) \in \Theta \times [T]} \sum_{i=1}^t \log \left( \frac{\mathbb{P}_{\bar{\theta}}^{\pi^i}(\tau^i)}{\mathbb{P}_{\theta^*}^{\pi^i}(\tau^i)} \right) > c(H(S^2A + SO) \log(TSAOH) + \log(T/\delta)) \right) \leq \delta. \quad \blacksquare$$

### C.2. Proof of Proposition 14

The proof in this section largely follows Agarwal et al. (2020), which is inspired by Zhang (2006).

We start by recalling the meta-algorithm (Algorithm 3): in each iteration  $t \in [T]$ , we pick  $\pi^t$  deterministically based on  $\{(\pi^i, \tau^i)\}_{i=1}^{t-1}$  and sample  $\tau^t$  from  $\mathbb{P}_{\theta^*}^{\pi^t}$ . Now consider a fixed  $t \in [T]$ ; let  $D = \{(\pi^i, \tau^i)\}_{i=1}^t$  denote the sequence of policy-trajectory pairs observed within the first  $t$  iterations, and denote by  $\tilde{D}$  a ‘‘tangent’’ sequence  $\{(\pi^i, \tilde{\tau}^i)\}_{i=1}^t$  where  $\tilde{\tau}^i \sim \mathbb{P}_{\theta^*}^{\pi^i}$ . Note that  $D$  and  $\tilde{D}$  share the same policy parts and their trajectories are independently sampled from the same distributions.

**Lemma 15** *Let  $\ell := \ell(\pi, \tau)$  be a real-valued function that maps a policy  $\pi$  and a trajectory  $\tau$  to  $\mathbb{R}$ . Let  $L(D) = \sum_{i=1}^t \ell(\pi^i, \tau^i)$ , and  $L(\tilde{D}) = \sum_{i=1}^t \ell(\pi^i, \tilde{\tau}^i)$ . Then,*

$$\mathbb{E} \left[ \exp \left( L(D) - \log \mathbb{E}[\exp(L(\tilde{D})) \mid D] \right) \right] = 1.$$

**Proof** Before proving the lemma, we remark that all the expectations considered above and below always exist because each  $\tau_i$  and  $\pi_i$  can only take finitely many different values. It is direct to see  $\tau_i$  only has finitely many possibilities since the number of observations and actions are finite. As for  $\pi_i$ , notice that  $\pi_i$  is a deterministic function of  $\{(\pi^j, \tau^j)\}_{j=1}^{i-1}$  in Algorithm 3, so it follows by induction that  $\pi_i$  also only takes finitely many different values.

Now we prove Lemma 15. Define  $E_i = \mathbb{E}[\exp(\ell(\pi^i, \tau^i)) \mid \pi^i]$  and

$$u_t = \mathbb{E} \left[ \exp \left( L(D) - \log \mathbb{E}[\exp(L(\tilde{D})) \mid D] \right) \right] = \mathbb{E} \left[ \frac{\exp \left( \sum_{i=1}^t \ell(\pi^i, \tau^i) \right)}{\mathbb{E}[\exp \left( \sum_{i=1}^t \ell(\pi^i, \tilde{\tau}^i) \right) \mid D]} \right].$$

Our goal is to show that  $u_t = 1$ . Owing to the definition of  $\{\tilde{\tau}_i\}_i$ , some calculation gives

$$\mathbb{E} \left[ \exp \left( \sum_{i=1}^t \ell(\pi^i, \tilde{\tau}^i) \right) \mid D \right] = \prod_{i=1}^t E_i.$$

Plugging this into the last expression obtained for  $u_t$  and using the tower rule with

$$D_{t-1} = (\pi^1, \tau^1, \dots, \pi^{t-1}, \tau^{t-1}, \pi^t)$$

we get

$$u_t = \mathbb{E} \left[ \mathbb{E} \left[ \frac{\exp \left( \sum_{i=1}^t \ell(\pi^i, \tau^i) \right)}{\prod_{i=1}^t E_i} \mid D_{t-1} \right] \right] = \mathbb{E} \left[ \frac{\mathbb{E}[\exp \left( \sum_{i=1}^t \ell(\pi^i, \tau^i) \right) \mid D_{t-1}]}{\prod_{i=1}^t E_i} \right],$$

where we used that  $E_1, \dots, E_t$  are  $\sigma(D_{t-1})$ -measurable. Now,

$$\mathbb{E} \left[ \exp \left( \sum_{i=1}^t \ell(\pi^i, \tau^i) \right) \mid D_{t-1} \right] = \mathbb{E} \left[ \exp \left( \sum_{i=1}^{t-1} \ell(\pi^i, \tau^i) \right) \mid D_{t-1} \right] E_t.$$

Plugging this back into our previous expression,

$$\begin{aligned}
 u_t &= \mathbb{E} \left[ \frac{\mathbb{E} \left[ \exp(\sum_{i=1}^{t-1} \ell(\pi^i, \tau^i)) \mid D_{t-1} \right] \cancel{E_t}}{\left( \prod_{i=1}^{t-1} E_i \right) \cancel{E_t}} \right] = \mathbb{E} \left[ \mathbb{E} \left[ \frac{\exp(\sum_{i=1}^{t-1} \ell(\pi^i, \tau^i))}{\prod_{i=1}^{t-1} E_i} \mid D_{t-1} \right] \right] \\
 &= \mathbb{E} \left[ \frac{\exp(\sum_{i=1}^{t-1} \ell(\pi^i, \tau^i))}{\prod_{i=1}^{t-1} E_i} \right] \quad (\text{by the tower rule}) \\
 &= u_{t-1} = \dots = u_0 = 1,
 \end{aligned}$$

where the second equality used that  $E_1, \dots, E_{t-1}$  is  $\sigma(D_{t-1})$ -measurable.  $\blacksquare$

For any  $\bar{\theta} \in \bar{\Theta}$ , we define

$$\ell_{\bar{\theta}}(\pi, \tau) := \begin{cases} \frac{1}{2} \log \left( \frac{\mathbb{P}_{\bar{\theta}}^{\pi}(\tau)}{\mathbb{P}_{\theta^*}^{\pi}(\tau)} \right), & \mathbb{P}_{\theta^*}^{\pi}(\tau) \neq 0, \\ 0, & \text{otherwise,} \end{cases}$$

and

$$L_{\bar{\theta}}(D) = \sum_{i=1}^t \ell_{\bar{\theta}}(\pi^i, \tau^i), \quad L_{\bar{\theta}}(\tilde{D}) = \sum_{i=1}^t \ell_{\bar{\theta}}(\pi^i, \tilde{\tau}^i).$$

By Lemma 15, Chernoff's method and the union bound, with probability at least  $1 - \delta$ , for all  $\bar{\theta} \in \bar{\Theta}$  we have

$$\begin{aligned}
 -\log \mathbb{E}_{\tilde{D}}[\exp(L_{\bar{\theta}}(\tilde{D})) \mid D] &< -L_{\bar{\theta}}(D) + \log(|\bar{\Theta}|/\delta) \\
 &\leq -L_{\bar{\theta}}(D) + \mathcal{O}(H(S^2 A + SO)) \log(TSAOH) + \log(1/\delta),
 \end{aligned}$$

Then, by the definition of  $L_{\bar{\theta}}$  and by the inequality  $-\log x \geq 1 - x$ ,

$$-\log \mathbb{E}_{\tilde{D}}[\exp(L_{\bar{\theta}}(\tilde{D})) \mid D] = -\sum_{i=1}^t \log \mathbb{E}_{\tau \sim \mathbb{P}_{\theta^*}^{\pi^i}} \left[ \sqrt{\frac{\mathbb{P}_{\bar{\theta}}^{\pi^i}(\tau)}{\mathbb{P}_{\theta^*}^{\pi^i}(\tau)}} \right] \geq \sum_{i=1}^t \left( 1 - \mathbb{E}_{\tau \sim \mathbb{P}_{\theta^*}^{\pi^i}} \left[ \sqrt{\frac{\mathbb{P}_{\bar{\theta}}^{\pi^i}(\tau)}{\mathbb{P}_{\theta^*}^{\pi^i}(\tau)}} \right] \right).$$

Recall that for any  $\theta \in \Theta$  and any policy  $\pi$ ,  $\|\mathbb{P}_{\theta}^{\pi} - \mathbb{P}_{\theta^*}^{\pi}\|_1 \leq 1/T$ . As a result, by algebra,

$$\begin{aligned}
 &\sum_{i=1}^t \left( 1 - \mathbb{E}_{\tau \sim \mathbb{P}_{\theta^*}^{\pi^i}} \left[ \sqrt{\frac{\mathbb{P}_{\bar{\theta}}^{\pi^i}(\tau)}{\mathbb{P}_{\theta^*}^{\pi^i}(\tau)}} \right] \right) \\
 &= \sum_{i=1}^t \left( 1 - \sum_{\tau} \sqrt{\mathbb{P}_{\theta^*}^{\pi^i}(\tau) \mathbb{P}_{\bar{\theta}}^{\pi^i}(\tau)} \right) \\
 &\geq \frac{1}{2} \sum_{i=1}^t \sum_{\tau} \left( \sqrt{\mathbb{P}_{\bar{\theta}}^{\pi^i}(\tau)} - \sqrt{\mathbb{P}_{\theta^*}^{\pi^i}(\tau)} \right)^2 - \frac{1}{2} \\
 &\geq \frac{1}{12} \sum_{i=1}^t \left[ \sum_{\tau} \left( \sqrt{\mathbb{P}_{\bar{\theta}}^{\pi^i}(\tau)} - \sqrt{\mathbb{P}_{\theta^*}^{\pi^i}(\tau)} \right)^2 \right] \left[ \sum_{\tau} \left( \sqrt{\mathbb{P}_{\bar{\theta}}^{\pi^i}(\tau)} + \sqrt{\mathbb{P}_{\theta^*}^{\pi^i}(\tau)} \right)^2 \right] - \frac{1}{2} \\
 &\geq \frac{1}{12} \sum_{i=1}^t \left( \sum_{\tau} |\mathbb{P}_{\bar{\theta}}^{\pi^i}(\tau) - \mathbb{P}_{\theta^*}^{\pi^i}(\tau)| \right)^2 - \frac{1}{2},
 \end{aligned}$$



where the last inequality follows from the Cauchy-Schwarz inequality. Putting all relations together, we have that with probability at least  $1 - \delta$ , for all  $\bar{\theta} \in \bar{\Theta}$

$$-L_{\bar{\theta}}(D) + \mathcal{O}(H(S^2 A + SO)) \log(TSAOH) + \log(1/\delta) \geq \frac{1}{12} \sum_{i=1}^t \left( \sum_{\tau} |\mathbb{P}_{\bar{\theta}}^{\pi^i}(\tau) - \mathbb{P}_{\theta^*}^{\pi^i}(\tau)| \right)^2 - \frac{1}{2}.$$

Finally, notice that for all  $\theta \in \Theta$ , we have

$$-L_{\bar{\theta}}(D) \leq -L_{\theta}(D).$$

Then, repeatedly using that  $\|\mathbb{P}_{\theta}^{\pi} - \mathbb{P}_{\theta^*}^{\pi}\|_1 \leq 1/T$  and  $\|\mathbb{P}_{\bar{\theta}}^{\pi} - \mathbb{P}_{\theta^*}^{\pi}\|_1 \leq 2 + 1/T$  for any policy  $\pi$ , we get

$$\sum_{i=1}^t \left( \sum_{\tau} |\mathbb{P}_{\bar{\theta}}^{\pi^i}(\tau) - \mathbb{P}_{\theta^*}^{\pi^i}(\tau)| \right)^2 \geq \sum_{i=1}^t \left( \sum_{\tau} |\mathbb{P}_{\theta}^{\pi^i}(\tau) - \mathbb{P}_{\theta^*}^{\pi^i}(\tau)| \right)^2 - 6.$$

We conclude that with probability at least  $1 - \delta$ , for all  $\theta \in \Theta$

$$-L(\theta, D) + \mathcal{O}(H(S^2 A + SO)) \log(SAOHT) + \log(1/\delta) \geq \frac{1}{12} \sum_{i=1}^t \left( \sum_{\tau} |\mathbb{P}_{\theta}^{\pi^i}(\tau) - \mathbb{P}_{\theta^*}^{\pi^i}(\tau)| \right)^2 - 1.$$

Taking a union bound for all  $t \in [T]$  completes the proof.

## Appendix D. $\ell_1$ -norm eluder Dimension

In this section, we introduce the framework of  $\ell_1$ -norm eluder dimension, and present the corresponding pigeonhole-style regret guarantee. All the proofs for this section are deferred to Appendix E.

### D.1. Definitions and properties of $\ell_1$ -norm eluder Dimension

To begin with, we define the  $\varepsilon$ -independence relation between a point and a set of points with respect to a function class under the  $\ell_1$ -norm.

**Definition 16 ( $\ell_1$ -norm  $\varepsilon$ -independence)** Let  $\mathcal{F}$  be a function class defined on  $\mathcal{X}$ , and let

$$z, x_1, x_2, \dots, x_n \in \mathcal{X}.$$

We say  $z$  is  $\varepsilon$ -independent of  $\{x_1, x_2, \dots, x_n\}$  with respect to  $\mathcal{F}$  if there exists  $f \in \mathcal{F}$  such that  $\sum_{i=1}^n |f(x_i)| \leq \varepsilon$ , but  $|f(z)| > \varepsilon$ .

When  $\mathcal{F}$  is clear from the context, we drop ‘‘with respect to  $\mathcal{F}$ ’’ for brevity. We say  $z$  is  $\varepsilon$ -dependent on  $\{x_1, x_2, \dots, x_n\}$  if it is *not*  $\varepsilon$ -independent of  $\{x_1, x_2, \dots, x_n\}$ .

**Definition 17 ( $\ell_1$ -norm  $\varepsilon$ -eluder sequence)** We say that  $\{x_i\}_{i=1}^n \subseteq \mathcal{X}$  is an  $\ell_1$ -norm  $\varepsilon$ -eluder sequence if for all  $i \in [n]$ ,  $x_i$  is  $\varepsilon$ -independent of  $\{x_1, \dots, x_{i-1}\}$ .

**Definition 18 ( $\ell_1$ -norm  $\varepsilon$ -eluder dimension)** Let  $\mathcal{F}$  be a function class defined on  $\mathcal{X}$ . The  $\ell_1$ -norm  $\varepsilon$ -eluder dimension,  $\dim_{\mathbb{E}}(\mathcal{F}, \varepsilon)$ , is the length of the longest  $\ell_1$ -norm  $\varepsilon'$ -eluder sequence with some  $\varepsilon' \geq \varepsilon$ .

If  $\{x_i\}_{i=1}^n \subseteq \mathcal{X}$  is an  $\ell_1$ -norm  $\varepsilon'$ -eluder sequence with  $n = \dim_{\mathbb{E}}(\mathcal{F}, \varepsilon)$  then we say that this sequence is witness to the  $\ell_1$ -norm eluder dimension of  $\mathcal{F}$ .

The difference between the  $\ell_1$ -norm eluder dimension defined here and the original eluder dimension (Russo and Van Roy, 2013) is that in the definition of independence our  $\ell_1$ -version evaluates the  $\ell_1$ -norm of the function on the dataset instead of the  $\ell_2$ -norm as in the original definition. Therefore, we will refer to the original eluder dimension as  $\ell_2$ -norm eluder dimension throughout this paper, while we will drop the “ $\ell_1$ -norm qualifier” whenever it is clear from the context.

**Proposition 19** *The  $\ell_1$ -norm  $\varepsilon$ -eluder dimension is always upper bounded by the  $\ell_2$ -norm  $\varepsilon$ -eluder dimension.*

By combining the upper bound for the  $\ell_2$ -norm eluder dimension of  $d$ -dimensional linear function classes with Proposition 19, we immediately obtain that the  $\ell_1$ -norm eluder dimension of any bounded  $d$ -dimensional linear function class is at most  $\tilde{O}(d)$ .

**Corollary 20** *The  $\ell_1$ -norm  $\varepsilon$ -eluder dimension of*

$$\mathcal{F} = \{f_\theta : f_\theta(x) = \langle x, \theta \rangle, x \in B_{R_2}^d(0), \theta \in B_{R_1}^d(0)\}$$

*is at most  $\tilde{O}(d \log(1 + R_1 R_2 / \varepsilon))$ .*

## D.2. The pigeonhole principle for the $\ell_1$ -norm eluder dimension

Similar to  $\ell_2$ -norm eluder dimension, we can also prove a pigeonhole-style regret guarantee for the  $\ell_1$  version. This will play a key role in deriving the final regret bound from the MLE guarantee.

**Proposition 21** *Let  $\Phi$  be a set of real-valued functions sharing the domain  $\mathcal{X}$  and bounded by  $C > 0$ . Suppose sequence  $\{\phi_k\}_{k=1}^K \subset \Phi$  and  $\{x_k\}_{k=1}^K \subset \mathcal{X}$  satisfy that for all  $k \in [K]$ ,*

$$\sum_{t=1}^{k-1} |\phi_k(x_t)| \leq \beta.$$

*Then for all  $k \in [K]$  and  $\omega > 0$ ,*

$$\sum_{t=1}^k |\phi_t(x_t)| \leq (d+1)C + d\beta \log(C/\omega) + k\omega,$$

*where  $d = \dim_{\mathbb{E}}(\Phi, \omega)$  is the  $\ell_1$ -norm  $\omega$ -eluder dimension*

Finally, we instantiate Proposition 21 on a linear function class with Corollary 20, and obtain the following  $\ell_1$ -norm pigeonhole regret bound.

**Proposition 22** *Suppose  $\{w_{k,j}\}_{(k,j) \in [K] \times [m]}, \{x_{k,i}\}_{(k,i) \in [K] \times [n]} \subset \mathbb{R}^d$  satisfy*

$$\begin{cases} \sum_{t=1}^{k-1} \sum_{i=1}^n \sum_{j=1}^m |w_{k,j}^\top x_{t,i}| \leq \gamma_k \\ \sum_{i=1}^n \|x_{k,i}\|_2 \leq R_x \\ \sum_{j=1}^m \|w_{k,j}\|_2 \leq R_w \end{cases} \quad \text{for all } k \in [K]. \quad (15)$$

*Then we have*

$$\sum_{t=1}^k \sum_{i=1}^n \sum_{j=1}^m |w_{t,j}^\top x_{t,i}| = \mathcal{O}\left(d \left( R_w R_x + \max_{t \leq k} \gamma_t \right) \log^2(Kn)\right) \quad \text{for all } k \in [K]. \quad (16)$$

## Appendix E. Proofs for $\ell_1$ -norm eluder dimension

In this section, we provide the proofs for the results in Appendix D.

### E.1. Proof of Proposition 19

**Proof** [Proof of Proposition 19] Let  $x_1, \dots, x_n$  be an  $\ell_1$ -norm  $\varepsilon$ -independent sequence. By definition, there exist  $f_1, \dots, f_n \in \mathcal{F}$  such that for all  $k \in [n]$

$$\begin{cases} \sum_{i=1}^{k-1} |f_k(x_i)| \leq \varepsilon, \\ |f_k(x_k)| \geq \varepsilon. \end{cases} \quad (17)$$

Since  $\ell_1$ -norm is always an upper bound for the  $\ell_2$ -norm, we also have for all  $k \in [n]$

$$\begin{cases} \sqrt{\sum_{i=1}^{k-1} |f_k(x_i)|^2} \leq \varepsilon, \\ |f_k(x_k)| \geq \varepsilon. \end{cases} \quad (18)$$

Therefore,  $x_1, \dots, x_n$  is also an  $\ell_2$ -norm  $\varepsilon$ -independent sequence. Hence, the  $\ell_2$ -eluder dimension of  $\mathcal{F}$  is at least as large as the  $\ell_1$ -eluder dimension of  $\mathcal{F}$  since any witness of the  $\ell_1$ -eluder dimension at some scale is a witness of the  $\ell_2$ -eluder dimension for the same scale. ■

### E.2. Proof of Proposition 21

The proofs in this subsection follow the arguments developed for the analogous statements for the  $\ell_2$ -norm eluder dimension in Appendix C of Russo and Van Roy (2013). We first prove a few auxiliary claims.

**Claim 1** *Assume that the conditions of Proposition 21 hold. Let  $\varepsilon > 0$ . Then, if for some  $k \in [K]$  we have  $|\phi_k(x_k)| > \varepsilon$ , then  $x_k$  is  $\varepsilon$ -dependent with respect to  $\Phi$  on at most  $\beta/\varepsilon$  disjoint subsequences in  $\{x_1, \dots, x_{k-1}\}$ .*

**Proof** Pick  $k \in [K]$  as in the claim. Let  $\{z_1, \dots, z_\ell\}$  be a subsequence of  $\{x_1, \dots, x_{k-1}\}$  such that  $x_k$  is  $\varepsilon$ -dependent on  $\{z_1, \dots, z_\ell\}$  with respect to  $\Phi$ . Since from  $\sum_{i=1}^{\ell} |\phi_k(z_i)| \leq \varepsilon$  it follows that  $x_k$  is  $\varepsilon$ -independent on  $\{z_1, \dots, z_\ell\}$  with respect to  $\Phi$ , which we assumed not to hold, it follows that  $\sum_{i=1}^{\ell} |\phi_k(z_i)| > \varepsilon$ . This implies that if  $x_k$  is  $\varepsilon$ -dependent on  $L$  disjoint subsequences in  $\{x_1, \dots, x_{k-1}\}$ , we have

$$\beta \geq \sum_{t=1}^{k-1} |\phi_k(x_t)| > L\varepsilon,$$

which results in  $L < \beta/\varepsilon$ . ■

**Claim 2** *Let  $\Phi$  be an arbitrary class of real valued functions sharing the common domain  $\mathcal{X}$ . For any  $\varepsilon > 0$ , sequence  $\{z_1, \dots, z_\kappa\} \subseteq \mathcal{X}$ , there exists  $j \in [\kappa]$  such that  $z_j$  is  $\varepsilon$ -dependent on at least  $L = \lfloor (\kappa - 1)/d_E(\Phi, \varepsilon) \rfloor$  disjoint subsequences in  $\{z_1, \dots, z_{j-1}\}$ .*

**Proof** [Proof of Claim 2] Let  $L$  be defined as in the statement of the claim. If  $L = 0$ , there is nothing to be proven. Otherwise we argue as follows.

Initialize the sequences  $B_1 = \{z_1\}, \dots, B_L = \{z_L\}$ , and let  $j = L + 1$ . Now, consider the following process: if at any point,  $z_j$  is  $\varepsilon$ -dependent on all of the  $B_1, \dots, B_L$ , then the claim is proven and we terminate the process. Otherwise, we pick an  $i \in [L]$  such that  $z_j$  is  $\varepsilon$ -independent of  $B_i$  and update  $B_i := B_i \cup \{z_j\}$ . Then we increment  $j$  by 1 and continue this process if  $j \leq \kappa$ .

It remains to be proven that the process is terminated before  $j$  gets to  $\kappa + 1$ . We prove this by contradiction: Assume that  $j$  gets to  $\kappa + 1$ . Let  $j$  be the index of a sequence from  $\{B_i\}_{i \in [L]}$  that has the most elements (in case of ties, chose arbitrarily). From  $\kappa = \sum_{i=1}^L |B_i| \leq L|B_j|$ ,  $|B_j| \geq \kappa/L \geq \frac{\kappa}{\kappa-1} d_E(\Phi, \varepsilon) > d_E(\Phi, \varepsilon)$ . But  $B_j$  is an  $\varepsilon$ -eluder sequence by construction, which contradicts that  $d_E(\Phi, \varepsilon)$  is the length of longest  $\varepsilon'$ -eluder sequences with  $\varepsilon' \geq \varepsilon$ . ■

Equipped with Claims 1 and 2, we can prove the following lemma, which bounds the frequency of large values in  $\{|\phi_1(x_1)|, \dots, |\phi_k(x_k)|\}$ .

**Lemma 23** *Under the same condition of Proposition 21, for all  $k \in [K]$ ,*

$$\sum_{t=1}^k \mathbf{1}\{|\phi_t(x_t)| > \varepsilon\} \leq \left(\frac{\beta}{\varepsilon} + 1\right) d_E(\Phi, \varepsilon) + 1.$$

**Proof** Fix  $k \in [K]$  and let  $\{z_1, \dots, z_\kappa\}$  be the subsequence of  $\{x_1, \dots, x_k\}$  consisting of elements for which  $|\phi_t(x_t)| > \varepsilon$ . By Claim 2, we know there exists  $j \in [\kappa]$  such that  $z_j$  is  $\varepsilon$ -dependent on  $\lfloor (\kappa - 1)/d_E(\Phi, \varepsilon) \rfloor$  disjoint subsequences of  $\{z_1, \dots, z_{j-1}\}$ . By Claim 1,  $z_j$  is  $\varepsilon$ -dependent on at most  $\beta/\varepsilon$  disjoint subsequences of  $\{z_1, \dots, z_{j-1}\}$ . Therefore, we have

$$\lfloor (\kappa - 1)/d_E(\Phi, \varepsilon) \rfloor \leq \beta/\varepsilon,$$

which implies

$$\kappa \leq \left(\frac{\beta}{\varepsilon} + 1\right) d_E(\Phi, \varepsilon) + 1,$$

completing the proof. ■

With this, we are ready to prove Proposition 21.

**Proof** [Proof of Proposition 21] Fix  $k \in [K]$ ; let  $d = d_E(\Phi, \omega)$ . Noting that  $|\phi_t(x_t)| \leq C$ , we have

$$\begin{aligned} \sum_{t=1}^k |\phi_t(x_t)| &= \sum_{t=1}^k \int_0^C \mathbf{1}\{|\phi_t(x_t)| > y\} dy \\ &\leq k\omega + \sum_{t=1}^k \int_\omega^C \mathbf{1}\{|\phi_t(x_t)| > y\} dy \\ &= k\omega + \int_\omega^C \left( \sum_{t=1}^k \mathbf{1}\{|\phi_t(x_t)| > y\} \right) dy \\ &\leq k\omega + \int_\omega^C \left( \left\{ \frac{\beta}{y} + 1 \right\} d + 1 \right) dy \\ &\leq (d + 1)C + d\beta \log(C/\omega) + k\omega, \end{aligned}$$

where the one but last inequality follows from Lemma 23 and the monotonicity of  $\ell_1$ -norm eluder dimension.  $\blacksquare$

### E.3. Proof of Proposition 22

**Proof** Given  $v \in [Kn]$ , denote  $p_v = \lfloor \frac{v-1}{n} \rfloor + 1$  and  $q_v = v - n \lfloor \frac{v-1}{n} \rfloor$ . Note that  $1 \leq p_v \leq K$  and  $1 \leq q_v \leq n$ . We introduce the following two auxiliary sequences:

$$\begin{cases} \zeta_{v,j} = w_{p_v,j} \\ z_v = x_{p_v,q_v} \end{cases} \quad (v,j) \in [Kn] \times [m]. \quad (19)$$

By the definition of  $(\zeta_{v,j}, z_v)$  and precondition (15),

$$\sum_{t=1}^{v-1} \sum_{j=1}^m |\zeta_{v,j}^\top z_t| = \sum_{k=1}^{p_v-1} \sum_{i=1}^n \sum_{j=1}^m |w_{p_v,j}^\top x_{k,i}| + \sum_{i=1}^{q_v-1} \sum_{j=1}^m |w_{p_v,j}^\top x_{p_v,i}| \leq \gamma_{p_v} + R_w R_x. \quad (20)$$

Notice that for all  $k \in [K]$ ,

$$\sum_{v=1}^{kn} \sum_{j=1}^m |\zeta_{v,j}^\top z_v| = \sum_{t=1}^k \sum_{i=1}^n \sum_{j=1}^m |w_{t,j}^\top x_{t,i}|.$$

So in order to prove Proposition 22, it suffices to upper bound the value of the following optimization problem:

$$\begin{aligned} & \max_{\zeta, z} \sum_{v=1}^{kn} \sum_{j=1}^m |\langle \zeta_{v,j}, z_v \rangle| \\ \text{s.t. for any } v : & \begin{cases} \sum_{t=1}^{v-1} \sum_{j=1}^m |\zeta_{v,j}^\top z_t| \leq \gamma_{p_v} + R_w R_x, \\ \sum_{j=1}^m \|\zeta_{v,j}\|_2 \leq R_w, \\ \|z_v\|_2 \leq R_x. \end{cases} \end{aligned} \quad (21)$$

Now, we make the observation that the above optimization problem has the same optimal value as the following one (we assume all vectors in this section have the same dimension):

$$\begin{aligned} & \max_{\zeta, z} \sum_{v=1}^{kn} |\langle \zeta_v, z_v \rangle| \\ \text{s.t. for any } v : & \begin{cases} \sum_{t=1}^{v-1} |\zeta_v^\top z_t| \leq \gamma_{p_v} + R_w R_x, \\ \|\zeta_v\|_2 \leq R_w, \\ \|z_v\|_2 \leq R_x. \end{cases} \end{aligned} \quad (22)$$

It is direct to see the optimal value of Problem (22) is always no larger than that of Problem (21). For the other direction, suppose  $\{\zeta_{v,j}^*\}_{(v,j) \in [Kn] \times [m]}$  and  $\{z_v^*\}_{v \in [Kn]}$  are optimal solution to Problem (21). Consider

$$\zeta_v^* = \sum_{j=1}^m \zeta_{v,j}^* \times \text{sign}(\zeta_{v,j}^{*\top} z_v^*) \quad \text{for } k \in [K].$$

One can easily verify  $\{\phi_k^*\}_{k \in [K]}$  and  $\{z_v^*\}_{v \in [Kn]}$  are feasible solution to Problem (21) by triangle inequality. Moreover, its objective value is

$$\sum_{v=1}^{kn} \left| \sum_{j=1}^m \left( \zeta_{v,j}^* \times \text{sign}(\zeta_{v,j}^{*\top} z_v^*) \right)^\top z_v^* \right| = \sum_{v=1}^{kn} \sum_{j=1}^m \left| \zeta_{v,j}^{*\top} z_v^* \right|,$$

which is equal to the optimal value of Problem (21). Therefore, it suffices to upper bound the optimal value of Problem (22).

Finally, by applying Proposition 21 to  $d$ -dimensional linear function class, we can upper bound the optimal value of Problem (22) by

$$\min_{\omega \leq C} \{d_E(\omega)C + d_E(\omega)\lambda \log(C/\omega) + kn\omega\}$$

where

$$\begin{cases} d_E(\omega) = d \log(e + R_w R_x / \omega) \\ C = R_w R_x \\ \lambda = \gamma_{p_v} + R_w R_x. \end{cases}$$

We conclude the proof by choosing  $\omega = R_w R_x / (kn)$ . ■

## Appendix F. Proofs for Undercomplete POMDPs

In this section, we prove Theorem 4 with a specific polynomial dependency as stated in the following theorem.

**Theorem 24 (Regret of OMLE)** *There exists an absolute constant  $c > 0$  such that for any  $\delta \in (0, 1]$  and  $S, A, O, H, K \in \mathbb{N}$ , if we choose  $\beta = c(H(S^2 A + SO) \log(SAOHK) + \log(K/\delta))$  in Algorithm 1, then, for any POMDP with  $S$  states,  $A$  actions,  $O$  observations and horizon  $H$  and satisfying Assumption 1, with probability at least  $1 - \delta$ ,*

$$\text{Regret}(k) \leq \tilde{O} \left( \frac{S^2 AO}{\alpha^2} \sqrt{k(S^2 A + SO)} \times \text{poly}(H) \right) \quad \text{for all } k \in [K].$$

To begin with, we recall (or introduce) some notations that will be needed in the proof.

**Notation** Denote by  $\tau_h = (o_1, a_1, \dots, o_h, a_h)$  a trajectory up to step  $h$ . (That is, for a natural number  $h$ ,  $\tau_h \in (\mathcal{O} \times \mathcal{A})^h$ , by convention.) Given a policy  $\pi$  and a trajectory  $\tau_h$ , define  $\pi(\tau_h) = \prod_{h'=1}^h \pi(a_{h'} | \tau_{h'-1}, o_{h'})$ . This is the part of the probability of  $\tau_h$  that can be attributed to the policy  $\pi$ . We use  $\theta^* = (\mathbb{T}, \mathbb{O}, \mu_1)$  to denote the parameters of the groundtruth POMDP that the learner interacts with, and denote by  $\mathbf{B}$  the observable operators corresponding to  $\theta^*$  and by  $\mathbf{b}_0$  the initial observation distribution. Formally,

$$\begin{cases} \mathbf{b}_0 = \mathbb{O}_1 \mu_1 \in \mathbb{R}^O, \\ \mathbf{B}_h(o, a) = \mathbb{O}_{h+1} \mathbb{T}_{h,a} \text{diag}(\mathbb{O}_h(o | \cdot)) \mathbb{O}_h^\dagger \in \mathbb{R}^{O \times O}, \end{cases} \quad (23)$$

where  $\mathbb{O}_h(o | \cdot)$  is the  $o^{\text{th}}$  row of the observation matrix  $\mathbb{O}_h$  and  $\text{diag}(\mathbb{O}_h(o | \cdot))$  is a diagonal matrix with diagonals equal to  $\mathbb{O}_h(o | \cdot)$ . We use  $\theta^k$  to denote the optimistic estimate of the

POMDP model in the  $k^{\text{th}}$  episode of Algorithm 1, and denote by  $\pi^k$  the optimal policy of model  $\theta^k$ , i.e.,  $(\theta^k, \pi^k) = \operatorname{argmax}_{\theta \in \mathcal{B}^{k-1}, \pi} V_1^\pi(\theta)$ . Finally, we will use  $\mathbf{B}^k$  and  $\mathbf{b}_0^k$  to denote the operators corresponding to  $\theta^k$ .

For a matrix  $\mathbf{A}$ ,  $p \geq 1$ , let  $\|\mathbf{A}\|_p = \sup_{x: \|x\|_p \leq 1} \|\mathbf{A}x\|_p$ . We start with a lemma that bounds  $\|\mathbf{B}_h(o, a)\|_1$  as a function of  $\alpha$ :

**Lemma 25** *Under Assumption 1 it holds that for any  $(o, a, h, k) \in \mathcal{O} \times \mathcal{A} \times [H-1] \times \mathbb{N}$ ,*

$$\|\mathbf{B}_h(o, a)\|_1, \|\mathbf{B}_h^k(o, a)\|_1 \leq \sqrt{S}/\alpha.$$

**Proof** By definition,

$$\begin{aligned} \|\mathbf{B}_h(o, a)\|_1 &= \|\mathbb{O}_{h+1} \mathbb{T}_{h,a} \operatorname{diag}(\mathbb{O}_h(o | \cdot)) \mathbb{O}_h^\dagger\|_1 \\ &\leq \|\mathbb{O}_{h+1}\|_1 \times \|\mathbb{T}_{h,a}\|_1 \times \|\operatorname{diag}(\mathbb{O}_h(o | \cdot))\|_1 \times \|\mathbb{O}_h^\dagger\|_1 \\ &\leq \|\mathbb{O}_h^\dagger\|_1 = \sup_{x: \|x\|_1 \leq 1} \|\mathbb{O}_h^\dagger x\|_1 \leq \sqrt{S} \sup_{x: \|x\|_2 \leq 1} \|\mathbb{O}_h^\dagger x\|_2 \leq \sqrt{S}/\alpha, \end{aligned}$$

where the last inequality uses Assumption 1. The same arguments also give  $\|\mathbf{B}_h^k(o, a)\|_1 \leq \sqrt{S}/\alpha$ .  $\blacksquare$

**Computing probabilities using the operators** As it is well known (e.g., Jin et al., 2020a), given a set of POMDP parameters  $\theta$ , its corresponding operators satisfy the following relations for any policy  $\pi$  and trajectory  $\tau_h = (o_1, a_1, \dots, o_h, a_h)$ :

$$\begin{cases} \mathbb{P}_\theta^\pi(\tau_h) = \pi(\tau_h) \cdot (\mathbf{e}_{o_h}^\top \mathbf{B}_{h-1}(o_{h-1}, a_{h-1}; \theta) \cdots \mathbf{B}_1(o_1, a_1; \theta) \mathbf{b}_0(\theta)) \in \mathbb{R}, \\ \mathbb{P}_\theta^\pi(\tau_h, o_{h+1}) = \pi(\tau_h) \cdot (\mathbf{e}_{o_{h+1}}^\top \mathbf{B}_h(o_h, a_h; \theta) \cdots \mathbf{B}_1(o_1, a_1; \theta) \mathbf{b}_0(\theta)) \in \mathbb{R}, \\ \mathbb{P}_\theta^\pi(o_1) = \mathbf{e}_{o_1}^\top \mathbf{b}_0(\theta) \in \mathbb{R}^{\mathcal{O}}, \end{cases} \quad (24)$$

where  $\mathbb{P}_\theta^\pi(\tau_h)$  denotes the probability of observing  $\tau_h$  under policy  $\pi$  in the POMDP model defined by  $\theta$ , and  $(\{\mathbf{B}_h(o, a; \theta)\}_{(o,a,h) \in \mathcal{O} \times \mathcal{A} \times [H-1]}, \mathbf{b}_0(\theta))$  denote the operators corresponding to  $\theta$  (Equation (5) in the main body is the first case in Equation (24).)

**Proof roadmap** The proof consists of three main steps. In the first step, we upper bound the distribution estimation error by the operator estimation error. The proof of this follows arguments similar to those used by Jin et al. (2020a). Then we leverage the MLE guarantee developed in Appendix B to derive certain constraints satisfied by the operator estimation error. Finally, we use the  $\ell_1$ -norm eluder dimension framework developed in Appendix D to bridge the results in the first two steps and obtain the desired regret guarantee.

### F.1. Step 1: bound the regret by the error of operator estimates

We first present a lemma that upper bounds the cumulative suboptimality of  $\pi_1, \dots, \pi_k$  by the cumulative density estimation error.

**Lemma 26** *In Algorithm 1 and 2, if we choose  $\beta$  according to Theorem 4 and 7 respectively, then with probability at least  $1 - \delta$ ,*

$$\sum_{t=1}^k V^{\pi^t}(\theta^t) - V^{\pi^t}(\theta^*) \leq H \sum_{t=1}^k \sum_{\tau_H} |\mathbb{P}_{\theta^t}^{\pi^t}(\tau_H) - \mathbb{P}_{\theta^*}^{\pi^t}(\tau_H)|. \quad (25)$$

**Proof** By the choice of  $\beta$  and Proposition 13, we have  $\theta^* \in \mathcal{B}^t$  for all  $t \in [K]$  with probability at least  $1 - \delta$ . In what follows, we assume that the event  $\theta^* \in \cap_{t \in [K]} \mathcal{B}^t$  holds. On this event, by the optimism of  $\theta^t$  and  $\pi^t$  for  $t \in [k]$ ,

$$\sum_{t=1}^k V^*(\theta^*) - V^{\pi^t}(\theta^*) \leq \sum_{t=1}^k V^{\pi^t}(\theta^t) - V^{\pi^t}(\theta^*). \quad (26)$$

Because the cumulative reward of each trajectory is bounded by  $H$ , we conclude

$$\sum_{t=1}^k V^{\pi^t}(\theta^t) - V^{\pi^t}(\theta^*) \leq H \sum_{t=1}^k \sum_{\tau_H} |\mathbb{P}_{\theta^t}^{\pi^t}(\tau_H) - \mathbb{P}_{\theta^*}^{\pi^t}(\tau_H)|. \quad (27)$$

■

As a result, to prove Theorem 4, it suffices to upper bound the RHS of Equation (25).

To begin with, we represent the probability of observing  $\tau_H$  by the product of operators using Equation (24), which gives

$$\begin{aligned} & \sum_{t=1}^k \sum_{\tau_H} \left| \mathbb{P}_{\theta^t}^{\pi^t}(\tau_H) - \mathbb{P}_{\theta^*}^{\pi^t}(\tau_H) \right| \\ &= \sum_{t=1}^k \sum_{\tau_H} \left| \mathbf{e}_{o_H}^\top \left( \prod_{h=1}^{H-1} \mathbf{B}_h^t(o_h, a_h) \right) \mathbf{b}_0^t - \mathbf{e}_{o_H}^\top \left( \prod_{h=1}^{H-1} \mathbf{B}_h(o_h, a_h) \right) \mathbf{b}_0 \right| \times \pi^t(\tau_H) \\ &= \sum_{t=1}^k \sum_{\tau_{H-1}} \left\| \left( \prod_{h=1}^{H-1} \mathbf{B}_h^t(o_h, a_h) \right) \mathbf{b}_0^t - \left( \prod_{h=1}^{H-1} \mathbf{B}_h(o_h, a_h) \right) \mathbf{b}_0 \right\|_1 \times \pi^t(\tau_{H-1}). \end{aligned}$$

By using the lemma below, we can further control the difference between two products of operators by the difference between each pair of operators. The proof of a more general version of this lemma is given in Section G.4.

**Lemma 27** For any  $k \in \mathbb{N}$ ,  $h \in [H - 1]$  and policy  $\pi$

$$\begin{aligned} & \sum_{\tau_h} \left\| \mathbf{B}_h^k(o_h, a_h) \cdots \mathbf{B}_1^k(o_1, a_1) \mathbf{b}_0^k - \mathbf{B}_h(o_h, a_h) \cdots \mathbf{B}_1(o_1, a_1) \mathbf{b}_0 \right\|_1 \times \pi(\tau_h) \\ & \leq \frac{\sqrt{S}}{\alpha} \left( \sum_{j=1}^h \sum_{\tau_j} \left\| (\mathbf{B}_j^k(o_j, a_j) - \mathbf{B}_j(o_j, a_j)) \mathbf{b}(\tau_{j-1}) \right\|_1 \times \pi(\tau_j) + \|\mathbf{b}_0^k - \mathbf{b}_0\|_1 \right), \end{aligned}$$

where for a trajectory  $\tau_h = (o_1, a_1, \dots, o_h, a_h)$ ,  $\mathbf{b}(\tau_h) = \left( \prod_{h'=1}^h \mathbf{B}_{h'}(o_{h'}, a_{h'}) \right) \mathbf{b}_0$ .

As noted in the main text, this lemma abuses notations in a few ways: In the innermost sum over the observation-action trajectories  $\tau_j$  of length  $j$ ,  $\tau_{j-1}$  refers to the prefix of  $\tau_j$  where the last observation-action is dropped. Also, in this sum,  $(o_j, a_j)$  refer to the last observation-action pair of  $\tau_j$ .



Putting things together, we get

$$\begin{aligned}
 & \sum_{t=1}^k \sum_{\tau_H} |\mathbb{P}_{\theta^t}^{\pi^t}(\tau_H) - \mathbb{P}_{\theta^*}^{\pi^t}(\tau_H)| \\
 & \leq \sum_{t=1}^k \sum_{\tau_{H-1}} \left\| \left( \prod_{h=1}^{H-1} \mathbf{B}_h^t(o_h, a_h) \right) \mathbf{b}_0^t - \left( \prod_{h=1}^{H-1} \mathbf{B}_h(o_h, a_h) \right) \mathbf{b}_0 \right\|_1 \times \pi^t(\tau_{H-1}) \\
 & \leq \frac{\sqrt{S}}{\alpha} \left( \sum_{t=1}^k \sum_{h=1}^{H-1} \sum_{\tau_h} \|(\mathbf{B}_h(o_h, a_h) - \mathbf{B}_h^t(o_h, a_h)) \mathbf{b}(\tau_{h-1})\|_1 \times \pi^t(\tau_h) + \|\mathbf{b}_0 - \mathbf{b}_0^t\|_1 \right). \tag{28}
 \end{aligned}$$

## F.2. Step 2: derive constraints for the operator estimates from OMLE

By the construction of  $\mathcal{B}^k$ ,  $\sum_{i=1}^{k-1} \log \left( \frac{\mathbb{P}_{\theta^*}^{\pi^i}(\tau^i)}{\mathbb{P}_{\theta^k}^{\pi^i}(\tau^i)} \right) \leq \beta$ . Therefore, by Proposition 14 and the choice of  $\beta$ , we have with probability at least  $1 - \delta$ ,

$$\text{for all } k \in [K]: \quad \sum_{t=1}^{k-1} \left( \sum_{\tau_H} |\mathbb{P}_{\theta^k}^{\pi^t}(\tau_H) - \mathbb{P}_{\theta^*}^{\pi^t}(\tau_H)| \right)^2 = \mathcal{O}(\beta).$$

In the rest of the proof, we assume that the event above is true.

By the Cauchy-Schwarz inequality,

$$\sum_{t=1}^{k-1} \sum_{\tau_H} |\mathbb{P}_{\theta^k}^{\pi^t}(\tau_H) - \mathbb{P}_{\theta^*}^{\pi^t}(\tau_H)| = \mathcal{O}(\sqrt{\beta k}), \quad \text{for all } k \in [K].$$

Since marginalizing two distributions will not increase their distance, from the previous inequality we also have that for all  $(k, h) \in [K] \times [H - 1]$ ,

$$\begin{cases} \sum_{t=1}^{k-1} \sum_{\tau_h, o_{h+1}} |\mathbb{P}_{\theta^k}^{\pi^t}(\tau_h, o_{h+1}) - \mathbb{P}_{\theta^*}^{\pi^t}(\tau_h, o_{h+1})| = \mathcal{O}(\sqrt{\beta k}), \\ \sum_{t=1}^{k-1} \sum_{o_1} |\mathbb{P}_{\theta^k}^{\pi^t}(o_1) - \mathbb{P}_{\theta^*}^{\pi^t}(o_1)| = \mathcal{O}(\sqrt{\beta k}). \end{cases}$$

By equation (24), we can replace the probability with the product of operators and obtain for all  $(k, h) \in [K] \times [H - 1]$

$$\begin{cases} \sum_{t=1}^{k-1} \sum_{\tau_h} \pi^t(\tau_h) \times \|\mathbf{b}^k(\tau_h) - \mathbf{b}(\tau_h)\|_1 = \mathcal{O}(\sqrt{\beta k}), \\ \|\mathbf{b}_0^k - \mathbf{b}_0\|_1 = \mathcal{O}(\sqrt{\beta/k}), \end{cases} \tag{29}$$

where we used that the initial distributions over the observations are independent of the policies used and where  $\mathbf{b}(\tau_h) = \left( \prod_{h'=1}^h \mathbf{B}_{h'}(o_{h'}, a_{h'}) \right) \mathbf{b}_0$  and  $\mathbf{b}^k(\tau_h) = \left( \prod_{h'=1}^h \mathbf{B}_{h'}^k(o_{h'}, a_{h'}) \right) \mathbf{b}_0^k$ .

The second part of Equation (29) shows that the parameter estimates are improving in terms of how well the initial observation distribution is estimated (which is unsurprising). Our goal now is to use the first part of this equation to show that the estimated observable operators are also getting

closer to their groundtruth counterparts. To show this, first notice that by the triangle inequality, for any  $(k, h) \in [K] \times [H - 1]$ ,

$$\begin{aligned} & \sum_{t=1}^{k-1} \sum_{\tau_h} \pi^t(\tau_h) \times \left\| \left( \mathbf{B}_h^k(o_h, a_h) - \mathbf{B}_h(o_h, a_h) \right) \mathbf{b}(\tau_{h-1}) \right\|_1 \\ & \leq \sum_{t=1}^{k-1} \sum_{\tau_h} \pi^t(\tau_h) \times \left\| \mathbf{B}_h^k(o_h, a_h) \left( \mathbf{b}^k(\tau_{h-1}) - \mathbf{b}(\tau_{h-1}) \right) \right\|_1 \\ & \quad + \sum_{t=1}^{k-1} \sum_{\tau_h} \pi^t(\tau_h) \times \left\| \mathbf{B}_h^k(o_h, a_h) \mathbf{b}^k(\tau_{h-1}) - \mathbf{B}_h(o_h, a_h) \mathbf{b}(\tau_{h-1}) \right\|_1. \end{aligned}$$

In the last display, the second term is at most  $\mathcal{O}(\sqrt{k\beta})$  by (29). For the first term, we can apply Lemma 32 and obtain

$$\sum_{t=1}^{k-1} \sum_{\tau_h} \pi^t(\tau_h) \times \left\| \mathbf{B}_h^k(o_h, a_h) \left( \mathbf{b}^k(\tau_{h-1}) - \mathbf{b}(\tau_{h-1}) \right) \right\|_1 = \mathcal{O} \left( \frac{\sqrt{S}}{\alpha} \sqrt{k\beta} \right).$$

Putting things together, we conclude that with probability at least  $1 - \delta$ :

$$\sum_{t=1}^{k-1} \sum_{\tau_h} \pi^t(\tau_h) \times \left\| \left( \mathbf{B}_h^k(o_h, a_h) - \mathbf{B}_h(o_h, a_h) \right) \mathbf{b}(\tau_{h-1}) \right\|_1 = \mathcal{O} \left( \frac{\sqrt{S}}{\alpha} \sqrt{k\beta} \right) \quad (30)$$

for all  $(k, h) \in [K] \times [H - 1]$ .

### F.3. Step 3: bridge Step 1 and 2 via $\ell_1$ -norm eluder dimension

By Step 1 (in particular, by Equation (28)), to control the total estimation error of the distributions over trajectories, it suffices to upper bound the following quantity for all  $k \in [K]$ :

$$\mathbf{Target:} \sum_{t=1}^k \left( \sum_{h=1}^{H-1} \sum_{\tau_h} \left\| \left( \mathbf{B}_h^t(o_h, a_h) - \mathbf{B}_h(o_h, a_h) \right) \mathbf{b}(\tau_{h-1}) \right\|_1 \times \pi^t(\tau_h) + \left\| \mathbf{b}_0^t - \mathbf{b}_0 \right\|_1 \right).$$

By Step 2, we have that with probability at least  $1 - \delta$ : for all  $(k, h) \in [K] \times [H - 1]$

$$\mathbf{Condition:} \begin{cases} \left\| \mathbf{b}_0^k - \mathbf{b}_0 \right\|_1 = \mathcal{O}(\sqrt{\beta/k}), \\ \sum_{t=1}^{k-1} \sum_{\tau_h} \left\| \left( \mathbf{B}_h^k(o_h, a_h) - \mathbf{B}_h(o_h, a_h) \right) \mathbf{b}(\tau_{h-1}) \right\|_1 \times \pi^t(\tau_h) = \mathcal{O} \left( \frac{\sqrt{S}}{\alpha} \sqrt{k\beta} \right). \end{cases}$$

From now on, we assume that the two conditions of the last display hold and will show that this suffices to bound the ‘‘target’’ above. By using the first inequality in the condition, the summation of the second term in the target is upper bounded by  $\mathcal{O}(\sum_{t=1}^k \sqrt{\beta/t}) = \mathcal{O}(\sqrt{k\beta})$ . It remains to control the summation of the first term in the target.

Now let us consider an *arbitrary fixed* triple  $(o, a, h) \in \mathcal{O} \times \mathcal{A} \times [H - 1]$ . Denote by  $\mathbf{X}_l$  the  $l^{\text{th}}$  row of matrix  $\mathbf{X}$ . The condition implies that

$$\sum_{t=1}^{k-1} \sum_{\tau_h: (o_h, a_h) = (o, a)} \sum_{l \in [O]} \left| \left[ \left( \mathbf{B}_h^k(o, a) - \mathbf{B}_h(o, a) \right) \mathbb{O}_h \right]_l \mathbb{O}_h^\dagger \mathbf{b}(\tau_{h-1}) \right| \times \pi^t(\tau_h) = \mathcal{O} \left( \frac{\sqrt{S}}{\alpha} \sqrt{k\beta} \right). \quad (31)$$

To further simplify the notations, for  $(t, l) \in [K] \times [O]$ , let  $w_{t,l} := [(\mathbf{B}_h^t(o, a) - \mathbf{B}_h(o, a)) \mathbb{O}_h]_l$  and denote the  $n := (OA)^{h-1}$  not necessarily distinct elements in the sequence  $\{\mathbb{O}_h^\dagger \mathbf{b}(\tau_{h-1}) \times \pi^t(\tau_h) : \tau_h = (o_1, a_1, \dots, o_{h-1}, a_{h-1}, o, a) \in (\mathcal{O} \times \mathcal{A})^{h-1} \times (o, a)\}$  by  $x_{t,1}, \dots, x_{t,n}$ . Using the newly defined notations, Equation (31) is equivalent to

$$\sum_{t=1}^{k-1} \sum_{l=1}^O \sum_{i=1}^n |w_{k,l}^\top x_{t,i}| = \mathcal{O} \left( \frac{\sqrt{S}}{\alpha} \sqrt{k\beta} \right). \quad (32)$$

We have three observations about the  $x, w$  sequences:

- The vectors  $\{x_{t,i}\}_{i=1}^n$  satisfy  $\sum_{i=1}^n \|x_{t,i}\|_1 \leq 1$  for all  $t$  because

$$\begin{aligned} \sum_{\tau_h: (o_h, a_h) = (o, a)} \|\mathbb{O}_h^\dagger \mathbf{b}(\tau_{h-1}) \pi^t(\tau_h)\|_1 &\leq \sum_{\tau_h: (o_h, a_h) = (o, a)} \|\mathbb{O}_h^\dagger \mathbf{b}(\tau_{h-1}) \pi^t(\tau_{h-1})\|_1 \\ &= \sum_{\tau_{h-1}} \|\mathbb{O}_h^\dagger \mathbf{b}(\tau_{h-1}) \pi^t(\tau_{h-1})\|_1 = 1, \end{aligned}$$

where the final equality follows from  $\|\mathbb{O}_h^\dagger \mathbf{b}(\tau_{h-1}) \pi(\tau_{h-1})\|_1 = \mathbb{P}_{\theta^*}^\pi(\tau_{h-1})$ .

- The vectors  $\{w_{t,l}\}_{l=1}^O$  satisfy  $\sum_{l=1}^O \|w_{t,l}\|_1 \leq 1$  for all  $t$ . By definition and Lemma 25,

$$\sum_{l=1}^O \|w_{t,l}\|_1 = \|(\mathbf{B}_h(o, a) - \mathbf{B}_h^t(o, a)) \mathbb{O}_h\|_1 \leq S(\|\mathbf{B}_h(o, a)\|_1 + \|\mathbf{B}_h^t(o, a)\|_1) \leq \frac{2S^{1.5}}{\alpha}.$$

As a result, we can invoke Proposition 22 and obtain

$$\sum_{t=1}^k \sum_{l=1}^O \sum_{i=1}^n |w_{t,l}^\top x_{t,i}| = \tilde{\mathcal{O}} \left( \frac{S^{1.5} H^2}{\alpha} \sqrt{k\beta} \right), \quad \text{for all } k \in [K]. \quad (33)$$

Translating the above inequality back to the operator language gives that for an *arbitrary fixed* tuple  $(o, a, h) \in \mathcal{O} \times \mathcal{A} \times [H-1]$

$$\sum_{t=1}^k \sum_{\tau_h: (o_h, a_h) = (o, a)} |(\mathbf{B}_h^t(o, a) - \mathbf{B}_h(o, a))_l \mathbf{b}(\tau_{h-1})| \times \pi^t(\tau_h) = \tilde{\mathcal{O}} \left( \frac{S^{1.5} H^2}{\alpha} \sqrt{k\beta} \right), \quad \text{for all } k \in [K].$$

Therefore, we conclude that with probability at least  $1 - \delta$ , for all  $k \in [K]$ ,

$$\sum_{t=1}^k \left( \sum_{h=1}^{H-1} \sum_{\tau_h} \|(\mathbf{B}_h^t(o_h, a_h) - \mathbf{B}_h(o_h, a_h)) \mathbf{b}(\tau_{h-1})\|_1 \times \pi^t(\tau_h) + \|\mathbf{b}_0^t - \mathbf{b}_0\|_1 \right) = \tilde{\mathcal{O}} \left( \frac{S^{1.5} OAH^3}{\alpha} \sqrt{k\beta} \right).$$

Combining the above equality with Step 1 gives:

$$\text{Regret}(k) \leq \tilde{\mathcal{O}} \left( \frac{S^2 AO}{\alpha^2} \sqrt{k(S^2 A + SO)} \times \text{poly}(H) \right) \quad \text{for all } k \in [K].$$

**Remark 28** We remark that the regret above can be improved by a factor of  $O$  if we adopt a more active strategy for exploration in *OMLE*. The modified algorithm is a special case of Algorithm 2 with  $m = 1$ , which we describe in details in Section 4.2.

## Appendix G. Proofs for Overcomplete POMDPs

In this section, we prove Theorem 7 with a specific polynomial dependency as stated in the following theorem.

**Theorem 29 (Total suboptimality of multi-step OMLE)** *There exists an absolute constant  $c > 0$  such that for any  $\delta \in (0, 1]$  and  $S, A, O, K, H \in \mathbb{N}$ , if we choose parameter  $\beta$  in Algorithm 2 as  $\beta = c(H(S^2A + SO) \log(SAOH) + \log(KH/\delta))$ , then, for any POMDP with  $S$  states,  $A$  actions,  $O$  observations and horizon  $H$  and satisfying Assumption 2, with probability at least  $1 - \delta$ ,*

$$\sum_{t=1}^k (V^* - V^{\pi^t}) \leq \tilde{\mathcal{O}} \left( \frac{S^2 A^{3m-2}}{\alpha^2} \sqrt{k(S^2 A + SO)} \times \text{poly}(H) \right) \quad \text{for all } k \in [K].$$

The proof follows a similar three-step recipe to the undercomplete case: first we show that to control the regret it suffices to control the operator estimation error, then we derive some constraints for the operator estimation from the MLE guarantee, and finally we utilize the constraints to upper bound the estimation error with the help of the  $\ell_1$ -norm eluder dimension framework. The main difference to the previous proof is only that the data generated for building the confidence sets is based on adding exploration to the policies  $\pi^1, \dots, \pi^K$ .

The reader may find it useful to first read Appendix F before reading this section as we will heavily reuse the notations and techniques developed there.

**Notation** As usual, we start by introducing (or recalling) the definitions of several notations that will be used in the proofs. Recall we define the  $m$ -step emission-action matrices  $\{\mathbb{M}_h \in \mathbb{R}^{(A^{m-1}O^m) \times S}\}_{h \in [H-m+1]}$  as follows: For an observation sequence  $\mathbf{o}$  of length  $m$ , initial state  $s$  and action sequence  $\mathbf{a}$  of length  $m-1$ , we let  $[\mathbb{M}_h]_{(\mathbf{a}, \mathbf{o}), s}$  be the probability of receiving  $\mathbf{o}$  provided that the action sequence  $\mathbf{a}$  is used from state  $s$  and step  $h$ : for all  $(\mathbf{a}, \mathbf{o}) \in \mathcal{A}^{m-1} \times \mathcal{O}^m$  and  $s \in \mathcal{S}$

$$[\mathbb{M}_h]_{(\mathbf{a}, \mathbf{o}), s} = \mathbb{P}(o_{h:h+m-1} = \mathbf{o} \mid s_h = s, a_{h:h+m-2} = \mathbf{a}). \quad (34)$$

With slight abuse of notations, we reuse  $\mathbf{B}$  and  $\mathbf{b}$  to denote the operators for the general setting. Specifically, the operators corresponding to the groundtruth model  $\theta^*$  is defined as:

$$\begin{cases} \mathbf{b}_0 = \mathbb{M}_1 \mu_1 \in \mathbb{R}^{A^{m-1}O^m}, \\ \mathbf{B}_h(o, a) = \mathbb{M}_{h+1} \mathbb{T}_{h,a} \text{diag}(\mathbb{O}_h(o \mid \cdot)) \mathbb{M}_h^\dagger \in \mathbb{R}^{(A^{m-1}O^m) \times (A^{m-1}O^m)}, \end{cases} \quad (35)$$

where the only difference from the undercomplete version (equation (23)) is that we replace the single-step emission matrices with the  $m$ -step emission-action matrices.

We use  $\theta^k$  to denote the optimistic estimate of the POMDP model in the  $k^{\text{th}}$  iteration of Algorithm 2, and denote by  $\pi^k$  the optimal policy for model  $\theta^k$ , i.e.,  $(\theta^k, \pi^k) = \arg\max_{\theta \in \mathcal{B}^{k-1}, \pi} V_1^\pi(s_1; \theta)$ . We use  $\mathbf{B}^k$  and  $\mathbf{b}_0^k$  to denote the operators corresponding to  $\theta^k$ .

To simplicity the notation, throughout this section, we will use  $\mathbf{u} \in (\mathcal{O} \times \mathcal{A})^{m-1} \times \mathcal{O}$  to denote a length- $(2m-1)$  observation-action sequence (the observations are interleaved with actions in  $\mathbf{u}$ ). We use  $\mathbf{u}_a \in \mathcal{A}^{m-1}$  to refer to the action part of  $\mathbf{u}$  and  $\mathbf{u}_o \in \mathcal{O}^m$  for the observation part. We further define  $\mathbf{U} := (\mathcal{O} \times \mathcal{A})^{m-1} \times \mathcal{O}$  to be the collections of all the  $(2m-1)$ -length action-observation sequences.

**Computing probabilities using the operators** Similar to the undercomplete case, several useful relations hold for the newly defined operators. In particular, for any policy  $\pi$ , trajectory  $\tau_h$  (with  $h \leq H - m$ ) and observation-action sequence  $\mathbf{u} \in \mathbf{U}$ , it holds that

$$\begin{cases} \mathbb{P}_\theta^\pi(\tau_h) \cdot \mathbb{P}_\theta(o_{h+1:h+m} = \mathbf{u}_o \mid \tau_h, a_{h+1:h+m-1} = \mathbf{u}_a) \\ \quad = \pi(\tau_h) \cdot (\mathbf{e}_\mathbf{u}^\top \mathbf{B}_h(o_h, a_h; \theta) \cdots \mathbf{B}_1(o_1, a_1; \theta) \mathbf{b}_0(\theta)), \\ \mathbb{P}_\theta(o_{1:m} = \mathbf{u}_o \mid a_{1:m-1} = \mathbf{u}_a) = \mathbf{e}_\mathbf{u}^\top \mathbf{b}_0(\theta), \end{cases} \quad (36)$$

where  $\mathbf{e}_\mathbf{u}$  is the standard basis vector in  $\mathbb{R}^{|\mathbf{U}|}$  corresponding to  $\mathbf{u}$ ,

$$(\{\mathbf{B}_h(o, a; \theta)\}_{(o,a,h) \in \mathcal{O} \times \mathcal{A} \times [H-m]}, \mathbf{b}_0(\theta))$$

denote the operators corresponding to  $\theta$ , and  $\mathbb{P}_\theta(o_{h+1:h+m} = \mathbf{u}_o \mid \tau_h, a_{h+1:h+m-1} = \mathbf{u}_a)$  denotes the probability of observing  $\mathbf{u}_o$  at steps  $h + 1$  to  $h + m$  conditioning on the first- $h$ -steps history  $\tau_h$  and the player following fixed action sequence  $\mathbf{u}_a$  from step  $h + 1$  to  $h + m - 1$ .

For a matrix  $\mathbf{A}$ ,  $p \geq 1$ , let  $\|\mathbf{A}\|_p = \sup_{x: \|x\|_p \leq 1} \|\mathbf{A}x\|_p$ . We start with a lemma that bounds  $\|\mathbf{B}_h(o, a)\|_1$  as a function of  $\alpha$ :

**Lemma 30** *Under Assumption 2 it holds that for any  $(o, a, h, k) \in \mathcal{O} \times \mathcal{A} \times [H - m] \times \mathbb{N}$ ,*

$$\|\mathbf{B}_h(o, a)\|_1, \|\mathbf{B}_h^k(o, a)\|_1 \leq A^{m-1} \sqrt{S} / \alpha.$$

We omit the proof of Lemma 30 here since it is basically the same as that of Lemma 25.

### G.1. Step 1: bound the regret by the error of operator estimates

To begin with, we upper bound the cumulative regret by the summation of density estimation errors using Lemma 26,

$$\sum_{t=1}^k V^{\pi^t}(\theta^t) - V^{\pi^t}(\theta^*) \leq H \sum_{t=1}^k \sum_{\tau_H} |\mathbb{P}_{\theta^t}^{\pi^t}(\tau_H) - \mathbb{P}_{\theta^*}^{\pi^t}(\tau_H)|. \quad (37)$$

As a result, to prove Theorem 7, it suffices to upper bound the RHS of equation (37).

Using the first relation in equation (36), we can represent the probability of observing a trajectory by the product of operators, which gives

$$\begin{aligned} & \sum_{t=1}^k \sum_{\tau_H} \left| \mathbb{P}_{\theta^t}^{\pi^t}(\tau_H) - \mathbb{P}_{\theta^*}^{\pi^t}(\tau_H) \right| \\ &= \sum_{t=1}^k \sum_{\tau_H} \left| \mathbf{e}_{(o_{H-m+1}, \dots, a_{H-1}, o_H)}^\top \mathbf{B}_{H-m}^t(o_{H-m}, a_{H-m}) \cdots \mathbf{B}_1^t(o_1, a_1) \mathbf{b}_0^t \right. \\ & \quad \left. - \mathbf{e}_{(o_{H-m+1}, \dots, a_{H-1}, o_H)}^\top \mathbf{B}_{H-m}(o_{H-m}, a_{H-m}) \cdots \mathbf{B}_1(o_1, a_1) \mathbf{b}_0 \right| \times \pi^t(\tau_H) \\ &\leq \sum_{t=1}^k \sum_{\tau_{H-m}} \left\| \mathbf{B}_{H-m}^t(o_{H-m}, a_{H-m}) \cdots \mathbf{B}_1^t(o_1, a_1) \mathbf{b}_0^t \right. \\ & \quad \left. - \mathbf{B}_{H-m}(o_{H-m}, a_{H-m}) \cdots \mathbf{B}_1(o_1, a_1) \mathbf{b}_0 \right\|_1 \times \pi^t(\tau_{H-m}). \end{aligned}$$

By Lemma 31 (a general version of Lemma 27), we can control the difference between the products of operators by the difference between each pair of operators.

$$\begin{aligned} & \sum_{\tau_{H-m}} \left\| \mathbf{B}_{H-m}^t(o_{H-m}, a_{H-m}) \cdots \mathbf{B}_1^t(o_1, a_1) \mathbf{b}_0^t - \mathbf{B}_{H-m}(o_{H-m}, a_{H-m}) \cdots \mathbf{B}_1(o_1, a_1) \mathbf{b}_0 \right\|_1 \times \pi^t(\tau_{H-m}) \\ & \leq \frac{A^{m-1} \sqrt{S}}{\alpha} \left( \sum_{h=1}^{H-m} \sum_{\tau_h} \left\| (\mathbf{B}_h^t(o_h, a_h) - \mathbf{B}_h(o_h, a_h)) \mathbf{b}(\tau_{h-1}) \right\|_1 \times \pi^t(\tau_h) + \left\| \mathbf{b}_0^t - \mathbf{b}_0 \right\|_1 \right), \end{aligned}$$

where  $\mathbf{b}(\tau_h) = \left( \prod_{h'=1}^h \mathbf{B}_{h'}(o_{h'}, a_{h'}) \right) \mathbf{b}_0$ .

## G.2. Step 2: derive constraints for the operator estimates from OMLE

For simplicity of notation, we denote by  $\mu^{k,h}$  the policy executed in the  $h^{\text{th}}$  inner loop of the  $k^{\text{th}}$  outer loop, i.e.,

$$\mu^{k,h} = \pi_{1:h}^k \circ [\text{Uniform}(\mathcal{A})]_{h+1:H}$$

By Proposition 14, the definition of the confidence set and the choice of  $\beta$ , we have with probability at least  $1 - \delta$ ,

$$\text{for all } k \in [K] \quad \sum_{t=1}^{k-1} \sum_{h=0}^{H-m} \left( \sum_{\tau_H} \left| \mathbb{P}_{\theta^k}^{\mu^{t,h}}(\tau_H) - \mathbb{P}_{\theta^*}^{\mu^{t,h}}(\tau_H) \right| \right)^2 = \mathcal{O}(\beta).$$

In what follows we assume the above equation holds. By the Cauchy-Schwarz inequality,

$$\sum_{t=1}^{k-1} \sum_{h=0}^{H-m} \sum_{\tau_H} \left| \mathbb{P}_{\theta^k}^{\mu^{t,h}}(\tau_H) - \mathbb{P}_{\theta^*}^{\mu^{t,h}}(\tau_H) \right| = \mathcal{O}(\sqrt{\beta k H}).$$

Because marginalizing distributions cannot increase their distance, and using the definition of  $\mu^{t,h}$ , the above relation implies that for any  $(k, h) \in [K] \times [H - m]$ ,

$$\left\{ \begin{array}{l} \frac{1}{A^{m-1}} \sum_{t=1}^{k-1} \sum_{\tau_h, \mathbf{u} \in \mathbf{U}} \left| \mathbb{P}_{\theta^k}^{\pi^t}(\tau_h) \mathbb{P}_{\theta^k}(o_{h+1:h+m} = \mathbf{u}_o \mid \tau_h, a_{h+1:h+m-1} = \mathbf{u}_a) \right. \\ \quad \left. - \mathbb{P}_{\theta^*}^{\pi^t}(\tau_h) \mathbb{P}_{\theta^*}(o_{h+1:h+m} = \mathbf{u}_o \mid \tau_h, a_{h+1:h+m-1} = \mathbf{u}_a) \right| = \mathcal{O}(\sqrt{\beta k H}), \\ \frac{1}{A^m} \sum_{t=1}^{k-1} \sum_{\tau_h, \mathbf{u} \in \mathbf{U}} \left| \mathbb{P}_{\theta^k}^{\pi^t}(\tau_{h-1}) \mathbb{P}_{\theta^k}(o_{h+1:h+m} = \mathbf{u}_o \mid \tau_h, a_{h+1:h+m-1} = \mathbf{u}_a) \right. \\ \quad \left. - \mathbb{P}_{\theta^*}^{\pi^t}(\tau_{h-1}) \mathbb{P}_{\theta^*}(o_{h+1:h+m} = \mathbf{u}_o \mid \tau_h, a_{h+1:h+m-1} = \mathbf{u}_a) \right| = \mathcal{O}(\sqrt{\beta k H}), \\ \frac{1}{A^{m-1}} \sum_{t=1}^{k-1} \sum_{\mathbf{u} \in \mathbf{U}} \left| \mathbb{P}_{\theta^k}(o_{1:m} = \mathbf{u}_o \mid a_{1:m-1} = \mathbf{u}_a) \right. \\ \quad \left. - \mathbb{P}_{\theta^*}(o_{1:m} = \mathbf{u}_o \mid a_{1:m-1} = \mathbf{u}_a) \right|_1 = \mathcal{O}(\sqrt{\beta k H}). \end{array} \right.$$

By using the two relations in Equation (36), we can replace the probability with the product of operators and obtain that for any  $(k, h) \in [K] \times [H - m]$

$$\left\{ \begin{array}{l} \sum_{t=1}^{k-1} \sum_{\tau_h} \pi^t(\tau_h) \times \left\| \mathbf{b}^k(\tau_h) - \mathbf{b}(\tau_h) \right\|_1 = \mathcal{O}(A^{m-1} \sqrt{\beta k H}), \\ \sum_{t=1}^{k-1} \sum_{\tau_h} \pi^t(\tau_{h-1}) \times \left\| \mathbf{b}^k(\tau_h) - \mathbf{b}(\tau_h) \right\|_1 = \mathcal{O}(A^m \sqrt{\beta k H}), \\ \left\| \mathbf{b}_0^k - \mathbf{b}_0 \right\|_1 = \mathcal{O}(A^{m-1} \sqrt{\beta H/k}), \end{array} \right. \quad (38)$$

where  $\mathbf{b}(\tau_h) = \left( \prod_{h'=1}^h \mathbf{B}_{h'}(o_{h'}, a_{h'}) \right) \mathbf{b}_0$  and  $\mathbf{b}^k(\tau_h) = \left( \prod_{h'=1}^h \mathbf{B}_{h'}^k(o_{h'}, a_{h'}) \right) \mathbf{b}_0^k$ .

Now, we are ready to derive the final guarantee for each individual operator estimate  $\mathbf{B}_h^k(o, a)$ . For all  $(k, h) \in [K] \times [H - m]$ ,

$$\begin{aligned} & \sum_{t=1}^{k-1} \sum_{\tau_h} \pi^t(\tau_{h-1}) \times \left\| \left( \mathbf{B}_h^k(o_h, a_h) - \mathbf{B}_h(o_h, a_h) \right) \mathbf{b}(\tau_{h-1}) \right\|_1 \\ & \leq \sum_{t=1}^{k-1} \sum_{\tau_h} \pi^t(\tau_{h-1}) \times \left\| \mathbf{B}_h^k(o_h, a_h) \left( \mathbf{b}^k(\tau_{h-1}) - \mathbf{b}(\tau_{h-1}) \right) \right\|_1 \\ & \quad + \sum_{t=1}^{k-1} \sum_{\tau_h} \pi^t(\tau_{h-1}) \times \left\| \mathbf{B}_h^k(o_h, a_h) \mathbf{b}^k(\tau_{h-1}) - \mathbf{B}_h(o_h, a_h) \mathbf{b}(\tau_{h-1}) \right\|_1. \end{aligned} \quad (39)$$

For the first term in the RHS of Equation (39):

$$\begin{aligned} & \sum_{t=1}^{k-1} \sum_{\tau_h} \pi^t(\tau_{h-1}) \times \left\| \mathbf{B}_h^k(o_h, a_h) \left( \mathbf{b}^k(\tau_{h-1}) - \mathbf{b}(\tau_{h-1}) \right) \right\|_1 \\ & = A \sum_{t=1}^{k-1} \sum_{\tau_h} \pi^{t, h-1}(\tau_h) \times \left\| \mathbf{B}_h^k(o_h, a_h) \left( \mathbf{b}^k(\tau_{h-1}) - \mathbf{b}(\tau_{h-1}) \right) \right\|_1 \\ \text{by Lemma 32} & \leq \frac{A^m \sqrt{S}}{\alpha} \times \sum_{t=1}^{k-1} \sum_{\tau_{h-1}} \pi^{t, h-1}(\tau_{h-1}) \times \left\| \mathbf{b}^k(\tau_{h-1}) - \mathbf{b}(\tau_{h-1}) \right\|_1 \\ & = \frac{A^m \sqrt{S}}{\alpha} \times \sum_{t=1}^{k-1} \sum_{\tau_{h-1}} \pi^t(\tau_{h-1}) \times \left\| \mathbf{b}^k(\tau_{h-1}) - \mathbf{b}(\tau_{h-1}) \right\|_1 \\ & = \mathcal{O} \left( \frac{A^{2m-1} \sqrt{S}}{\alpha} \sqrt{\beta k H} \right), \end{aligned}$$

where the last equality uses the first relation in Equation (38).

For the second term in the RHS of Equation (39): by the second relation in Equation (38).

$$\sum_{t=1}^{k-1} \sum_{\tau_h} \pi^t(\tau_{h-1}) \times \left\| \mathbf{B}_h^k(o_h, a_h) \mathbf{b}^k(\tau_{h-1}) - \mathbf{B}_h(o_h, a_h) \mathbf{b}(\tau_{h-1}) \right\|_1 = \mathcal{O}(A^m \sqrt{\beta k H}).$$

Therefore, we conclude that for all  $(k, h) \in [K] \times [H - m]$ ,

$$\sum_{t=1}^{k-1} \sum_{\tau_h} \pi^t(\tau_{h-1}) \times \left\| \left( \mathbf{B}_h^k(o_h, a_h) - \mathbf{B}_h(o_h, a_h) \right) \mathbf{b}(\tau_{h-1}) \right\|_1 = \mathcal{O} \left( \frac{A^{2m-1} \sqrt{S}}{\alpha} \sqrt{\beta k H} \right).$$

### G.3. Step 3: bridge Step 1 and 2 via $\ell_1$ -norm eluder dimension

By Step 1, to control the density estimation error, it suffices to upper bound the following quantity for all  $k \in [K]$ :

$$\text{Target: } \sum_{t=1}^k \left( \sum_{h=1}^{H-m} \sum_{\tau_h} \left\| \left( \mathbf{B}_h^t(o_h, a_h) - \mathbf{B}_h(o_h, a_h) \right) \mathbf{b}(\tau_{h-1}) \right\|_1 \times \pi^t(\tau_{h-1}) + \left\| \mathbf{b}_0^t - \mathbf{b}_0 \right\|_1 \right).$$

By Step 2, we have with probability at least  $1 - \delta$ : for all  $(k, h) \in [K] \times [H - m]$

$$\text{Condition: } \begin{cases} \|\mathbf{b}_0^k - \mathbf{b}_0\|_1 = \mathcal{O}(A^{m-1}\sqrt{\beta H/k}), \\ \sum_{t=1}^{k-1} \sum_{\tau_h} \|(\mathbf{B}_h^k(o_h, a_h) - \mathbf{B}_h(o_h, a_h)) \mathbf{b}(\tau_{h-1})\|_1 \times \pi^t(\tau_{h-1}) = \mathcal{O}\left(\frac{A^{2m-1}\sqrt{S}}{\alpha} \sqrt{\beta k H}\right). \end{cases}$$

Below, we will carry on the proof conditioning on the event that the above relations hold.

The summation of the second term in the target can be upper bounded by using the first condition:

$$\mathcal{O}\left(\sum_{t=1}^k A^{m-1}\sqrt{\beta H/t}\right) = \mathcal{O}\left(A^{m-1}\sqrt{\beta k H}\right).$$

It remains to control the summation of the first term.

Let's consider an *arbitrary fixed*  $h \in [H - m]$ . Denote by  $\mathbf{X}_l$  the  $l^{\text{th}}$  row of matrix  $\mathbf{X}$ . The second condition is equivalent to

$$\begin{aligned} & \sum_{t=1}^{k-1} \sum_{l=1}^{O^m A^{m-1}} \sum_{o,a} \sum_{\tau_{h-1}} \left| \left[ (\mathbf{B}_h^k(o, a) - \mathbf{B}_h(o, a)) \mathbb{M}_h \right]_l \mathbb{M}_h^\dagger \mathbf{b}(\tau_{h-1}) \right| \times \pi^t(\tau_{h-1}) \\ & = \mathcal{O}\left(\frac{A^{2m-1}\sqrt{S}}{\alpha} \sqrt{\beta k H}\right). \end{aligned} \quad (40)$$

To further simplify the notations, for  $(t, l, o, a) \in [K] \times [O^m A^{m-1}] \times \mathcal{O} \times \mathcal{A}$ , let  $w_{t,l,o,a} := \left[ (\mathbf{B}_h^t(o, a) - \mathbf{B}_h(o, a)) \mathbb{M}_h \right]_l$  and denote the  $n := (OA)^{h-1}$  not necessarily distinct elements in the sequence  $\{\mathbb{M}_h^\dagger \mathbf{b}(\tau_{h-1}) \times \pi^t(\tau_h) : \tau_h = (o_1, a_1, \dots, o_{h-1}, a_{h-1}, o, a) \in (\mathcal{O} \times \mathcal{A})^{h-1} \times (o, a)\}$  by  $x_{t,1}, \dots, x_{t,n}$ . Using the newly defined notations, Equation (40) is equivalent to

$$\sum_{t=1}^{k-1} \sum_{l=1}^{O^m A^{m-1}} \sum_{o,a} \sum_{i=1}^n |w_{k,l,o,a}^\top x_{t,i}| = \mathcal{O}\left(\frac{A^{2m-1}\sqrt{S}}{\alpha} \sqrt{\beta k H}\right). \quad (41)$$

Now, we invoke Proposition 22 with Equation (41) and the following normalization conditions: for all  $t \in \mathbb{N}$

- The vectors  $\{x_{t,i}\}_{i=1}^n$  satisfy  $\sum_{i=1}^n \|x_{t,i}\|_1 \leq 1$ , which follows from the same argument as in Appendix F.3),
- For the vectors  $\{w_{t,l,o,a}\}_{l,o,a}$ , we have

$$\sum_{l=1}^{O^m A^{m-1}} \sum_{o,a} \|w_{t,l,o,a}\|_1 = \sum_{(o,a) \in \mathcal{O} \times \mathcal{A}} \sum_{s=1}^S \|(\mathbf{B}_h(o, a) - \mathbf{B}_h^t(o, a)) \mathbb{M}_h \mathbf{e}_s\|_1 \leq \frac{2S^{1.5} A^{2m-1}}{\alpha},$$

where the inequality uses  $\|\mathbb{M}_h \mathbf{e}_s\|_1 = A^{m-1}$  and Lemma 32 with  $\pi = \text{Uniform}(\mathcal{A})$ .

As a result, we obtain

$$\sum_{t=1}^k \sum_{l=1}^{O^m A^{m-1}} \sum_{i=1}^n |w_{t,l}^\top x_{t,i}| = \tilde{\mathcal{O}}\left(\frac{S^{1.5} H^2 A^{2m-1}}{\alpha} \sqrt{\beta k H}\right), \text{ for all } k \in [K]. \quad (42)$$



Translating the above guarantee back using the operator language gives that for all  $k \in [K]$ ,

$$\begin{aligned} & \sum_{t=1}^k \left( \sum_{h=1}^{H-1} \sum_{\tau_h} \left\| (\mathbf{B}_h^t(o_h, a_h) - \mathbf{B}_h(o_h, a_h)) \mathbf{b}(\tau_{h-1}) \right\|_1 \times \pi^t(\tau_{h-1}) + \left\| \mathbf{b}_0^t - \mathbf{b}_0 \right\|_1 \right) \\ &= \tilde{\mathcal{O}} \left( \frac{S^{1.5} H^3 A^{2m-1}}{\alpha} \sqrt{\beta k H} \right). \end{aligned}$$

Combining the above equality with Step 1 gives:

$$\sum_{t=1}^k \left( V^* - V^{\pi^t} \right) \leq \tilde{\mathcal{O}} \left( \frac{S^2 A^{3m-2}}{\alpha^2} \sqrt{k(S^2 A + SO)} \times \text{poly}(H) \right) \quad \text{for all } k \in [K].$$

#### G.4. Auxiliary lemmas

We prove the following auxiliary lemma for general POMDPs. The definitions of the notations can be found in the beginning of Appendix G.

**Lemma 31** *Suppose Assumption 2 holds. For any  $k \in \mathbb{N}$ ,  $h \in [H - m]$  and policy  $\pi$*

$$\begin{aligned} & \sum_{\tau_h} \left\| \mathbf{B}_h^k(o_h, a_h) \cdots \mathbf{B}_1^k(o_1, a_1) \mathbf{b}_0^k - \mathbf{B}_h(o_h, a_h) \cdots \mathbf{B}_1(o_1, a_1) \mathbf{b}_0 \right\|_1 \times \pi(\tau_h) \\ & \leq \frac{A^{m-1} \sqrt{S}}{\alpha} \left( \sum_{j=1}^h \sum_{\tau_j} \left\| (\mathbf{B}_j^k(o_j, a_j) - \mathbf{B}_j(o_j, a_j)) \mathbf{b}(\tau_{j-1}) \right\|_1 \times \pi(\tau_j) + \left\| \mathbf{b}_0^k - \mathbf{b}_0 \right\|_1 \right), \end{aligned}$$

where  $\mathbf{b}(\tau_h) = \left( \prod_{h'=1}^h \mathbf{B}_{h'}(o_{h'}, a_{h'}) \right) \mathbf{b}_0$ .

**Proof** [Proof of Lemma 31] To reduce clutter, we abbreviate  $\mathbf{B}_h^k(o_h, a_h)$  as  $\mathbf{B}_h^k$  and denote  $\mathbf{B}_{h:j}^k := \mathbf{B}_h^k \times \cdots \times \mathbf{B}_j^k$ . By triangle inequality,

$$\begin{aligned} & \sum_{\tau_h} \left\| \mathbf{B}_h^k \cdots \mathbf{B}_1^k \mathbf{b}_0^k - \mathbf{B}_h \cdots \mathbf{B}_1 \mathbf{b}_0 \right\|_1 \times \pi(\tau_h) \\ & \leq \sum_{j=1}^h \sum_{\tau_h} \left\| \mathbf{B}_{h:j+1}^k (\mathbf{B}_j^k - \mathbf{B}_j) \mathbf{b}(\tau_{j-1}) \right\|_1 \times \pi(\tau_h) + \sum_{\tau_h} \left\| \mathbf{B}_{h:1}^k (\mathbf{b}_0^k - \mathbf{b}_0) \right\|_1 \times \pi(\tau_h). \end{aligned}$$

So it suffices prove the following Lemma:

**Lemma 32** *For any index  $j < h \leq H - m$ , trajectory  $\tau_j \in (\mathcal{O} \times \mathcal{A})^j$ , vector  $x \in \mathbb{R}^{O^m A^{m-1}}$ , policy  $\pi$ , and operator  $\tilde{\mathbf{B}} \in \{\mathbf{B}, \mathbf{B}^k\}$ , we have*

$$\sum_{\tau_{h:j+1} \in (\mathcal{O} \times \mathcal{A})^{h-j}} \left\| \tilde{\mathbf{B}}_{h:j+1} x \right\|_1 \times \pi(\tau_{h:j+1} | \tau_j) \leq \frac{A^{m-1} \sqrt{S}}{\alpha} \|x\|_1.$$

**Proof** [Proof of Lemma 32] We only prove the lemma for  $\tilde{\mathbf{B}} = \mathbf{B}^k$  since the other case follows exactly the same arguments. By the definition of  $\mathbf{B}_{j+1}^k$ , we have that the row space of  $\mathbf{B}_{h:j+1}^k$  belongs to the column space of  $\mathbb{M}_{j+1}^k$ , which implies

$$\mathbf{B}_{h:j+1}^k x = \mathbf{B}_{h:j+1}^k \mathbb{M}_{j+1}^k (\mathbb{M}_{j+1}^k)^\dagger x.$$

Moreover, for any standard basis  $\mathbf{e}_i \in \mathbb{R}^S$

$$\begin{aligned} & \sum_{\tau_{h:j+1}} \|\mathbf{B}_{h:j+1}^k \mathbb{M}_j^k \mathbf{e}_i\|_1 \times \pi(\tau_{h:j+1} \mid \tau_j) \\ &= \sum_{\mathbf{a} \in \mathcal{A}^{m-1}} \sum_{\mathbf{o} \in \mathcal{O}^m} \sum_{\tau_{h:j+1}} \mathbb{P}_{\theta^k}^\pi(\tau_{h:j+1} \mid s_{j+1} = i, \tau_j) \cdot \mathbb{P}_{\theta^k}(o_{h+1:h+m} = \mathbf{o} \mid \tau_h, a_{h+1:h+m-1} = \mathbf{a}) \\ &= A^{m-1}. \end{aligned}$$

Combining all relations together, we obtain

$$\sum_{\tau_{h:j+1}} \|\mathbf{B}_{h:j+1}^k x\|_1 \times \pi^k(\tau_{h:j+1} \mid \tau_j) \leq A^{m-1} \|(\mathbb{M}_{j+1}^k)^\dagger x\|_1 \leq \frac{A^{m-1} \sqrt{S}}{\alpha} \|x\|_1,$$

where the final inequality follows from Assumption 2. ■

## Appendix H. Proofs for Lower Bounds

In this section, we prove the two lower bounds presented in Section 4. ■

### H.1. Proof of Theorem 6

**Proof** [Proof of Theorem 6] We construct the hard instance based on combinatorial lock, which we define formally as follows.

1. STATE: We have 2 states, with 1 “good state”  $s_g$  and 1 “bad state”  $s_b$ .
2. OBSERVATION: There are three observations  $o_g, o_b$  and  $o_{\text{dummy}}$ . The emission matrices in the first  $H - 1$  steps are

$$\begin{pmatrix} \alpha & 0 \\ 0 & \alpha \\ 1 - \alpha & 1 - \alpha \end{pmatrix},$$

while the emission matrix at step  $H$  is

$$\begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{pmatrix}.$$

In other words, at each step  $h \in [H - 1]$ , with probability  $\alpha$  we observe the current latent state, and with probability  $1 - \alpha$  we receive a dummy observation. That is, with probability

$\alpha$  we find out what state we are currently at, and with probability  $1 - \alpha$ , we learn nothing. But at step  $H$ , we always directly observe the current latent state. By standard linear algebra, one can verify  $\min_h \sigma_S(\mathbb{O}_h) \geq \alpha$ .

3. **ACTION AND TRANSITION:** There are  $A$  different actions and the initial state at step 1 is fixed as  $s_g$ . We now define the transitions as follows. For every  $h \in [H - 1]$ , we designate one of these actions to be “good”, the others “bad”. If one is currently in the “good” state and takes the “good” action, then the agent will transition to the “good” state. Otherwise, one will always transition to the “bad” state. For each  $h \in [H - 1]$ , the good action is chosen *uniformly at random* from  $\mathcal{A}$ . The current episode terminates immediately after  $o_H$  is received.
4. **REWARD:** Viewing every observation at the first  $H - 1$  steps yields a reward of 0. At step  $H$ , viewing  $o_g$  yields reward 1 while viewing  $o_b$  yields reward 0. As a result, the agent receives reward 1 only if the state  $s_g$  is reached at step  $H$ , i.e., if the good action is taken at every step in the POMDP.

**Showing large number of samples are necessary:** Suppose that we attempt to learn this POMDP with an algorithm  $\mathcal{A}$ , where we are allowed to interact with the POMDP for  $K \leq \lfloor \frac{1}{2\alpha H} \rfloor$  episodes. Now, consider the probability that in these  $K$  episodes, *both  $s_g$  and  $s_b$  only emit dummy observations* in the first  $H - 1$  steps and the agent always gets reward 0. We can write this as

$$\begin{aligned} & \mathbb{P}(\text{only emit dummy observations in the first } H - 1 \text{ steps and get reward 0 in all episodes}) \\ &= \mathbb{P}(\text{get reward 0 in all episodes} \mid \text{only emit dummy observations in the first } H - 1 \text{ steps}) \\ & \quad \times \mathbb{P}(\text{only emit dummy observations in the first } H - 1 \text{ steps}). \end{aligned}$$

The probability that both  $s_g$  and  $s_b$  only emit dummy observations in the first  $H - 1$  steps of all  $K$  episodes is upper bounded by  $(1 - \alpha)^{1/\alpha}$  because  $s_g$  and  $s_b$  together can emit at most  $2 \times \lfloor \frac{1}{2\alpha H} \rfloor \times (H - 1)$  observations in the first  $H - 1$  steps of  $K$  episodes.

Now, conditioned on the event that both  $s_g$  and  $s_b$  only emit dummy observations in the first  $H - 1$  steps of all  $K$  episodes, in each episode, the best strategy we can use is to randomly guess the sequence of optimal actions. In particular, the probability that  $\mathcal{A}$  fails to guess the optimal sequence correctly, given that we have  $K$  guesses, is  $\binom{A^{H-1}-1}{K} / \binom{A^{H-1}}{K} = \frac{A^{H-1}-K}{A^{H-1}}$ . Then, for  $K \leq A^{H-1}/10$ , this is at least  $9/10$ . Thus, with probability at least  $0.9 \times (1 - \alpha)^{1/\alpha} \geq 1/6$ , the agent learns nothing except that the action sequences it chose are incorrect, and because it only sampled at most one tenth of the total possible action sequences by assumption, the best policy it can output is just to randomly guess from the remaining action sequences, which is worse than  $(1/2)$ -optimal, which completes the proof.  $\blacksquare$

**Lemma 33** For  $0 < \alpha \leq \frac{1}{2}$ ,  $(1 - \alpha)^{1/\alpha} \geq \frac{1}{5}$ .

**Proof** It suffices to prove

$$1 - \alpha \geq \frac{1}{5^\alpha} \Leftrightarrow \alpha + \left(\frac{1}{5^\alpha} - 1\right) \leq 0.$$

We note that the LHS is an increasing function and that for  $\alpha = \frac{1}{2}$ , the LHS is

$$\frac{1}{\sqrt{5}} - \frac{1}{2} \leq 0,$$

which suffices to show the desired inequality for all  $\alpha \leq \frac{1}{2}$ . ■

## H.2. Lower bound for the exponential dependency on $m$

**Proof** [Proof of Theorem 9] We construct the hard instance based on combinatorial lock, which we define formally as follows. We choose the episode length  $H$  equal to  $m$ .

1. STATE: We have 2 states, with 1 “good state”  $s_g$  and 1 “bad state”  $s_b$ .
2. OBSERVATION AND REWARD: We have three observations.
  - (a) Dummy Observation: at the first  $m - 1$  steps, we view this observation with probability 1 regardless of our current latent state, which gives reward 0.
  - (b) Reward 1 Observation: at step  $m$  and state  $s_g$  we view this observation with probability 1, which gives reward 1.
  - (c) Reward 0 Observation: at step  $m$  and state  $s_b$  we view this observation with probability 1, which gives reward 0.
3. ACTIONS AND TRANSITIONS: *We define the actions and transitions identically as in the proof of Theorem 6.*

It is easy to verify the  $m$ -step emission-action matrix at step  $h = 1$  has rank 2 and minimum singular value no smaller than 1 (by using Proposition 3 with the optimal action sequence). Therefore, it satisfies Assumption 2 with  $\alpha \geq 1$ . (Since  $H = m$ , Assumption 2 only requires  $\sigma_S(\mathbb{M}_1) \geq 1$ .) Moreover, we claim that the agent must take  $\Omega(A^{m-1})$  samples in order to learn this POMDP. This is because in the first  $m - 1$  steps, the reward is always 0 and the observation is always dummy, so that no knowledge is learned. The only useful information is the reward feedback in the final step. Therefore, this is equivalent to a multi-arm bandit problem with  $A^{m-1}$  arms. ■

## Appendix I. Proofs for Weakly Revealing Conditions

### I.1. Proof of Proposition 2

**Proof** By the definition of minimum singular value,  $\sigma_S(\mathbb{O}_h) \neq 0$  implies that for any  $\mu \neq \mu'$ ,  $\mathbb{O}_h(\mu - \mu') \neq \mathbf{0}$ . On the other hand, if  $\sigma_S(\mathbb{O}_h) = 0$ , then there exists  $z \neq \mathbf{0}$  such that  $\mathbb{O}_h z = \mathbf{0}$ . Let  $z^+ = \max\{z, 0\}$  and  $z^- = -\min\{z, 0\}$ , where the max and min are taken entry-wisely. By definition,  $\mathbb{O}_h z^+ = \mathbb{O}_h z^-$ . Moreover,  $\|z^+\|_1 = \|\mathbb{O}_h z^+\|_1 = \|\mathbb{O}_h z^-\|_1 = \|z^-\|_1$ . As a result, we have  $\mathbb{O}_h \frac{z^+}{\|z^+\|_1} = \mathbb{O}_h \frac{z^-}{\|z^-\|_1}$  where  $\frac{z^+}{\|z^+\|_1}$  and  $\frac{z^-}{\|z^-\|_1}$  are two disjoint state distributions that induce the same distribution over observations. ■

### I.2. Proof of Proposition 3

**Proof** We prove Proposition 3 by showing that for any fixed  $h$ ,  $\sigma_S(\mathbb{M}_h) \geq \max_{\mathbf{a} \in \mathcal{A}^{m-1}} \sigma_S(\mathbb{M}_{h,\mathbf{a}})$ . By the Courant-Fischer-Weyl min-max principle, for matrix  $U \in \mathbb{R}^{n \times S}$  with  $n \geq S$ ,

$$\sigma_S^2(U) = \min_{y \in \mathbb{R}^S: \|y\|_2=1} \|Uy\|_2^2 = \min_{y \in \mathbb{R}^S: \|y\|_2=1} \sum_{i=1}^n (U_i \cdot y)^2,$$

where  $U_i$  stands for the  $i$ th row of  $U$ . Fix  $h \in [H]$  and by abusing notation let  $\mathbb{M} = \mathbb{M}_h$ . Then,

$$\begin{aligned} \sigma_S^2(\mathbb{M}) &= \min_{y \in \mathbb{R}^S: \|y\|_2=1} \sum_{\mathbf{a}} \sum_{\mathbf{o}} (\mathbb{M}_{(\mathbf{a},\mathbf{o})} \cdot y)^2 \geq \sum_{\mathbf{a}} \min_{y \in \mathbb{R}^S: \|y\|_2=1} \sum_{\mathbf{o}} (\mathbb{M}_{(\mathbf{a},\mathbf{o})} \cdot y)^2 \\ &\geq \max_{\mathbf{a}} \sigma_S^2(\mathbb{M}_{h,\mathbf{a}}). \end{aligned}$$

■

### I.3. Relation between the $\gamma$ -observability (Golowich et al., 2022) and the weakly revealing conditions

**Lemma 34** Suppose  $\phi, \psi \in \mathbb{R}^d$  satisfy  $\sum_i \phi_i = 0$  and  $(\max_i \psi_i) \times (\min_i \psi_i) \geq 0$ , then

$$\sum_i |\phi_i + \psi_i| \geq \frac{1}{2} \max\{\|\phi\|_1, \|\psi\|_1\}.$$

**Proof** WLOG, assume  $\min_i \psi_i \geq 0$ . By triangle inequality,  $\sum_i |\phi_i + \psi_i| \geq |\sum_i \phi_i + \sum_i \psi_i| = \|\psi\|_1$ . Moreover,  $\sum_i |\phi_i + \psi_i| \geq \sum_{i: \phi_i \geq 0} (\phi_i + \psi_i) \geq \frac{1}{2} \|\phi\|_1$ . ■

**Lemma 35** Suppose  $\gamma$  and  $\alpha$  are the largest real numbers that satisfy

1. For any  $v_1, v_2 \in \Delta_S$ ,  $\|\mathbb{O}_h(v_1 - v_2)\|_1 \geq \gamma \|v_1 - v_2\|_1$ ,
2.  $\sigma_S(\mathbb{O}_h) \geq \alpha$ .

Then  $\frac{\alpha}{\sqrt{S}} \leq \gamma \leq 4\sqrt{O}\alpha$ .

**Proof** Suppose  $\sigma_S(\mathbb{O}_h) \geq \alpha$ . We have

$$\|\mathbb{O}_h(v_1 - v_2)\|_1 \geq \|\mathbb{O}_h(v_1 - v_2)\|_2 \geq \alpha \|v_1 - v_2\|_2 \geq \frac{\alpha}{\sqrt{S}} \|v_1 - v_2\|_1,$$

which implies  $\gamma \geq \frac{\alpha}{\sqrt{S}}$ .

Suppose for any  $v_1, v_2 \in \Delta_S$ ,  $\|\mathbb{O}_h(v_1 - v_2)\|_1 \geq \gamma \|v_1 - v_2\|_1$ . Consider an arbitrary  $z \in \mathbb{R}^O$  and decompose  $z$  as  $z = z^+ - z^- + \bar{z}$ , where  $\bar{z} = \frac{\sum_i z_i}{O} \times \mathbf{1}$ ,  $z^+ = \max\{z - \bar{z}, 0\}$  and  $z^- = \max\{\bar{z} - z, 0\}$ . Invoking Lemma 34 with  $\phi = \mathbb{O}_h(z^+ - z^-)$  and  $\psi = \mathbb{O}_h \bar{z}$ , we obtain

$$\|\mathbb{O}_h z\|_2 \geq \frac{1}{2\sqrt{O}} \max\{\|\mathbb{O}_h(z^+ - z^-)\|_1, \|\mathbb{O}_h \bar{z}\|_1\}.$$

Note that  $\frac{z^+}{\|z^+\|_1}, \frac{z^-}{\|z^-\|_1} \in \Delta_S$  and  $\|z^+\|_1 = \|z^-\|_1$ , so we have  $\|\mathbb{O}_h(z^+ - z^-)\|_1 \geq \gamma\|z^+ - z^-\|_2$ . Besides, by the definition of  $\bar{z}$  and  $\mathbb{O}_h$ ,  $\|\mathbb{O}_h\bar{z}\|_1 = \|\bar{z}\|_1$ , which implies  $\|\mathbb{O}_h\bar{z}\|_1 \geq \|\bar{z}\|_2$ . As a result, we conclude that

$$\|\mathbb{O}_h z\|_2 \geq \frac{\gamma}{2\sqrt{O}} \max\{\|z^+ - z^-\|_2, \|\bar{z}\|_2\} \geq \frac{\gamma}{4\sqrt{O}} \|z\|_2.$$

■