PAC Bounds for Imitation and Model-based Batch Learning of Contextual Markov Decision Processes

Yash Nair
Harvard University
Cambridge, MA 02138
yashnair@college.harvard.edu

Finale Doshi-Velez
Harvard University
Cambridge, MA 02138
finale@seas.harvard.edu

Abstract

We consider the problem of batch multi-task reinforcement learning with observed context descriptors, motivated by its application to personalized medical treatment. In particular, we study two general classes of learning algorithms: direct policy learning (DPL), an imitation-learning based approach which learns from expert trajectories, and model-based learning. First, we derive sample complexity bounds for DPL, and then show that model-based learning from expert actions can, even with a finite model class, be impossible. After relaxing the conditions under which the model-based approach is expected to learn by allowing for greater coverage of state-action space, we provide sample complexity bounds for model-based learning with finite model classes, showing that there exist model classes with sample complexity exponential in their statistical complexity. We then derive a sample complexity upper bound for model-based learning based on a measure of concentration of the data distribution. Our results give formal justification for imitation learning over model-based learning in this setting.

1 Introduction

Families of context-dependent tasks are common in many real-world settings. For example, controlling a UAV might depend on factors such as the parameters of the specific UAV's weight and wingspan. After successfully controlling several different UAVs, one might hope to be able to control a new UAV quickly. Similarly, managing hypotension well may depend on some specific properties of the patient; after treating many distinct patients, one may hope to manage a new patient well.

The question of efficiently learning a collection of related, context-dependent tasks has been studied in the reinforcement learning (RL) literature under many names such as lifelong RL, multi-task RL, and, more generally, transfer learning (see, e.g., Isele et al. [2017], D'Eramo et al. [2019]), and Taylor and Stone [2009]). Even more specifically, the question of learning to *generalize* from a collection of data has been considered, for example, in Brunskill and Li [2013] and Lazaric and Restelli [2011]. These works consider the problem in an online setting and develop algorithmic contributions in the batch setting, respectively.

In this work, motivated by the problem of treating patients with personalized strategies, we consider the following setting and question: Suppose we are given a batch of trajectories obtained from experts in multiple contexts, where each context's transition is parametrized by some observed parameter θ—a framework called a Contextual MDP (CMDP). Will it be more sample efficient to directly learn a policy from these data (that is, imitate the expert), or to learn the transition function, parametrized by θ, and then plan according to it? We derive upper and lower sample-complexity bounds for direct policy learning (DPL); our upper bound for DPL is polynomial in all the relevant parameters.

Along the way, we prove impossibility results for learning certain transitions in the model-based paradigm, while our sample complexity upper bound for direct policy learning holds in a more general sense. Next, we show that, under a relaxed data generation process which affords greater coverage of state-action space, their exist hard families of CMDPs for model-based learning; our results extending those of Chen and Jiang [2019] and Krishnamurthy et al. [2016] to the multiple context setting. Finally, we derive a distribution-dependent sample complexity upper bound for model-based learning.

Our theory provides a formal justification for why imitation may be more successful than model-based learning in these settings, confirming trends observed in several more empirical and application-oriented works, including Yao et al. [2018] and Yang et al. [2019]. In particular, we note that while both DPT and model-based learning depend linearly on the complexity of the hypothesis class, model-based learning requires an extra dependence on the concentratability of the data distribution which can, in some cases, render learning hard. Our results extend the sample complexity bounds derived for data collected online in Brunskill and Li [2013] to the batch imitation learning and model-based context, as well as some of the single MDP results of Piot et al. [2013] to CMDPs.

2 Related Work

Some papers have considered the problem when the context is observable; that is, the learner is allowed to see some high-level labeling of each MDP in the training set as well as that of the target MDP before being asked to return a policy for the target MDP. In particular, Isele et al. [2017] describe the problem as zero-shot transfer learning and provide an algorithm which, under certain linearity assumptions regarding the task descriptions, is able to perform zero-shot transfer on large classes of MDPs. While they consider the problem both empirically and theoretically, they provide only convergence results and not a finite-sample analysis. This framework has also been studied in a more applied setting by Sohn et al. [2018], who, under the further assumption that the learner is given access to not only task descriptions, but also a graph describing relationships among tasks, embed the graph and use a non-parametric gradient-based policy to obtain a policy for the target MDP. Both works differ most notably from ours in that we make no assumption on the structure of the task descriptors nor the relationships between them and we also derive sample complexity results.

The problem has also been cast as an imitation learning problem. Osa et al. [2018] describe *behavior cloning*, the analog of DPL on a single MDP. They describe various algorithms and supervised learning techniques with which the learner can use to learn the expert policy, but do not give any sample complexity bounds. Furthermore, Piot et al. [2013] consider Apprenticeship Learning, whereby the learner has access to a set of states and expert actions and is tasked with learning a policy. They upper bound the difference in the value function of the learner with that of the expert in terms of the classification error of the learner's hypothesis. Our work most notably differs from these two in that we consider the contextual setting: DPL reduces to these works, in the case when $|\Theta| = 1$, and thus only a single MDP is being considered.

Other works have viewed the problem in the domain where the task descriptors are not available to the learner (i.e. the parameter θ is latent). In particular, Lazaric and Restelli [2011] consider the problem of transfer learning between MDPs in the batch setting, where they design the *All-sample Transfer* and *Best Average Transfer* algorithms, both Q-value approximation algorithms designed to learn from a batch of trajectories to plan on a new MDP with potentially different transitions. We approach the problem of model-based learning in a different way from their approach, and instead extend off the analysis of Chen and Jiang [2019]. In the case where transfer is not considered, Sun et al. [2018] consider the model-based approach on contextual decision processes (CDPs). They show a poly($L, 1/\epsilon, \log(1/\delta)$) sample complexity is attainable for the model-based approach, in the batch setting. Our problem of transfer learning on multiple MDPs reduces to theirs in the case when $|\Theta| = 1$. Finally, Yao et al. [2018] consider both a direct policy approach to the unobservable problem—direct policy transfer (DPT)—as well as a model-based approach in an empirical setting. They evaluate both learning strategies on a 2D Navigation task, Acrobot, and a simulated HIV treatment domain, and find that, empirically, DPT is more sample-efficient than the model-based approach.

3 Notation and Background

CMDPs A CMDP is a tuple $\langle S, A, R, T, L, \gamma, \Theta, \mathcal{P}_{s_0}, \mathcal{P}_\Theta \rangle$ where $S \subset \mathbb{R}^n$ and A denote the state and (discrete) action space, respectively; $R : S \times A \to \mathbb{R}$ is a reward function; γ, a discount factor used in evaluating long-term return; $\Theta \subset \mathbb{R}^d$; and $T : S \times A \times \Theta \to \Delta(S)$, is a conditional density over next states given the current state, action, and value of θ parametrizing the transition. The initial state s_0 is drawn from $\mathcal{P}_{s_0}$; at the start of an episode, the task parameter θ is drawn from some $\mathcal{P}_\theta$. This is a slight simplification of the standard definition of CMDP of Hallak et al. [2015] in that we assume the reward and initial distribution remain the same.

Setting In this paper, we shall consider the case in which the learner is given a batch of m trajectories of length L, each labeled with its associated context θ; that is, for each trajectory, the learner is allowed to see the value of θ corresponding to the MDP from which it was drawn. There is one trajectory per parameter setting θ, corresponding to real settings in which one only gets to treat each patient once. This is in contrast to the setting of Hidden Parameter MDPs (HiP-MDPs), introduced by Doshi-Velez and Konidaris [2016], in which the parameter θ is latent. We will further assume that the learner has access to the reward function, R.

Throughout this paper, we will assume that the trajectories are formed via following a deterministic (but possibly time-dependent) expert policy π that is α-optimal; that is, $v_\pi^L \geq v_{\pi^*}^L - \alpha$, where π^* is the optimal deterministic time-dependent policy and $v_\pi^L := \mathbb{E}[\sum_{l=0}^{L-1} r_l]$ is the expected undiscounted value associated with the first L rewards, where here the expectation also includes the randomness with respect to the draw of $\theta \sim \mathcal{P}_\Theta$. We also define $V_\pi^L(s; \theta) = \mathbb{E}[\sum_{l=0}^{L-1} r_l | s_0 = s, \theta]$.

Goal Under this data generation process, the learner's goal is to return a policy $\hat{\pi} : S \times \Theta \times \{0, \ldots, L-1\} \to A$, such that its value $v_{\hat{\pi}}^l$ is maximized. Specifically, we will define the error in value of a hypothesis below:

Definition 3.0.1. *We define the error (of reinforcement learning) of a policy $\hat{\pi}$ to be the difference in undiscounted value $v_{\pi^*}^L - v_{\hat{\pi}}^L$.*

Formally, the goal is thus to return a policy with small error in value.

Following learning theory terminology, we shall also refer to $\hat{\pi}$ as the *hypothesis* returned by the learner.

4 Sample Complexity Bounds for DPL

We now turn to our learning problems. One approach to learning a hypothesis $\hat{\pi}$ above is simply to treat the problem as a supervised learning problem and directly learn the association between the inputs—the states s and the task parameters θ—and the expert's action a. In the following, we assume the learner is allowed to return any hypothesis from some hypothesis class, $\mathcal{H}$ with

$$h : S \times \Theta \times \{0, \ldots, L-1\} \to A, \forall h \in \mathcal{H}.$$

In particular, just as in Yao et al. [2018], who consider a related algorithm called Direct Policy Transfer, we assume that DPL is agnostic to the reward sequence of the expert.

4.1 DPL Sample Complexity Upper Bound

We now derive a sample complexity upper bound for DPL. Our analysis is similar to that of the standard agnostic PAC learning upper bound, except that, in this setting, the batch of data are not i.i.d, but rather, come from a Markov chain. This, however, can be remedied, by simply replacing one of the key concentration inequalities in the standard setting (McDiarmid's) with an analogous concentration inequality which applies to Markov chains, shown in Paulin et al. [2015]. We define notions of classification error in this setting.

Definition 4.0.1 (nth Marginal Error)**.** *Let $n \in \mathbb{N}$ be at most $L - 1$. Define the nth marginal error of a hypothesis h to be*

$$\mathcal{L}_\pi^{marginal, n}(h) = P_\pi(h(s_n, n) \neq \pi(s_n, n)),$$

where the probability is being taken with respect to the randomness of drawing $\theta \sim \mathcal{P}_\Theta$, drawing $s_0 \sim \mathcal{P}_{s_0}$, and the randomness from following the deterministic policy π for n steps under the (potentially stochastic) transitions indexed by θ.

We now arrive at a definition of true error in this setting.

Definition 4.0.2. *Define the true error of the hypothesis h to be*

$$\mathcal{L}_{\mathcal{D},\pi}^L(h) = \frac{1}{L} \sum_{l=0}^{L-1} \mathcal{L}_\pi^{marginal,l}(h).$$

Finally, the definition of empirical error is same as in the standard setting:

Definition 4.0.3. *Letting S consist of mL tuples, as outlined in the data generation procedure above, we define the empirical error to be*

$$\mathcal{L}_S^{m,L}(h) = \frac{|\{(i,j) \in \{1,\ldots,m\} \times \{0,\ldots,L-1\} : h(s_j^{\theta_i}, \theta_i, j) \neq \pi(s_j^{\theta_i}, \theta_i, j)\}|}{mL}.$$

With these definitions in place, we are able to derive a sample complexity of classification upper bound for DPL using any hypothesis class $\mathcal{H}$. Our analysis follows the standard agnostic PAC analysis (see, e.g., Shalev-Shwartz and Ben-David [2014]) but applied to the non-i.i.d. setting, in which we use a concentration result of Paulin et al. [2015]. We state the agnostic sample complexity of classification upper bound in terms of the Natarajan dimension of the hypothesis class Ndim($\mathcal{H}$) (for proof of the theorem and definition of Natarajan dimension, see Appendix Section A.1)

Theorem 4.1. *Let the concept class $\mathcal{H}$ have Natarajan dimension d. There exists a learning algorithm $\mathcal{A}$ such that for any distribution over the data, there exists m with*

$$m = O\left(\frac{d}{\epsilon^2}\left(\log\left(\frac{d}{\epsilon L}\right) + \frac{d}{L}(\log(L) + \log(|A|)) + L^2\log(1/\delta)\right)\right),$$

such that if $\mathcal{A}$ receives at least m L-long trajectories in the batch, then $\mathcal{A}$ returns a hypothesis in $\mathcal{H}$ which has classification error at most ϵ greater than the true error minimizer of $\mathcal{H}$ with probability at least $1 - \delta$.

Before considering how the classification error of π affects its error in value, we make two standard assumptions regarding the hypothesis class $\mathcal{H}$ as well as the reward structure:

Assumption 4.1 (Realizability)**.** *The hypothesis class $\mathcal{H}$ contains the expert's labeling, π.*

Assumption 4.2. *All rewards are in the range $[0,1]$.*

We use an extension of a result of Ross and Bagnell [2010] who bound the error of a time-independent policy in terms of its classification error for a single sequential decision-making task. By carefully comparing the marginal distribution of states under $\hat{\pi}$ and π at each time step, they are able to bound the difference in values by $L^2\epsilon$, where ϵ is the classification error of $\hat{\pi}$ with respect to π. The proof in our multi-task setting with time-dependent policies follows theirs (for details, see Ross and Bagnell [2010]). With this result as well as our above sample complexity of classification bound, we now show that DPL requires only a $poly(\text{Ndim}(\mathcal{H}), \frac{1}{\epsilon}, \log(\frac{1}{\delta}), L, \log(|A|))$ number of samples.

Theorem 4.2. *Let the concept class $\mathcal{H}$ have Natarajan dimension d. There exists a learning algorithm $\mathcal{A}$ such that for any distribution over the data, there exists m that is*

$$O\left(\frac{L^4 d}{\epsilon^2}\left(\log\left(\frac{Ld}{\epsilon}\right) + \frac{d}{L}(\log(L) + \log(|A|)) + L^2\log(1/\delta)\right)\right)$$

receives at least m L-long trajectories in the batch, then $\mathcal{A}$ returns a hypothesis in $\mathcal{H}$ which has error (in terms of undiscounted value) at most $\epsilon + \alpha$ with probability at least $1 - \delta$.

4.2 DPL Sample Complexity Lower Bound

We derive a DPL lower bound by constructing a family of CMDPs for which the problem reduces to a standard PAC learning problem that must be learned to error ϵ/L with confidence δ. To do so, we

essentially put all the decision making power on the action taken at the first state. That is, let $\mathcal{P}_{s_0}$ be fully concentrated on the state s_0, and let all transitions be deterministic so that there are $|A|$ distinct potential next states. For each value of θ, call the state s, that satisfies $T(s|s_0, \pi(s_0, \theta), \theta) = 1$ good, and call the other $|A| - 1$ states bad. The good state satisfies the condition that it and all subsequent states give reward 1 independent of the actions taken in them, while any bad state and all subsequent states give reward 0, again, independently from actions. The learning problem is thus reduced to learning the first action to error at most ϵ/L with probability at least $1 - \delta$. Choosing $\mathcal{H}$ so that there is a shattered set of the form $\{(\theta_1, s_0), \ldots, (\theta_m, s_0)\}$, the following lower bound follows immediately from the standard i.i.d PAC lower bound.

Theorem 4.3. *There exist a family of CMDPs, a hypothesis class, $\mathcal{H}$ of Natarajan dimension d, and value of m with*

$$m = \Omega\left(\frac{L(d + \log(1/\delta))}{\epsilon}\right),$$

such that any learning algorithm given at most m L-long trajectories returns a policy whose error (in terms of value) is at most ϵ with probability at most $1 - \delta$.

Our results thus give a $\tilde{O}\left(\frac{L^2 d}{\epsilon} + \frac{L^5}{\epsilon}\right)$ separation between upper and lower bounds, in particular, highlighting how poorly the bounds scale with L, but how slowly they separate with δ and A; we discuss this further in Section 6.

5 Model-based Approach

In contrast to DPL, the model-based approach does not attempt to directly learn the expert policy, but rather attempts to learn the transition function parametrized by θ, and then plan according to the transition and reward function. One might believe that learning models would be more general than trying to directly learn policies, because one can use them to explore counterfactuals. Indeed, model-based learning is often the go-to approach in low data regimes (see, e.g. Rasmussen [2003] versus Deisenroth et al. [2013], Kamthe and Deisenroth [2017], Kocijan et al. [2004], Ko et al. [2007]). However, we first show that the paradigm of learning from expert actions is, in general settings, impossible—even if direct policy learning is possible in these settings. Thus, we relax the data generation process under which the model-based method is expected to learn and introduce a process which allows for greater coverage of $S \times A$.

Throughout our analysis, for simplicity, we will assume the model-based approach has access to an oracle called `PLAN`, which, upon receiving a transition function, T for a single MDP, and reward function, R (which, as mentioned in Section 3, the learner has access to), returns an optimal deterministic (possible time-dependent) policy under T and R in that MDP. Hence, if the model-based algorithm returns a hypothesis transition function h, its corresponding value is simply $v^L_{\text{PLAN}(h(\theta),R)}$, where $h(\theta)$ denotes the transition function restricted to θ. Note that evaluating a model-based algorithm in this sense implicitly defines the policy, $\hat{\pi}_h$, corresponding to the hypothesis transition function h, returned as $\hat{\pi}_h(s, \theta, l) = (\text{PLAN}(h(\theta), R))(s, l)$.

5.1 Impossibility of Model-Based Learning via Expert Actions

The primary issue a model-based approach having access only to an expert's trajectory is the lack of coverage of state-action space. In particular, it may be the case that a large portion of the learner's hypothesis class always agree on the subset of state-action space traversed by the expert, making it hard—and, as we show below, in some cases impossible—to output a low-error hypothesis with high confidence. We now give a construction formalizing the above which demonstrates an infinite sample complexity for the model-based approach. Under our construction, the learner is given a finite model hypothesis class from which to return its hypothesis transition function. This eliminates the possibility of the infinite sample complexity being explained away by the hypothesis class having infinite statistical complexity, and instead highlights the fundamental difficulty model-based learning faces when attempting to learn from a single expert's actions.

Theorem 5.1. *There exist classes of CMDPs which, if the learner must learn from expert trajectories, require infinite sample complexity even with a finite hypothesis class containing the true model.*

Proof. Our construction is inspired by that of Chen and Jiang [2019] who use a similar setup to derive an information-theoretic lower bound in the single MDP batch setting. First, let $\Theta = [|A|^{L-1} - 1]$ and $\mathcal{P}_\Theta = \text{Unif}(\Theta)$. Now, let a CMDP be represented as a rooted complete tree with branching factor $|A|$ and depth L. Each non-leaf node represents a distinct state giving 0 reward, while each leaf node represents two possible states; we furthermore restrict $\mathcal{P}_{s_0}$ to concentrate all its mass on some single initial state. Taking action $a_{|A|}$ (the rightmost action) from the rightmost node in the penultimate layer transitions to a state with reward 0 with probability $\frac{1-\epsilon}{2}$ and reward 1 with probability $\frac{1+\epsilon}{2}$. Finally, let all other edges leaving the penultimate layer, except the θth edge from the left, transition to a state with reward 0 with probability $1/2$ and a state with reward 1 with probability $1/2$, while the θth edge transitions to a state with reward 0 with probability $\frac{1}{2} - \frac{4}{3}\epsilon$ and reward 1 with probability $\frac{1}{2} + \frac{4}{3}\epsilon$ (so our result will hold for $\epsilon < \frac{3}{8}$). We will call this edge, leading to the highest rewarding leaf, in expectation, the special edge, and we will let the expert's policy be that which continually chooses the rightmost action—notice that the expert is $\frac{5}{6}\epsilon$-optimal. Call a CMDP of the above form a tree CMDP: Notice that a tree CMDP is fully characterized by the map taking θ to the special edge—in this case, the identity map. It is thus clear that any permutation of Θ will yield a different special edge map (i.e. the permutation σ says that θ corresponds to the MDP with special edge given by the $\sigma(\theta)$th edge of the penultimate layer). So, the learner is given as its hypothesis model class two hypotheses: the identity permutation and some derangement σ. Noticing that any algorithm must determine precisely the target CMDP in order to return an ϵ-optimal hypothesis (since otherwise, PLAN gives a policy with value $1/2$ because if $h \in \mathcal{H}$ is not equal to the target concept, it will disagree with the target concept on every value of θ, since we chose σ to be a derangement) and that the expert's trajectory gives no information about which member of $\mathcal{H}$ is the target CMDP, any learning algorithm returns an ϵ-optimal hypothesis with probability at most $\frac{1}{2}$, and, in particular, cannot guarantee ϵ-optimality with δ confidence for any $\delta, \epsilon < \frac{1}{2}$, thus indicating the impossibility of learning this model class. $\square$

The above construction illustrates that learning from a single expert's trajectory can be impossible for any model-based learning algorithm. In particular, while DPL, with sufficiently many samples, is able to return a policy which is nearly $\frac{5}{6}\epsilon$-optimal, any model-based learning algorithm will, with probability at least $\frac{1}{2}$ (regardless of the size of m), return a policy which has suboptimality larger than ϵ. This highlights the key difficulty with model-based learning in a CMDP from an expert's trajectories: while DPL can always achieve suboptimality arbitrarily close to that of the expert (as long as the Natarajan dimension of its hypothesis class is finite), the same cannot be said for model-based learning even when its hypothesis class is finite.

5.2 Hardness of Model-based Learning under strictly-positive visitation distributions

The impossibility result above motivates the use of the following more standard framework under which we expect batch model-based approaches to learn:

Definition 5.1.1 (Model-based Learning Data Generation Process). *Let μ be some distribution over $S \times A$ which assigns non-zero mass/density to every $(s, a) \in S \times A$. For each value of $\theta \sim \mathcal{P}_\Theta$, the model-based approach draws L pairs $(s, a) \overset{i.i.d}{\sim} \mu$ and, for each pair, draws $s' \sim T(\cdot|s, a, \theta)$ and $r = R(s, a)$. The model-based approach then has access to each of these one-step trajectories, labelled by θ. We also make the assumption that every element of $S \times A$ is reachable in at most L steps.*

Under the framework defined in Definition 5.1.1, we show that there are still classes for which model-based learning is hard. In fact, the construction of Chen and Jiang [2019] for single MDPs gives a lower bound in our setting as well since we can simply consider the CMDP which concentrates all its mass on a single value of θ, thus reducing to the single MDP case. Their construction is very similar to the above tree construction: States are the nodes of a complete tree with branching factor $|A|$, and all leaf nodes give $\text{Bern}(1/2)$ rewards while the leaf node corresponding to the special edge gives $\text{Bern}(1/2 + 3\epsilon/2)$ reward. In our setting, the construction of Krishnamurthy et al. [2016] yields a sample complexity lower bound of $\Omega\left(\frac{|A|^L}{L\epsilon^2}\right)$ when active exploration is allowed. This sample complexity is super-polynomial in statistical complexity of the hypothesis class, which has cardinality $|A|^L$. In Appendix Section A.2, we give a construction of a family of MDPs, motivated by contextual bandits and the constructions of Krishnamurthy et al. [2016] and Auer et al. [2003], which

yields a sample complexity lower bound of $\Omega\left(\frac{|A|}{\epsilon^2}\right)$, allowing exploration for a hypothesis class of cardinality $|A|$. Our bound, on the surface, is asymptotically lower than that of Krishnamurthy et al. [2016]; however, since the hypothesis class as well as the size of each CMDP in the class have size only $O(|A|)$, our bound exhibits a stronger dependence on the size of the model class. That is, while the original bound of Krishnamurthy et al. [2016] precludes $poly(\log(|\mathcal{H}|), \frac{1}{\delta}, \frac{1}{\epsilon})$ sample complexity, our second bound gives a $\Omega\left(\frac{|\mathcal{H}|}{\epsilon^2}\right)$ dependence on the size of the hypothesis class, rather than the immediate $\Omega\left(\frac{|\mathcal{H}|}{L\epsilon^2}\right)$ of Krishnamurthy et al. [2016].

The above constructions show how the sample complexity for model-based learning can scale poorly with both horizon as well as the size and statistical complexity of the hypothesis class, but fail to show that model-based learning can scale poorly with $|\Theta|$ when it is finite. Modi et al. [2017], who consider online CMDPs, essentially suggest constructing hard CMDPs by making the MDP for each context, which is drawn uniformly from Θ, hard and disallowing any information corresponding to one context be useful to another. That is, having knowledge of the target transition function for all values in $\Theta \backslash \{\theta\}$ provides no information about the target transition for θ. While this technique gives a generic way to increase any hard MDP lower bound by a multiplicative factor of $O(|\Theta|)$ in expectation for CMDPs, it unfortunately makes the hypothesis class have cardinality exponential in $|\Theta|$, therefore explaining away the factor of $O(|\Theta|)$ by the complexity of the hypothesis class. We give a construction of a class of CMDPs in Appendix Section A.2, extending off of Krishnamurthy et al. [2016], which yields the same lower bound, but does so in a way that scales linearly with $|\Theta|$, while having a hypothesis class of cardinality $|\Theta|$.

We now state lower bounds both in the setting where active exploration is allowed and also when it is not. The latter bound is in terms of, C, the concentratability coefficient of μ, which measures how much μ covers reachable state-action pairs (for a definition, see Section 5.3.1).

Theorem 5.2. *There exist hard families of CMDPs which are subject to the following sample complexity lower bounds (all are asymptotically at most $\Omega\left(\frac{|A|^L}{L\epsilon^2}\right)$): $\Omega\left(\frac{|A|^L}{L\epsilon^2}\right)$, $\Omega\left(\frac{|\mathcal{H}|}{\epsilon^2}\right)$, and $\Omega\left(\frac{|\Theta|}{L\epsilon^2}\right)$ with $|\Theta| = |\mathcal{H}|$. Furthermore, when active exploration is not allowed, we have the following lower bound, in expectation, with respect to the randomness of drawing a leaf state from μ: $\Omega\left(\frac{C|A|^L}{L\epsilon^2}\right)$*

Thus, to summarize, we have shown not only that the sample complexity for model-based learning can grow poorly with $|\Theta|$ and L even when $\mathcal{H}$ is not too large, but also that the sample complexity depends on the concentratability coefficient at least linearly.

5.3   Model-based Learning Sample Complexity Upper Bound

We now derive an upper bound for model-based learning to contrast the above lower bound and, in particular, show that the dependence on C is in fact linear. To do so, we extend off of the work of Chen & Jiang (2019), who derive an upper bound for Fitted Q-Iteration (FQI). The FQI sample complexity upper bound then immediately yields a sample complexity upper bound for any model-based approach with a finite model class. In their approach, they consider an infinite-horizon discounted setting on a single MDP, whereas in our setting, we consider a finite-horizon undiscounted setting on CMDPs, in which time-dependent policies are allowed. First, we show how FQI can be applied in this setting, define a complexity measure which controls μ, the distribution from which state-action pairs are drawn, and then state a sample complexity result for FQI from which a sample complexity upper bound for model-based learning follows as an immediate upper bound.

5.3.1   FQI on Finite-Horizon Undiscounted CMDPs

We give a brief outline of finite-horizon FQI on a CMDP below; it is essentially the same as FQI on a single MDP except that Bellman backups are done with respect to the context, θ. We first define this back-up and give the algorithm below:

Definition 5.2.1. *Define the lth Bellman backup of $f : S \times \Theta \times \{0, \ldots, L\} \to \mathbb{R}$ with respect to θ to be*

$$(\mathcal{T}_l(\theta)f)(s, a) = R(s, a) + \mathbb{E}_{s' \sim T(\cdot|s,a,\theta)}V_f(s', \theta, l),$$

where $V_f(s, \theta, l) = \max_{a \in A} f(s, a, \theta, l)$.

We will also assume FQI has access to a family of time-indexed Q functions. That is we have a set $\mathcal{F}$ which contains Q-value functions of the form $Q : S \times A \times \Theta \times \{0, \ldots, L\}$, where $Q(s, a, \theta, l) = (\mathcal{T}_{l-1}(\theta)Q)(s, a)$ for $l \geq 1$, and $Q(s, a, \theta, 0) = 0$. FQI on CMDPs operates in essentially the same way as the finite horizon case except that all backups and value functions are additionally parametrized by θ (see Appendix Section A.3 for pseudocode).

We now give the definition of admissible distribution and the assumption of concentratability of the data disitribution μ which extends that of Chen & Jiang (2019).

Definition 5.2.2 (Admissible Distribution). *A conditional distribution ν over $S \times A$ given $\theta \in \Theta$ is said to be admissible if there exists $0 \leq l \leq L-1$ if there exists a possibly time-dependent stochastic policy π such that $(\nu(\theta))(s, a) = P[s_l = s, a_h = a | \theta, s_0 \sim \mathcal{P}_{s_0}, \pi]$*

Assumption 5.1 (Concentratability). *We assume that there exists some $C < \infty$ such that, for any admissible distribution ν,*

$$\frac{(\nu(\theta))(s, a)}{\mu(s, a)} \leq C, \forall (s, a, \theta) \in S \times A \times \Theta.$$

With these definitions, we now state a sample complexity upper bound for FQI, and then derive, as an immediate corollary, a sample complexity upper bound for model-based learning; again, we assume realizability for the hypothesis class $\mathcal{H}$ and thus of the class $\mathcal{F}$. For proofs, see Appendix A.3.

Theorem 5.3 (FQI Upper Bound). *There exists m with*

$$m = O\left(\frac{CL^6 \log(L|\mathcal{F}|/\delta)}{\epsilon^2}\right),$$

such that if FQI receives at least m samples of L one-step trajectories under μ, then it returns a policy with error at most ϵ with probability at least $1 - \delta$.

Corollary 5.3.1 (Model-based Upper Bound). *Given the finite hypothesis model class $\mathcal{H}$, there exists a model-based learning algorithm $\mathcal{A}$ and m with*

$$m = O\left(\frac{CL^6 \log(L|\mathcal{H}|/\delta)}{\epsilon^2}\right),$$

such that if $\mathcal{A}$ receives at least m samples of L one-step trajectories under μ, then it returns a policy with error at most ϵ with probability at least $1 - \delta$.

6 Discussion

In this paper we investigate the sample complexities of an imitation learning-base approach for learning of CMDPs, DPL, as well as a model-based approach. We find that the upper bounds for each approach are, respectively, $\tilde{O}\left(\frac{L^4 d}{\epsilon^2}\left(\frac{d}{L} + L^2\right)\right)$ and $\tilde{O}\left(\frac{CL^6}{\epsilon^2}\right)$. Our results indicate that DPL is, theoretically, more sound than model-based approaches in the sense that the latter scales with respect to the concentratability coefficient of the distribution μ. In particular, while both upper bounds scale polynomially in all the relevant parameters—and, importantly, in the complexity of hypothesis class—our upper bound for model-based learning scales with C. As our lower bound for model-based learning shows, this additional dependence is, in fact, *necessary*: Data distributions μ which concentrate low mass to regions of $S \times A$ which differentiate hypotheses are harder to learn under. This highlights the importance of the data generation process for model-based learning: When data is gotten from expert trajectories, model-based learning can be impossible even with finite hypothesis classes, but even when data is drawn i.i.d. from the distribution μ, model-based learning depends greatly on the coverage of reachable state-action pairs.

We believe the following are primary interests for future work: Deriving general model-based sample complexity upper bounds which do not grow, even logarithmically with $|\mathcal{H}|$, but rather grow with some other complexity measure of the hypothesis class which can be finite even for infinite $\mathcal{H}$ (e.g. perhaps with an extension of *witness rank* introduced in Sun et al. [2018]); investigating a tighter relationship between the upper and lower bounds for DPL, in particular, bounds whose degree of separation scales more slowly with L; and understanding the sample complexity of similar imitation learning and model-based algorithms in the unobserved parameter setting of HiP-MDPs.

7 Broader Impact

The primary real-world impact of this research—and, in fact, the application by which the authors were motivated—is to the world of healthcare. In particular, this work serves as a first-step to formally understanding the more complex framework of HiP-MDPs introduced in Doshi-Velez and Konidaris [2016] to encapsulate learning to generalize from expert actions when the parametrization of the transition of each MDP is unknown and the parameter corresponding to each trajectory in the batch is unknown. In particular, the motivation for HiP-MDPs is to learn how to generalize from a healthcare professional's actions when treating patients whose response to the same treatment strategy may differ. Thus, we see our work as a first step in formally understanding what types of algorithms may be most sample-efficient for learning to generalize well in such settings. That said, this work is theoretical in nature, and makes standard theory assumptions such as Markovianity under the definition of state s, existence of the true function with the hypothesis class, etc. Thus, while we provide theoretical foundations and insights for empirical observations, any application of our work to real settings should be mindful of the assumptions we make.

Acknowledgment

References

Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E. Schapire. The nonstochastic multiarmed bandit problem. *SIAM J. Comput.*, 32(1):48–77, January 2003. ISSN 0097-5397. doi: 10.1137/S0097539701398375. URL https://doi.org/10.1137/S0097539701398375.

Emma Brunskill and Lihong Li. Sample complexity of multi-task reinforcement learning. *arXiv preprint arXiv:1309.6821*, 2013.

Jinglin Chen and Nan Jiang. Information-theoretic considerations in batch reinforcement learning. *arXiv preprint arXiv:1905.00360*, 2019.

Marc Peter Deisenroth, Dieter Fox, and Carl Edward Rasmussen. Gaussian processes for data-efficient learning in robotics and control. *IEEE transactions on pattern analysis and machine intelligence*, 37(2):408–423, 2013.

Carlo D'Eramo, Davide Tateo, Andrea Bonarini, Marcello Restelli, and Jan Peters. Sharing knowledge in multi-task deep reinforcement learning. In *International Conference on Learning Representations*, 2019.

Finale Doshi-Velez and George Konidaris. Hidden parameter markov decision processes: A semi-parametric regression approach for discovering latent task parametrizations. In *IJCAI: proceedings of the conference*, volume 2016, page 1432. NIH Public Access, 2016.

Assaf Hallak, Dotan Di Castro, and Shie Mannor. Contextual markov decision processes, 2015.

David Isele, Mohammad Rostami, and Eric Eaton. Using task descriptions in lifelong machine learning for improved performance and zero-shot transfer. *CoRR*, abs/1710.03850, 2017. URL http://arxiv.org/abs/1710.03850.

Sham Kakade and John Langford. Approximately optimal approximate reinforcement learning. In *ICML*, volume 2, pages 267–274, 2002.

Sanket Kamthe and Marc Peter Deisenroth. Data-efficient reinforcement learning with probabilistic model predictive control. *arXiv preprint arXiv:1706.06491*, 2017.

Taylor W Killian, Samuel Daulton, George Konidaris, and Finale Doshi-Velez. Robust and efficient transfer learning with hidden parameter markov decision processes. In *Advances in neural information processing systems*, pages 6250–6261, 2017.

Jonathan Ko, Daniel J Klein, Dieter Fox, and Dirk Haehnel. Gaussian processes and reinforcement learning for identification and control of an autonomous blimp. In *Proceedings 2007 ieee international conference on robotics and automation*, pages 742–747. IEEE, 2007.

Juš Kocijan, Roderick Murray-Smith, Carl Edward Rasmussen, and Agathe Girard. Gaussian process model based predictive control. In *Proceedings of the 2004 American control conference*, volume 3, pages 2214–2219. IEEE, 2004.

Akshay Krishnamurthy, Alekh Agarwal, and John Langford. Pac reinforcement learning with rich observations. In *Advances in Neural Information Processing Systems*, pages 1840–1848, 2016.

Alessandro Lazaric and Marcello Restelli. Transfer from multiple mdps. *CoRR*, abs/1108.6211, 2011. URL http://arxiv.org/abs/1108.6211.

Aditya Modi and Ambuj Tewari. Contextual markov decision processes using generalized linear models. *arXiv preprint arXiv:1903.06187*, 2019.

Aditya Modi, Nan Jiang, Satinder Singh, and Ambuj Tewari. Markov decision processes with continuous side information. *arXiv preprint arXiv:1711.05726*, 2017.

Takayuki Osa, Joni Pajarinen, Gerhard Neumann, J Andrew Bagnell, Pieter Abbeel, and Jan Peters. An algorithmic perspective on imitation learning. *arXiv preprint arXiv:1811.06711*, 2018.

Daniel Paulin et al. Concentration inequalities for markov chains by marton couplings and spectral methods. *Electronic Journal of Probability*, 20, 2015.

Bilal Piot, Matthieu Geist, and Olivier Pietquin. Learning from demonstrations: Is it worth estimating a reward function? In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 17–32. Springer, 2013.

Carl Edward Rasmussen. Gaussian processes in machine learning. In *Summer School on Machine Learning*, pages 63–71. Springer, 2003.

Bernardino Romera-Paredes and Philip Torr. An embarrassingly simple approach to zero-shot learning. In *International Conference on Machine Learning*, pages 2152–2161, 2015.

Stéphane Ross and Drew Bagnell. Efficient reductions for imitation learning. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 661–668, 2010.

Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.

Archit Sharma, Shixiang Gu, Sergey Levine, Vikash Kumar, and Karol Hausman. Dynamics-aware unsupervised discovery of skills. *arXiv preprint arXiv:1907.01657*, 2019.

Sungryull Sohn, Junhyuk Oh, and Honglak Lee. Hierarchical reinforcement learning for zero-shot generalization with subtask dependencies. In *Advances in Neural Information Processing Systems*, pages 7156–7166, 2018.

Wen Sun, Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, and John Langford. Model-based rl in contextual decision processes: Pac bounds and exponential improvements over model-free approaches. *arXiv preprint arXiv:1811.08540*, 2018.

Matthew E Taylor and Peter Stone. Transfer learning for reinforcement learning domains: A survey. *Journal of Machine Learning Research*, 10(Jul):1633–1685, 2009.

Jiachen Yang, Brenden Petersen, Hongyuan Zha, and Daniel Faissol. Single episode policy transfer in reinforcement learning. *arXiv preprint arXiv:1910.07719*, 2019.

Jiayu Yao, Taylor Killian, George Konidaris, and Finale Doshi-Velez. Direct policy transfer via hidden parameter markov decision processes. In *LLARLA Workshop, FAIM*, volume 2018, 2018.

A Proofs

A.1 DPL Sample Complexity of Classification Upper Bound

Our proofs below follow the analysis of Shalev-Shwartz and Ben-David [2014] and use similar notation.

Lemma A.0.1. *Let S be a batch of trajectories. Then $\mathbb{E}_S[\mathcal{L}_S^{m,L}(h)] = \mathcal{L}_{\mathcal{D},\pi}^L(h)$.*

Proof. Note that

$$\mathcal{L}_S^{m,L}(h) = \frac{1}{mL} \sum_{i=1}^m \sum_{l=0}^{L-1} \mathbb{1}_{\{h(s_l^{\theta_i},\theta_i,l) \neq \pi(s_l^{\theta_i},\theta_i,l)\}},$$

and notice that

$$\mathbb{E}\left[\mathbb{1}_{\{h(s_l^{\theta_i},\theta_i,l) \neq \pi(s_l^{\theta_i},\theta_i,l)\}}\right] = \mathcal{L}_\pi^{marginal,l},$$

and so

$$\mathbb{E}\mathcal{L}_S^{m,L}(h) = \frac{1}{mL} \sum_{i=1}^m \sum_{l=0}^{L-1} \mathcal{L}_\pi^{marginal,l} = \frac{1}{L} \sum_{l=0}^{L-1} \mathcal{L}_\pi^{marginal,l} = \mathcal{L}_{\mathcal{D},\pi}^{m,L}(h),$$

as desired. $\square$

Definition A.0.1. *A training set, S, of m L-long trajectories is ϵ-representative if*

$$\sup_{h \in \mathcal{H}} |\mathcal{L}_{\mathcal{D},\pi}^L(h) - \mathcal{L}_S^{m,L}(h)| \leq \epsilon.$$

Now we proceed with a lemma regarding bounds on ϵ-representativeness; it is essentially the same as that in Shalev-Shwartz and Ben-David [2014] save for a few minor modifications.

Lemma A.0.2. *Let S be a training set consisting of m L-long trajectories. Then*

$$\mathbb{E}_S\left[\sup_{h \in \mathcal{H}} \left(\mathcal{L}_{\mathcal{D},\pi}^L(h) - \mathcal{L}_S^{m,L}(h)\right)\right] \leq 2\mathbb{E}_S R(\mathcal{F} \circ S),$$

where R denotes the Rademacher complexity and

$$\mathcal{F} \circ S = \left\{ \left(\mathbb{1}_{\{h(s_0^{\theta_1},\theta_1,0) \neq \pi(s_0^{\theta_1},\theta_1,0)\}}, \ldots, \mathbb{1}_{\{h(s_{L-1}^{\theta_m},\theta_m,L-1) \neq \pi(s_{L-1}^{\theta_m},\theta_m,L-1)\}}\right) : h \in \mathcal{H} \right\}.$$

Proof. Let S and S' be two datasets both sampled according to the procedure above and recall that $\mathcal{L}_{\mathcal{D},\pi}^L(h) = \mathbb{E}_{S'}\mathcal{L}_{S'}^{m,L}(h), \forall h \in \mathcal{H}$. Thus, noting that the supremum of the expectation is at most the expectation of the supremum, we have

$$\sup_{h \in \mathcal{H}} \left(\mathcal{L}_{\mathcal{D},\pi}^L(h) - \mathcal{L}_S^{m,L}(h)\right) = \sup_{h \in \mathcal{H}} \mathbb{E}_{S'}\left[\mathcal{L}_{S'}^{m,L}(h) - \mathcal{L}_S^{m,L}(h)\right]$$

$$\leq \mathbb{E}_{S'}\left[\sup_{h \in \mathcal{H}}(\mathcal{L}_{S'}^{m,L}(h) - \mathcal{L}_S^{m,L}(h))\right],$$

so taking expectations gives

$$\mathbb{E}_S\left[\sup_{h \in \mathcal{H}} \left(\mathcal{L}_{\mathcal{D},\pi}^L(h) - \mathcal{L}_S^{m,L}(h)\right)\right] \leq \mathbb{E}_{S,S'}\left[\sup_{h \in \mathcal{H}}(\mathcal{L}_{S'}^{m,L}(h) - \mathcal{L}_S^{m,L}(h))\right]$$

$$= \frac{1}{mL}\mathbb{E}_{S,S'}\left[\sup_{h \in \mathcal{H}} \sum_{i=1}^m \sum_{l=0}^{L-1}(\mathbb{1}_{\{h(s_l^{\theta_i},\theta_i,l) \neq \pi(s_l^{\theta_i},\theta_i,l)\}} - \mathbb{1}_{\{h(s_l^{\theta_i'},\theta_i',l) \neq \pi(s_l^{\theta_i'},\theta_i',l)\}})\right].$$

Now, notice that since θ_i and θ_i' are independent and identically distributed and $s_l^{\theta_i}$ and $s_l^{\theta_i'}$ are independent and identically distributed, we see that

$$\mathbb{E}\Bigg[\sup_{h \in \mathcal{H}} \Bigg((\mathbb{1}_{\{h(s_l^{\theta_i'},\theta_i',l) \neq \pi(s_l^{\theta_i'},\theta_i',l)\}} - \mathbb{1}_{\{h(s_l^{\theta_i},\theta_i,l) \neq \pi(s_l^{\theta_i},\theta_i,l)\}})$$

$$+ \sum_{k \neq i} \sum_{j \neq l}(\mathbb{1}_{\{h(s_j^{\theta_k'},\theta_k',j) \neq \pi(s_j^{\theta_k'},\theta_k',j)\}} - \mathbb{1}_{\{h(s_j^{\theta_k},\theta_k,j) \neq \pi(s_j^{\theta_k},\theta_k,j)\}})\Bigg)\Bigg]$$

$$= \mathbb{E}\Bigg[\sup_{h \in \mathcal{H}} \Bigg((\mathbb{1}_{\{h(s_l^{\theta_i},\theta_i,l) \neq \pi(s_l^{\theta_i},\theta_i,l)\}} - \mathbb{1}_{\{h(s_l^{\theta_i'},\theta_i',l) \neq \pi(s_l^{\theta_i'},\theta_i',l)\}})$$

$$+ \sum_{k \neq i} \sum_{j \neq l}(\mathbb{1}_{\{h(s_j^{\theta_k'},\theta_k',j) \neq \pi(s_j^{\theta_k'},\theta_k',j)\}} - \mathbb{1}_{\{h(s_j^{\theta_k},\theta_k,j) \neq \pi(s_j^{\theta_k},\theta_k,j)\}})\Bigg)\Bigg],$$

so that if $\sigma_{i,l}$ is a random sign, we see that, by the law of total expectation,

$$\mathbb{E}_{S,S',\sigma_{i,l}}\Big[\sup_{h\in\mathcal{H}}\Big(\sigma_{i,l}(\mathbb{1}_{\{h(s_l^{\theta_i'},\theta_i',l)\neq\pi(s_l^{\theta_i'},\theta_i',l)\}}-\mathbb{1}_{\{h(s_l^{\theta_i},\theta_i,l)\neq\pi(s_l^{\theta_i},\theta_i,l)\}})$$

$$+\sum_{k\neq i}\sum_{j\neq l}(\mathbb{1}_{\{h(s_j^{\theta_k'},\theta_k',j)\neq\pi(s_j^{\theta_k'},\theta_k',j)\}}-\mathbb{1}_{\{h(s_j^{\theta_k},\theta_k,j)\neq\pi(s_j^{\theta_k},\theta_k,j)\}}))\Big]$$

$$=\mathbb{E}_{S,S'}\Big[\sup_{h\in\mathcal{H}}\Big((\mathbb{1}_{\{h(s_l^{\theta_i'},\theta_i',l)\neq\pi(s_l^{\theta_i'},\theta_i',l)\}}-\mathbb{1}_{\{h(s_l^{\theta_i},\theta_i,l)\neq\pi(s_l^{\theta_i},\theta_i,l)\}})$$

$$+\sum_{k\neq i}\sum_{j\neq l}(\mathbb{1}_{\{h(s_j^{\theta_k'},\theta_k',j)\neq\pi(s_j^{\theta_k'},\theta_k',j)\}}-\mathbb{1}_{\{h(s_j^{\theta_k},\theta_k,j)\neq\pi(s_j^{\theta_k},\theta_k,j)\}}))\Big)\Big],$$

which, repeating for all i,l, indicates that

$$\mathbb{E}_{S,S'}\Big[\sup_{h\in\mathcal{H}}\Big(\sum_{i=1}^{m}\sum_{l=0}^{L-1}(\mathbb{1}_{\{h(s_l^{\theta_k'},\theta_k',l)\neq\pi(s_l^{\theta_k'},\theta_k',l)\}}-\mathbb{1}_{\{h(s_l^{\theta_k},\theta_k,l)\neq\pi(s_l^{\theta_k},\theta_k,l)\}}))\Big)\Big]$$

$$=\mathbb{E}_{S,S',\sigma}\Big[\sup_{h\in\mathcal{H}}\Big(\sum_{i=1}^{m}\sum_{l=0}^{L-1}\sigma_{i,l}(\mathbb{1}_{\{h(s_l^{\theta_k'},\theta_k',l)\neq\pi(s_l^{\theta_k'},\theta_k',l)\}}-\mathbb{1}_{\{h(s_l^{\theta_k},\theta_k,l)\neq\pi(s_l^{\theta_k},\theta_k,l)\}}))\Big)\Big],$$

where σ denotes the $m\times L$ matrix of iid random signs. Noting that

$$\sup_{h\in\mathcal{H}}\Big(\sum_{i=1}^{m}\sum_{l=0}^{L-1}\sigma_{i,l}(\mathbb{1}_{\{h(s_l^{\theta_k'},\theta_k',l)\neq\pi(s_l^{\theta_k'},\theta_k',l)\}}-\mathbb{1}_{\{h(s_l^{\theta_k},\theta_k,l)\neq\pi(s_l^{\theta_k},\theta_k,l)\}})\Big)$$

$$\leq\sup_{h\in\mathcal{H}}\sum_{i=1}^{m}\sum_{l=0}^{L-1}\sigma_{i,l}\mathbb{1}_{\{h(s_l^{\theta_k'},\theta_k',l)\neq\pi(s_l^{\theta_k'},\theta_k',l)\}}+\sup_{h\in\mathcal{H}}\sum_{i=1}^{m}\sum_{l=0}^{L-1}-\sigma_{i,j}\mathbb{1}_{\{h(s_l^{\theta_k},\theta_k,l)\neq\pi(s_l^{\theta_k},\theta_k,l)\}},$$

we have that, since $\sigma\sim-\sigma$,

$$\mathbb{E}_{S,S',\sigma}\Big[\sup_{h\in\mathcal{H}}\Big(\sum_{i=1}^{m}\sum_{l=0}^{L-1}\sigma_{i,l}(\mathbb{1}_{\{h(s_l^{\theta_k'},\theta_k',l)\neq\pi(s_l^{\theta_k'},\theta_k',l)\}}-\mathbb{1}_{\{h(s_l^{\theta_k},\theta_k,l)\neq\pi(s_l^{\theta_k},\theta_k,l)\}})\Big)\Big]$$

$$\leq\mathbb{E}_{S,S',\sigma}\Big[\sup_{h\in\mathcal{H}}\sum_{i=1}^{m}\sum_{l=0}^{L-1}\sigma_{i,j}\mathbb{1}_{\{h(s_l^{\theta_k'},\theta_k',l)\neq\pi(s_l^{\theta_k'},\theta_k',l)\}}+\sup_{h\in\mathcal{H}}\sum_{i=1}^{m}\sum_{l=0}^{L-1}\sigma_{i,l}\mathbb{1}_{\{h(s_l^{\theta_k},\theta_k,l)\neq\pi(s_l^{\theta_k},\theta_k,l)\}}\Big]$$

$$=2mL\mathbb{E}_{S}[R(\mathcal{F}\circ S)],$$

implying that

$$\mathbb{E}_{S}\Big[\sup_{h\in\mathcal{H}}\Big(\mathcal{L}_{\mathcal{D},\pi}^{m,L}(h)-\mathcal{L}_{S}^{m,L}(h)\Big)\Big]\leq 2\mathbb{E}_{S}R(\mathcal{F}\circ S).$$

$\square$

Now, the classical approach would now rely on McDiarmid's concentration inequality. However, we are unable to do so due to the dependent structure on our sample S. Instead we use an analogous concentration inequality suited for Markov processes. To handle this setting, we define the notion of mixing time in the case of time-inhomogeneous Markov chains and then state the key lemma from Paulin et al. [2015].

Definition A.0.2. *Let $X_1,\ldots,X_N$ be a Markov chain on Polish state space $\Omega_1\times\cdots\times\Omega_N$. Let $\mathcal{S}(X_{i+t}|X_i=x)$ be the conditional distribution of X_{i+t} given $X_i=x$. Define*

$$\overline{d}(t)=\max_{1\leq i\leq N-t}\sup_{x,y\in\Omega_i}d_{TV}(\mathcal{S}(X_{i+t}|X_i=x),\mathcal{S}(X_{i+t}|X_i=y))\ \text{and}\ \tau(\epsilon)=\min\{t\in\mathbb{N}:\overline{d}(t)\leq\epsilon\}.$$

Furthermore, define

$$\tau_{\min}=\inf_{0\leq\epsilon<1}\tau(\epsilon)\cdot\Big(\frac{2-\epsilon}{1-\epsilon}\Big)^2.$$

Then, from Paulin et al. [2015], we use the following lemma:

Lemma A.0.3. *Let $X_1,\ldots,X_N$ be a Markov chain on Polish state space $\Lambda=\Lambda_1\times\cdots\times\Lambda_N$. Suppose that $f:\Lambda\to\mathbb{R}$ satisfies, for all $x,y\in\Lambda$ that $f(x)-f(y)\leq\sum_{i=1}^{N}c_i\mathbb{1}[x_i\neq y_i]$, for some $c_i\in\mathbb{R}_+$. Then, for any $t\geq 0$,*

$$P\left(|f(x)-\mathbb{E}f(X)|\geq t\right)\leq 2\exp\left(\frac{-2t^2}{\tau_{\min}\sum_{i=1}^{N}c_i^2}\right).$$

12

We rewrite the following immediate consequence of the claim of Paulin et al. [2015], stated above, in the following form:

Corollary A.0.1. *Let $f : \mathcal{X}^n \to \mathbb{R}$ satisfy bounded differences. That is, for all i and $\forall x_1, \ldots, x_n, x_i' \in \mathcal{X}$,*

$$|f(x_1, \ldots, x_n) - f(x_1, \ldots, x_{i-1}, x_i', x_{i+1}, \ldots, x_n)| \leq c.$$

Then with probability at least $1 - \delta$,

$$|f(s_1, \ldots, s_n) - \mathbb{E}[f(s_1, \ldots, s_n)]| \leq c\sqrt{\frac{n\tau_{\min}\log(2/\delta)}{2}},$$

where again the states are sampled according to the Markov process induced by π.

Finally, we arrive at the main lemma, which we will use to prove the main result: a sample complexity bound in this special setting.

Lemma A.0.4. *Let S, containing mL data points, be sampled according to the procedure described above. Then with probability at least $1 - \delta, \forall h \in \mathcal{H}$,*

$$\mathcal{L}_{\mathcal{D},\pi}^L(h) - \mathcal{L}_S^{m,L}(h) \leq 2\mathbb{E}_{S'} R(\mathcal{F} \circ S') + \sqrt{\frac{(3L+1)\tau_{\min}\log(2/\delta)}{2m}}.$$

Proof. Write S as $S = (\theta_1, 0, s_0^{\theta_1}, a_0^{\theta_1}, \ldots, \theta_m, \ldots, L-1, s_{L-1}^{\theta_m}, a_{L-1}^{\theta_m})$, define

$$\zeta(\theta_1, 0, s_0^{\theta_1}, a_0^{\theta_1}, \ldots, \theta_m, \ldots, L-1, s_{L-1}^{\theta_m}, a_{L-1}^{\theta_m}) = \sup_{h \in \mathcal{H}}(\mathcal{L}_{\mathcal{D},\pi}^L(h) - \mathcal{L}_S^{m,L}(h)).$$

We verify that ζ satisfies the bounded differences condition above. Letting S' be the vector equal to S except in one component (which may be a component containing either a θ or a s^θ or a time index or a^θ), we have that

$$|\zeta(S) - \zeta(S')| = |\sup_{h \in \mathcal{H}}(\mathcal{L}_{\mathcal{D},\pi}^L(h) - \mathcal{L}_S^{m,L}(h)) - \sup_{h \in \mathcal{H}}(\mathcal{L}_{\mathcal{D},\pi}^L(h) - \mathcal{L}_{S'}^{m,L}(h))|$$

$$\leq \sup_{h \in \mathcal{H}}|\mathcal{L}_{S'}^{m,L}(h) - \mathcal{L}_S^{m,L}(h)| \leq 1/m,$$

since, the worst case scenario is when the entry which disagrees is a θ entry, which will result in, at worst, incorrect classification for each state in the trajectory associated with that θ resulting in L mistakes. Therefore, by the above corollary, we get that, with probability at least $1 - \delta$,

$$\sup_{h \in \mathcal{H}}(\mathcal{L}_{\mathcal{D},\pi}^{m,L}(h) - \mathcal{L}_S^{m,L}(h)) \leq \mathbb{E} \sup_{h \in \mathcal{H}}(\mathcal{L}_{\mathcal{D},\pi}^{m,L}(h) - \mathcal{L}_S^{m,L}(h)) + \frac{1}{m}\sqrt{\frac{(m + 3mL)\tau_{\min}\log(2/\delta)}{2}},$$

which, by Lemma A.0.2 is at most, with probability at least $1 - \delta$,

$$2\mathbb{E}_{S'} R(\mathcal{F} \circ S') + \sqrt{\frac{(3L+1)\tau_{\min}\log(2/\delta)}{2m}},$$

implying the desired result. $\square$

Now, we will upper bound $R(\mathcal{F} \circ S)$ for any dataset S. First, we recall the definition of Natarajan dimension, Natarajan's lemma, and recall that our hypothesis is a classifier on $|A|$ different classes:

Definition A.0.3 (Natarajan Dimension for k-class classification). *A subset $S \subset \mathcal{X}$ is said to be shattered by $\mathcal{H}$ if there exist functions $f_0, f_1 : C \to [k]$ such that for all $B \subset S$, there exists $h \in \mathcal{H}$ such that h agrees with f_0 on B and agrees with f_1 on $S \backslash B$.*

Lemma A.0.5 (Natarajan's). *Let $\mathcal{H}$ be a hypothesis class from some finite set $\mathcal{X} \to [k]$. Then*

$$|\mathcal{H}| \leq |\mathcal{X}|^{\text{Ndim}(\mathcal{H})} \cdot k^{2\text{Ndim}(\mathcal{H})}.$$

From this lemma, we see that

$$\left|\left\{\left(\mathbb{1}_{\{h(s_0^{\theta_1},\theta_1,0)\neq\pi(s_0^{\theta_1},\theta_1,0)\}}, \ldots, \mathbb{1}_{\{h(s_{L-1}^{\theta_m},\theta_m,L-1)\neq\pi(s_{L-1}^{\theta_m},\theta_m,L-1)\}}\right) : h \in \mathcal{H}\right\}\right| \leq (mL)^{\text{Ndim}(\mathcal{H})} \cdot |A|^{2\text{Ndim}(\mathcal{H})}.$$

Now, we state Massart's Lemma:

Lemma A.0.6 (Massart). *Let $B = \{b_1, \ldots, b_N\}$ denote a finite set of vectors in $\mathbb{R}^m$. Define $\overline{b} = \frac{1}{N}\sum_{i=1}^N b_i$. Then*

$$R(B) \leq \max_{b \in B} ||b - \overline{b}|| \frac{\sqrt{2\log(N)}}{m}.$$

13

Thus, the above two lemmas imply that, letting $d = \text{Ndim}(\mathcal{H})$,

$$R(\mathcal{F} \circ S') \leq \sqrt{\frac{2d \left(\log(mL) + 2\log(|A|)\right)}{mL}}.$$

Thus, by a union-bound we have that with probability at least $1 - \delta$, for every $h \in \mathcal{H}$, we have

$$|\mathcal{L}_{\mathcal{D}, \pi}^{m, L}(h) - \mathcal{L}_S^{m, L}(h)| \leq 2\sqrt{\frac{2d \left(\log(mL) + 2\log(|A|)\right)}{mL}} + \sqrt{\frac{(3L+1)\tau_{\min}\log(4/\delta)}{2m}}.$$

So, in order to ensure that this is at most ϵ, we require that

$$2\sqrt{\frac{2d \left(\log(mL) + 2\log(|A|)\right)}{mL}} + \sqrt{\frac{(3L+1)\tau_{\min}\log(4/\delta)}{2m}} \leq \epsilon$$

$$\impliedby 2\sqrt{\frac{8d \left(\log(mL) + 2\log(|A|)\right)}{mL} + \frac{(3L+1)\tau_{\min}\log(4/\delta)}{2m}} \leq \epsilon$$

$$\impliedby 2\sqrt{\frac{8d \left(\log(mL) + 2\log(|A|)\right)/l + (3L+3)\tau_{\min}\log(4/\delta)}{m}} \leq \epsilon$$

$$\impliedby m \geq \frac{4}{\epsilon^2}\left(\frac{8d}{L}\left(\log(m) + \log(L) + 2\log(|A|)\right) + (3L+3)\tau_{\min}\log(4/\delta)\right),$$

which, by the below lemma (whose proof can be found in Shalev-Shwartz and Ben-David [2014]), is implied by

$$m \geq \frac{128d}{\epsilon^2 L}\log\left(\frac{64d}{\epsilon^2 L}\right) + \frac{8}{\epsilon^2}\left(\frac{8d}{L}(\log(L) + 2\log(|A|)) + (3L+3)\tau_{\min}\log(4/\delta)\right).$$

Lemma A.0.7. *Let $a \geq 1, b > 0$. Then $x \geq 4a\log(2a) + 2b \implies x \geq a\log(x) + b$.*

Replacing the ϵ above with $\epsilon/2$ gives the desired sample complexity upper bound:

Lemma A.0.8. *Let $\mathcal{A}$ denote an ERM learning algorithm. Then for any distribution over the data, if $\mathcal{A}$ receives*

$$m = O\left(\frac{d}{\epsilon^2}\left(\log\left(\frac{d}{\epsilon L}\right) + \frac{d}{L}(\log(L) + \log(|A|)) + L\tau_{\min}\log(1/\delta)\right)\right),$$

then $\mathcal{A}$ returns a hypothesis in $\mathcal{H}$ which has error at most ϵ greater than the true error minimizer in $\mathcal{H}$ with probability at least $1 - \delta$.

Finally, we return to the definition of $\tau_{\min}$:

Definition A.0.4. *Let $X_1, \ldots, X_N$ be a Markov chain on Polish state space $\Omega_1 \times \cdots \times \Omega_N$. Let $\mathcal{L}(X_{i+t}|X_i = x)$ be the conditional distribution of X_{i+t} given $X_i = x$. Define*

$$\overline{d}(t) = \max_{1 \leq i \leq N-t} \sup_{x,y \in \Omega_i} d_{TV}(\mathcal{L}(X_{i+t}|X_i = x), \mathcal{L}(X_{i+t}|X_i = y)) \text{ and } \tau(\epsilon) = \min\{t \in \mathbb{N} : \overline{d}(t) \leq \epsilon\}.$$

Furthermore, define

$$\tau_{\min} = \inf_{0 \leq \epsilon < 1} \tau(\epsilon) \cdot \left(\frac{2-\epsilon}{1-\epsilon}\right)^2.$$

Notice that, in our case, looking at the proof of Lemma A.0.4, $t(\epsilon) \leq L, \forall 0 \leq \epsilon \leq 2L$ (since the entire process 'restarts' with a new draw of θ), and so we have that $\tau_{\min} \leq 8L$, thus giving us the desired theorem on sample complexity:

Theorem A.1. *Let the concept class $\mathcal{H}$ have Natarajan dimension d. There exists a learning algorithm $\mathcal{A}$ such that for any distribution over the data, there exists m with*

$$m = O\left(\frac{d}{\epsilon^2}\left(\log\left(\frac{d}{\epsilon L}\right) + \frac{d}{L}(\log(L) + \log(|A|)) + L^2\log(1/\delta)\right)\right),$$

such that if $\mathcal{A}$ receives at least m L-long trajectories in the batch, then $\mathcal{A}$ returns a hypothesis in $\mathcal{H}$ which has classification error at most ϵ greater than the true error minimizer of $\mathcal{H}$ with probability at least $1 - \delta$.

A.2 Hard Family of CMDPs

A.2.1 An $\Omega(\frac{|\mathcal{H}|}{\epsilon^2})$ Lower Bound

We consider the following K-armed contextual bandit problem. The hypothesis class will be a set of functions $f : \{0, \ldots, L-1\} \to [K]$. The ith context vectors will simply be i for $i = 0, \ldots, L-1$; that is, the only information the context gives is the round number. $f^*(i)$ denotes the optimal arm for round i. For the first

round, the optimal arm gives $\text{Bern}(\frac{1}{2} + \frac{3}{2}\epsilon)$ reward and all other arms give $\text{Bern}(1/2)$ rewards. For all other contexts, all arms give $\text{Bern}(1/2)$ rewards. With this structure, we further restrict the hypothesis class so that for any $f \neq f' \in \mathcal{H}$, $f(i) = f'(i), \forall \theta \in \Theta, i = 1, \ldots, L-1$; that is, all hypotheses agree on all but the very first round. On the first round, the set $\{f(0) : f \in \mathcal{H}\} = A$. Therefore, we simply take $|\mathcal{H}| = |A|$. Notice that for each round, the final $L-1$ values seen provide no information. Thus, only the first context is important, and the problem reduces to the standard K-arm bandit problem of Krishnamurthy et al. [2016], and thus we get the lower bound $\Omega\left(\frac{K}{\epsilon^2}\right)$, as desired. Thus, we have shown that we can remove the factor of $\frac{1}{L}$ from the immediate application of the result of Krishnamurthy et al. [2016].

A.2.2 Linear Growth in $|\Theta|$ with Small Hypothesis Class

The standard construction of Chen and Jiang [2019] as well as Krishnamurthy et al. [2016] consider the problem for single MDPs and POMDPs, respectively, and note that it is information-theoretically equivalent to the best-arm identification identification problem, with number of arms equal to the size of the model class. Unfortunately, in our setting, this is not the case, however we are still able to perform an information-theoretic analysis which is very similar to that of Krishnamurthy et al. [2016] as well as Auer et al. [2003].

First, we consider the following family of K-arm bandit problems. Let S_K denote the set of permutations on $[K]$, and let Γ denote a maximal subset of S_K such that any two permutations in Γ disagree on all inputs. It is clear that $|\Gamma| \leq K$, and that equality can be achieved as follows: using the standard group theoretic notation, we simply let $\Gamma = \langle (1\,2\cdots K) \rangle$; that is, the subgroup of S_K generated by the permutation taking $i \mapsto i+1$ for $i = 1, \ldots, K-1$, and taking $K \mapsto 1$. It is clear that $|\Gamma| = K$ and that

$$\forall \sigma \neq \rho \in \Gamma, \sigma(i) \neq \rho(i), \forall i \in [K].$$

For $\sigma \in \Gamma$, let $\mathcal{M}_\sigma$ denote a family of K bandit problems, each indexed by $\theta \in [K]$; specifically, $\mathcal{M}_\sigma^\theta$ is a K-arm bandit with optimal arm $\sigma(\theta)$, giving reward from $\text{Bern}\left(\frac{1}{2} + \frac{3}{2}\epsilon\right)$ and all other arms giving reward from $\text{Bern}(1/2)$. The hypothesis class of bandit families is thus

$$\mathcal{H} = \{\mathcal{M}_\sigma | \sigma \in \Gamma\}.$$

When the target family is $\mathcal{M}_\sigma$, at each stage, a value of θ is drawn i.i.d from $\text{Unif}([K])$, and the agent is allowed to interact with the bandit $\mathcal{M}_\sigma^\theta$ for L iterations. We would like to lower bound m, the number of draws of θ required for any algorithm to choose an ϵ-optimal permutation (where, in this case, the value of the permutation σ is evaluated the same sense as in the standard CMDP setting: by taking $\mathbb{E}_\theta[R(\sigma(\theta))])$ with confidence δ. Notice that, in this setting, the only ϵ-optimal permutation is the optimal one: suppose that σ^* is the optimal permutation, and σ is some other permutation in Γ. Then since σ and σ^* disagree on all values in $[K]$, we will have that the value of σ will be $\frac{3}{2}\epsilon > \epsilon$ less than σ^*, since, for each value of θ, σ selects a suboptimal arm, thus implying that σ^* is indeed the only ϵ-optimal permutation.

To achieve a lower bound, we represent a deterministic algorithm as sequence of mappings $f_t : (\{0,1\} \times \Theta)^t \to \Gamma$ for $0 \leq t \leq mL$, where we will interpret $f_t(r^t)$ as the permutation the algorithm believes is optimal after seeing the reward sequence $r^t \in (\{0,1\} \times \Theta)^t$ (our reward sequences are 'labelled' by the value of θ which they correspond to; that is, the reward sequence is a t-long sequence of ordered pairs, where the ith pair is the pair containing the reward at time i and the value of θ at time i). In particular, this belief satisfies the property that for $(k-1)L+1 \leq t \leq kL$, $\left(f_t(r^t)\right)(\theta_k)$ is the action taken by the algorithm after seeing reward sequence r^t (where here the m values of θ drawn are $\theta_1, \ldots, \theta_m$). The return of the algorithm after seeing the reward sequence r^{mL} is thus the permutation $f_{mL}(r^{mL})$.

Following the notation of Krishnamurthy et al. [2016] we will then let $\mathbb{P}_{\sigma^*, f}$ denote the distribution over all mL rewards when σ^* is the optimal permutation and the algorithm and actions selected are done so according to f. We let $\mathbb{P}_{0,f}$ denote the same distribution as above, except in the setting when all permutations give the same family of bandits whose rewards are always drawn from $\text{Bern}(1/2)$.

Now, we have the following

$$|\mathbb{P}_{\sigma^*, f}(f_{mL} = \sigma^*) - \mathbb{P}_{0,f}(f_{mL} = \sigma^*)| \leq |\mathbb{P}_{\sigma^*, f} - \mathbb{P}_{0,f}|_{TV} \leq \sqrt{\frac{1}{2} KL(\mathbb{P}_{0,f} || \mathbb{P}_{\sigma^*, f})},$$

where the last inequality is Pinsker's. From the definition of KL divergence, we then have that

$$KL(\mathbb{P}_{0,f} || \mathbb{P}_{\sigma^*, f}) = \sum_{r^{mL} \in \{0,1\}^{mL}} \mathbb{P}_{0,f}(r^{mL}) \log\left(\frac{\mathbb{P}_{0,f}(r^{mL})}{\mathbb{P}_{\sigma^*, f}(r^{mL})}\right),$$

which, by the chain rule for KL divergence is

$$\sum_{t=1}^{mL} \sum_{r^t \in \{0,1\}^t} \mathbb{P}_{0,f}(r^t) \log\left(\frac{\mathbb{P}_{0,f}(r^t | r^{t-1})}{\mathbb{P}_{\sigma^*, f}(r^t | r^{t-1})}\right),$$

which, by noticing that if, at time t, any suboptimal action is selected, the log ratio will be zero, letting

$$\Lambda_{t-1} = \{r^{t-1} \in \{0,1\}^{t-1} : \left(f_{t-1}(r^{t-1})\right)(\theta_{\lfloor \frac{t-2}{L} \rfloor + 1}) = \sigma^*(\theta_{\lfloor \frac{t-2}{L} \rfloor + 1})\},$$

the above is simply

$$\sum_{t=1}^{mL} \sum_{r^{t-1} \in \Lambda_{t-1}} \mathbb{P}_{0,f}(r^{t-1}) \left(\sum_{x \in \{0,1\}} \mathbb{P}_{0,f}(x) \log \left(\frac{\mathbb{P}_{0,f}(r^t | a_t \text{ is optimal})}{\mathbb{P}_{\sigma^*,f}(r^t | a_t \text{ is optimal})} \right) \right)$$

$$= \sum_{t=1}^{mL} \sum_{r^{t-1} \in \Lambda_{t-1}} \mathbb{P}_{0,f}(r^{t-1}) \left(\frac{1}{2} \log \left(\frac{1/2}{1/2 - \frac{3}{2}\epsilon} \right) + \frac{1}{2} \log \left(\frac{1/2}{1/2 + \frac{3}{2}\epsilon} \right) \right)$$

$$= \left(-\frac{1}{2} \log \left(1 - 9\epsilon^2 \right) \right) \sum_{t=1}^{mL} \sum_{r^{t-1} \in \Lambda_{t-1}} \mathbb{P}_{0,f}(r^{t-1})$$

$$= \left(-\frac{1}{2} \log \left(1 - 9\epsilon^2 \right) \right) \sum_{t=1}^{mL} \mathbb{P}_{0,f}[\left(f_{t-1}(r^{t-1})\right)(\theta_{\lfloor \frac{t-2}{L} \rfloor + 1}) = \sigma^*(\theta_{\lfloor \frac{t-2}{L} \rfloor + 1})],$$

where in the above we have slightly abused notation, and any reference to Λ_{t-1} or $f_{t-1}(r^{t-1})$ for $t = 1$ refers to the initial action that the algorithm takes. Thus, again with the notation of Krishnamurthy et al. [2016], letting N_{σ^*} denote the random variable equal to the number of times the algorithm chose the arm given by σ^* in the bandit family where all arms are always equal, we have that

$$\mathbb{P}_{\sigma^*,f}(f_{mL} = \sigma^*) - \mathbb{P}_{0,f}(f_{mL} = \sigma^*) \leq \frac{1}{2} \sqrt{-\mathbb{E}_{0,f}[N_{\sigma^*}] \log \left(1 - 9\epsilon^2 \right)}.$$

Taking the expectation over all possible optimal permutations then gives that

$$\frac{1}{|\Gamma|} \sum_{\sigma^* \in \Gamma} \mathbb{P}_{\sigma^*,f}(f_{mL} = \sigma^*) \leq \frac{1}{|\Gamma|} \sum_{\sigma^* \in \Gamma} \mathbb{P}_{0,f}(f_{mL} = \sigma^*) + \frac{1}{2|\Gamma|} \sum_{\sigma^* \in \Gamma} \sqrt{-\mathbb{E}_{0,f}[N_{\sigma^*}] \log \left(1 - 9\epsilon^2 \right)},$$

which, by Jensen's inequality is at most

$$\frac{1}{|\Gamma|} \sum_{\sigma^* \in \Gamma} \mathbb{P}_{0,f}(f_{mL} = \sigma^*) + \frac{1}{2} \sqrt{-\frac{\log \left(1 - 9\epsilon^2 \right)}{|\Gamma|} \sum_{\sigma^* \in \Gamma} \mathbb{E}_{0,f}[N_{\sigma^*}]}.$$

Now, notice that we can write

$$\sum_{\sigma^* \in \Gamma} N_{\sigma^*} \leq \sum_{t=1}^{mL} |\{\sigma \in \Gamma : \left(f_{t-1}(r^{t-1})\right)(\theta_{\lfloor \frac{t-2}{L} \rfloor + 1}) = \sigma(\theta_{\lfloor \frac{t-2}{L} \rfloor + 1})\}|,$$

which, based on how we chose Γ, is at most mL. So, the above is at most

$$\frac{1}{|\Gamma|} + \frac{1}{2} \sqrt{-\frac{\log \left(1 - 9\epsilon^2 \right) mL}{|\Gamma|}}.$$

Noticing that $-\log(1-x) \leq 2x$ when $x \leq 1/2$, if $9\epsilon^2 K^2 \leq \frac{1}{2}$, then the above is at most

$$\frac{1}{|\Gamma|} + \frac{1}{2} \sqrt{\frac{18\epsilon^2 mL}{|\Gamma|}}.$$

Thus, if we take

$$m < \frac{|\Gamma|}{162\epsilon^2 L} = \frac{K}{162\epsilon^2 L},$$

the above will be at most $2/3$ since certainly $|\Gamma| \geq 2$, thus indicating that if m is less than the above, we cannot achieve suboptimality of at most ϵ for any $\delta < 2/3$.

The construction for achieving the $\Omega\left(\frac{|A|^L}{L\epsilon^2}\right)$ is then achieved in a similar way in Krishnamurthy et al. [2016]. This can be done by describing the MDP for the context θ as a tree with special edge given by $\sigma^*(\theta)$. The hypothesis class is then $\langle (1\,2\, \cdots\, |A|^{L-1}) \rangle \subset S_{|A|^{L-1}}$, and $\Theta = [|A|^{L-1}]$. Thus, our lower bound grows at the rate $\Omega\left(\frac{|\Theta|}{L\epsilon^2}\right)$ while having a hypothesis class which has size $|\Theta|$.

A.3 FQI Analysis

First we give pseudocode describing FQI for CMDPs:

Algorithm 1 CMDP-FQI

   Input: Dataset $\mathcal{D} = \{(\theta_i, s_l^{\theta_i}, a_l^{\theta_i}, r_l^{\theta_i}, s_l^{\theta_{i'}})\}_{i=1,l=0}^{i=m,l=L-1}$
   Initialize $Q : S \times A \times \Theta \times \{0, \dots, L\} \to \mathbb{R}$ with $Q(s, a, \theta, l) = 0, \forall s, a, \theta, l$
   for $l = 1, \dots, L$ **do**
      $\mathcal{D}_l \leftarrow \emptyset$
      for $(\theta_i, s_l^{\theta_i}, a_l^{\theta_i}, r_l^{\theta_i}, s_l^{\theta_{i'}}) \in \mathcal{D}$ **do**
         input $\leftarrow (s_l^{\theta_i}, a_l^{\theta_i}, \theta_i, l)$
         target $\leftarrow r_l^{\theta_i} + \max_{a \in A} Q(s_l^{\theta_{i'}}, a, \theta_i, l-1)$
         $\mathcal{D}_l \leftarrow \mathcal{D}_l \cup \{(\text{input}, \text{target})\}$
      end for
      $Q_l(\cdot, \cdot, \cdot, l) = \texttt{REGRESS}(\mathcal{D}_l)$
   end for
   return Q

In the above, the program $\texttt{REGRESS}(\mathcal{D}_l)$ returns the Q function in $\mathcal{F}$ which, for the lth timestep, minimizes the empirical ℓ_2 norm of the dataset. We use the notation of subscripting the Q function by l to indicate that the $l+1$th Q function need not be the exact Bellman backup of the lth.

Now, our setup and derivation follows the same structure as Chen and Jiang [2019]. First, just as in Chen and Jiang [2019], we define a semi-norm on real-valued functions with state-action inputs and also write FQI in terms of the backups $\mathcal{T}_l$.

Definition A.1.1. *Define a semi-norm $|| \cdot ||_{p, \nu \times \mathcal{P}_\Theta}$ on functions $f : S \times A \times \Theta \to \mathbb{R}$ by*

$$||f||_{p, \nu \times \mathcal{P}_\Theta} = \left(\mathbb{E}_{\theta \sim \mathcal{P}_\Theta, (s,a) \sim \nu(\theta)}[f(s, a, \theta)^{1/p}] \right)^p$$

Now, we can write the FQI above in terms of backups. For $l = 0, \dots, L-1$, we say that $f(\cdot, \cdot, \cdot, l+1) = \widehat{\mathcal{T}_l^{\mathcal{F}}} f(\cdot, \cdot, \cdot, l)$, where, for $f' \in \mathcal{F}$,

$$\widehat{\mathcal{T}_l^{\mathcal{F}}} f' = \text{argmin}_{f \in \mathcal{F}} \mathcal{L}_l(f; f'),$$

and

$$\mathcal{L}_l(f; f') = \frac{1}{|\mathcal{D}|} \sum_{(\theta, s^\theta, a_l^\theta, r^\theta, s^{\theta'}) \in \mathcal{D}} (f(s^\theta, a^\theta, \theta, l+1) - r^\theta - V_{f'}(s^{\theta'}, \theta, l))^2.$$

We now state and prove the lemmas needed for the sample complexity upper bound as in Chen and Jiang [2019].

Lemma A.1.1. *Let μ be the data distribution and ν an admissible distribution. Then*

$$|| \cdot ||_{2, \nu \times \mathcal{P}_\Theta} \leq \sqrt{C} || \cdot ||_{2, \mu \times \mathcal{P}_\Theta},$$

where $\nu \times \mathcal{P}_\Theta$ is the distribution over (θ, s, a) triples with $\theta \sim \mathcal{P}_\Theta$, and $(s, a) \sim \nu(\theta)$, and $\mu \times \Theta$ is the distribution over (θ, s, a) triples with $\theta \sim \mathcal{P}_\Theta$, and $(s, a) \sim \mu$.

Proof.

$$||f||_{2, \nu \times \mathcal{P}_\Theta} = \sqrt{\sum_{\theta \in \Theta} \mathcal{P}_\Theta(\theta) \sum_{(s,a) \in S \times A} (\nu(\theta))(s,a) f(s,a)^2}$$

$$\leq \sqrt{C \sum_{\theta \in \Theta} \mathcal{P}_\Theta(\theta) \sum_{(s,a) \in S \times A} \mu(s,a) f(s,a)^2} = \sqrt{C} ||f||_{2, \mu \times \mathcal{P}_\Theta}$$

$\square$

We now rederive a standard infinite-horizon lemma from Kakade and Langford [2002] for our finite-horizon setting.

Lemma A.1.2. *Let π^* denote an optimal possibly time-dependent deterministic policy and let $\hat{\pi}$ be some time-dependent deterministic policy. Letting $\mathcal{M}_l^\pi(s, \theta)$ denote the marginal distribution of the lth state, given θ and the initial state s by following policy π, then*

$$V_{\pi^*}^L(s; \theta) - V_{\hat{\pi}}^L(s; \theta) = \sum_{l=0}^{L-1} \mathbb{E}_{s_l \sim \mathcal{M}_l^{\hat{\pi}}(s, \theta)}[V_{\pi^*}^{L-l}(s_l; \theta) - Q^*(s_l, \hat{\pi}, \theta, L-l)].$$

Proof.

$$V_L^{\hat{\pi}}(s; \theta) = \sum_{l=0}^{L-1} \mathbb{E}_{s_l \sim \mathcal{M}_l^{\hat{\pi}}(s, \theta)}[R(s_l, \hat{\pi}(s_l, \theta, l)]$$

$$= \sum_{l=0}^{L-1} \mathbb{E}_{s_l \sim \mathcal{M}_l^{\hat{\pi}}(s, \theta)}[R(s_l, \hat{\pi}(s_l, \theta, l)) + V_{\pi^*}^{L-l}(s_l; \theta) - V_{\pi^*}^{L-l}(s_l; \theta)]$$

$$= \sum_{l=0}^{L-2} \mathbb{E}_{s_l \sim \mathcal{M}_l^{\hat{\pi}}(s, \theta), s_{l+1} \sim \mathcal{M}_{l+1}^{\hat{\pi}}(s, \theta)}[R(s_l, \hat{\pi}(s_l, \theta, l)) + V_{\pi^*}^{L-l-1}(s_{l+1}; \theta) - V_{\pi^*}^{L-l}(s_l; \theta)]$$

$$+ V_{\pi^*}^L(s_0; \theta) + \mathbb{E}_{s_{L-1} \sim \mathcal{M}_{L-1}^{\hat{\pi}}(s, \theta)}[R(s_{L-1}, \hat{\pi}(s_{L-1}, \theta, L-1)) - V_{\pi^*}^1(s_{L-1}; \theta)]$$

$$\implies V_L^{\hat{\pi}}(s; \theta) - V_{\pi^*}^L(s_0; \theta) = \sum_{l=0}^{L-1} \mathbb{E}_{s_l \sim \mathcal{M}_l^{\hat{\pi}}(s, \theta)}[Q^*(s_l, \hat{\pi}, \theta, L-l) - V_{\pi^*}^{L-l}(s_l; \theta)],$$

as desired. $\qquad \square$

With this, we now show how to control the difference in value of two policies in terms of the semi-norm of the difference of their Q-functions.

Lemma A.1.3. *Let $f : S \times A \times \Theta \times \{0, \ldots, L\}$ and let $\hat{\pi} = \pi_f$ denote the policy, which at time-step t is greedy with respect to $f(\cdot, \cdot, \theta, t)$ for all t. Then*

$$v_{\pi^*}^L - v_{\hat{\pi}}^L \leq \sum_{l=0}^{L-1} ||Q^*(\cdot, \cdot, \cdot, L-l) - f(\cdot, \cdot, \cdot, L-l)||_{2, \mathcal{M}_l^{\hat{\pi}} \times \pi^* \times \mathcal{P}_\Theta} + ||Q^*(\cdot, \cdot, \cdot, L-l) - f(\cdot, \cdot, \cdot, L-l)||_{2, \mathcal{M}_l^{\hat{\pi}} \times \hat{\pi} \times \mathcal{P}_\Theta},$$

where the notation $\mathcal{M}_l^\pi \times \pi' \times \mathcal{P}_\Theta$ denote the distribution over state action context triples in which $\theta \sim \mathcal{P}_\Theta$, then lth state, s, is drawn from $\mathcal{M}_l^\pi(\theta)$ (i.e. the marginal distribution conditioned just on θ) and, $a = \pi'(s, l)$.

Proof. From the above, we have that

$$v_{\pi^*}^L - v_{\hat{\pi}}^L = \sum_{l=0}^{L-1} \mathbb{E}_{\theta \sim \mathcal{P}_\Theta, s \sim \mathcal{P}_{s_0}, s_l \sim \mathcal{M}_l^{\hat{\pi}}(s, \theta)}[V_{\pi^*}^{L-l}(s_l; \theta) - Q^*(s_l, \hat{\pi}, \theta, L-l)]$$

$$\leq \sum_{l=0}^{L-1} \mathbb{E}_{\theta \sim \mathcal{P}_\Theta, s \sim \mathcal{P}_{s_0}, s_l \sim \mathcal{M}_l^{\hat{\pi}}(s, \theta)}[V_{\pi^*}^{L-l}(s_l; \theta) - f(s, \pi^*, \theta, L-l) + f(s, \hat{\pi}, \theta, L-l) - Q^*(s_l, \hat{\pi}, \theta, L-l)]$$

$$\leq \sum_{l=0}^{L-1} ||Q^*(\cdot, \cdot, \cdot, L-l) - f(\cdot, \cdot, \cdot, L-l)||_{1, \mathcal{M}_l^{\hat{\pi}} \times \pi^* \times \mathcal{P}_\Theta} + ||Q^*(\cdot, \cdot, \cdot, L-l) - f(\cdot, \cdot, \cdot, L-l)||_{1, \mathcal{M}_l^{\hat{\pi}} \times \hat{\pi} \times \mathcal{P}_\Theta}$$

$$\leq \sum_{l=0}^{L-1} ||Q^*(\cdot, \cdot, \cdot, L-l) - f(\cdot, \cdot, \cdot, L-l)||_{2, \mathcal{M}_l^{\hat{\pi}} \times \pi^* \times \mathcal{P}_\Theta} + ||Q^*(\cdot, \cdot, \cdot, L-l) - f(\cdot, \cdot, \cdot, L-l)||_{2, \mathcal{M}_l^{\hat{\pi}} \times \hat{\pi} \times \mathcal{P}_\Theta},$$

as desired. $\qquad \square$

Lemma A.1.4. *Let $f, f' : S \times A \times \Theta \times \{0, \ldots, L\} \to \mathbb{R}$ and $\pi_{f, f'}(s, \theta, l) = \arg\max_{a \in A} \max\{f(s, a, \theta, l), f'(s, a, \theta, l)\}$. Then, for all admissible distributions ν and all $l \in \{0, \ldots, L\}$,*

$$||V_f^l(\cdot; \cdot) - V_{f'}^l(\cdot; \cdot)||_{2, P(\nu) \times \mathcal{P}_\Theta} \leq ||f(\cdot, \cdot, \cdot, l) - f'(\cdot, \cdot, \cdot, l)||_{2, P(\nu) \times \pi_{f, f'}(\cdot, \cdot, l) \times \mathcal{P}_\Theta},$$

where $P(\nu) \times \mathcal{P}_\Theta$ denotes the distribution over context-state pairs (s', θ) given that $\theta \sim \mathcal{P}_\Theta$, $(s, a) \sim \nu(\theta)$, and $s' \sim T(\cdot|s, a, \theta)$, and $P(\nu) \times \pi_{f, f'}(\cdot, \cdot, l) \times \mathcal{P}_\Theta$ denotes the distribution over triples (s', a, θ) where $(s', \theta) \sim P(\nu) \times \mathcal{P}_\Theta$ and $a = \pi_{f, f'}(s, \theta, l)$.

Proof. We have that

$$||V_f^l(\cdot;\cdot) - V_{f'}^l(\cdot;\cdot)||_{2,P(\nu)\times\mathcal{P}_\Theta}^2$$

$$= \sum_{\theta\in\Theta}\mathcal{P}_\Theta(\theta)\sum_{(s,a)\in S\times A}(\nu(\theta))(s,a)\sum_{s'\in S}T(s'|s,a,\theta)\left(\max_{a\in A}f(s',a,\theta,l) - \max_{a\in A}f'(s',a,\theta,l)\right)^2$$

$$\leq \sum_{\theta\in\Theta}\mathcal{P}_\Theta\sum_{(s,a)\in S\times A}(\nu(\theta))(s,a)\sum_{s'\in S}T(s'|s,a,\theta)(f(s',\pi_{f,f'}(s',\theta,l),\theta,l) - f'(s',\pi_{f,f'}(s',\theta,l),\theta,l))^2$$

$$= ||f(\cdot,\cdot,\cdot,l) - f'(\cdot,\cdot,\cdot,l)||_{2,P(\nu)\times\pi_{f,f'}(\cdot,\cdot,l)\times\mathcal{P}_\Theta}^2,$$

as desired. $\square$

We now upper bound $||f - Q^*||_{2,\nu\times\mathcal{P}_\Theta}$.

Lemma A.1.5. *For any data distribution μ over $S\times A$ and admissible distribution ν, policy (which is potentially time-dependent), π, and $f_{l+1}, f_l : S\times A\times\Theta\times\{0,\ldots,L-1\}\to R$ gotten from FQI, then we have*

$$||f_{l+1}(\cdot,\cdot,\cdot,l+1) - Q^*(\cdot,\cdot,\cdot,l+1)||_{2,\nu\times\mathcal{P}_\Theta}$$

$$\leq \sqrt{C}||f_{l+1}(\cdot,\cdot,\cdot,l+1) - (\mathcal{T}_l(\cdot))f_l(\cdot,\cdot,\cdot,l)||_{2,\mu\times\mathcal{P}_\Theta} + ||f_l(\cdot,\cdot,\cdot,l) - Q^*(\cdot,\cdot,\cdot,l)||_{2,P(\nu)\times\pi_{f_l,Q^*}(\cdot,\cdot,l)\times\mathcal{P}_\Theta},$$

for all l.

Proof. We have

$$||f_{l+1}(\cdot,\cdot,\cdot,l+1) - Q^*(\cdot,\cdot,\cdot,l+1)||_{2,\nu\times\mathcal{P}_\Theta}$$

$$= ||f_{l+1}(\cdot,\cdot,\cdot,l+1) - (\mathcal{T}_l(\cdot))f_l(\cdot,\cdot,\cdot,l) + (\mathcal{T}_l(\cdot))f_l(\cdot,\cdot,\cdot,l) - Q^*(\cdot,\cdot,\cdot,l+1)||_{2,\nu\times\mathcal{P}_\Theta}$$

$$\leq ||f_{l+1}(\cdot,\cdot,\cdot,l+1) - (\mathcal{T}_l(\cdot))f_l(\cdot,\cdot,\cdot,l)||_{2,\nu\times\mathcal{P}_\Theta} + ||(\mathcal{T}_l(\cdot))f_l(\cdot,\cdot,\cdot,l) - (\mathcal{T}_l(\cdot))Q^*(\cdot,\cdot,\cdot,l)||_{2,\nu\times\mathcal{P}_\Theta}.$$

Now, notice that

$$||(\mathcal{T}_l(\cdot))f_l(\cdot,\cdot,\cdot,l) - (\mathcal{T}_l(\cdot))Q^*(\cdot,\cdot,\cdot,l)||_{2,\nu\times\mathcal{P}_\Theta}^2$$

$$= \mathbb{E}_{\theta\sim\mathcal{P}_\Theta,(s,a)\sim\nu(\theta)}[((\mathcal{T}_l(\theta))f_l(s,a,\theta,l) - (\mathcal{T}_l(\theta))Q^*(s,a,\theta,l))^2]$$

$$= \mathbb{E}_{\theta\sim\mathcal{P}_\Theta,(s,a)\sim\nu(\theta)}\left[\left(\mathbb{E}_{s'\sim T(\cdot|s,a,\theta)}[V_{f_l}^l(s';\theta) - V_{\pi^*}^l(s';\theta)]\right)^2\right]$$

$$\leq \mathbb{E}_{\theta\sim\mathcal{P}_\Theta,(s,a)\sim\nu(\theta),s'\sim T(\cdot|s,a,\theta)}\left[\left(V_{f_l}^l(s';\theta) - V_{\pi^*}^l(s';\theta)\right)^2\right] = ||V_{f_l}^l(\cdot;\cdot) - V_{\pi^*}^l(\cdot;\cdot)||_{2,P(\nu)\times\mathcal{P}_\Theta}^2.$$

Thus, the above is at most

$$||f_{l+1}(\cdot,\cdot,\cdot,l+1) - (\mathcal{T}_l(\cdot))f_l(\cdot,\cdot,\cdot,l)||_{2,\nu\times\mathcal{P}_\Theta} + ||V_{f_l}^l(\cdot;\cdot) - V_{\pi^*}^l(\cdot;\cdot)||_{2,P(\nu)\times\mathcal{P}_\Theta}$$

$$\leq \sqrt{C}||f_{l+1}(\cdot,\cdot,\cdot,l+1) - (\mathcal{T}_l(\cdot))f_l(\cdot,\cdot,\cdot,l)||_{2,\mu\times\mathcal{P}_\Theta} + ||V_{f_l}(\cdot,\cdot,l) - V_{\pi^*}^l(\cdot;\cdot)||_{2,P(\nu)\times\mathcal{P}_\Theta},$$

which, by the previous lemma, is at most

$$\sqrt{C}||f_{l+1}(\cdot,\cdot,\cdot,l+1) - (\mathcal{T}_l(\cdot))f_l(\cdot,\cdot,\cdot,l)||_{2,\mu\times\mathcal{P}_\Theta} + ||f_l(\cdot,\cdot,\cdot,l) - Q^*(\cdot,\cdot,\cdot,l)||_{2,P(\nu)\times\pi_{f_l,Q^*}(\cdot,\cdot,l)\times\mathcal{P}_\Theta},$$

as desired. $\square$

We now assume access to a function class $\mathcal{G}$ which approximates Bellman backups of $\mathcal{F}$. In particular, we define the error of approximation as follows:

Definition A.1.2. *Define $\epsilon_{\mathcal{F},\mathcal{G}}$ to be the smallest real number such that if we let*

$$g_f^*(l) = \mathrm{argmin}_{g\in\mathcal{G}}||g(\cdot,\cdot,\cdot,l+1) - (\mathcal{T}_l(\cdot))f(\cdot,\cdot,\cdot,l)||_{2,\mu\times\mathcal{P}_\Theta},$$

then

$$||(g_f^*(l))(\cdot,\cdot,\cdot,l+1) - (\mathcal{T}_l(\cdot))f(\cdot,\cdot,\cdot,l)||_{2,\mu\times\mathcal{P}_\Theta}^2 \leq \epsilon_{\mathcal{F},\mathcal{G}}, \forall l.$$

Lemma A.1.6. *Let the dataset $\mathcal{D}$ be generated as described in the text. Then, for any $l\in\{0,\ldots,L\}$ and $f\in\mathcal{F}$, we have that, with probability at least $1-\delta$,*

$$\mathbb{E}_\mathcal{D}[\mathcal{L}_l(\widehat{\mathcal{T}^\mathcal{G}}_l f; f) - \mathcal{L}_l(g_f^*(l); f)] \leq \frac{56L^2\ln\left(\frac{L|\mathcal{F}||\mathcal{G}|}{\delta}\right)}{3m} + \sqrt{\frac{32L^2\ln\left(\frac{L|\mathcal{F}||\mathcal{G}|}{\delta}\right)}{m}}\epsilon_{\mathcal{F},\mathcal{G}}$$

Proof. Fix l, and just as in Chen and Jiang [2019], define

$$X(g, f, g_f^*(l)) = (g(s, a, \theta, l+1) - r - V_f(s', \theta, l))^2 - ((g_f^*(l))(s, a, \theta, l+1) - r - V_f(s', \theta, l))^2.$$

Letting $X_{i,l}$ denote the evaluation of X on $(\theta_i, s_l^{\theta_i}, a_l^{\theta_i}, r_l^{\theta_i}, s_l^{\theta_i'})$ it is clear that

$$\frac{1}{mL} \sum_{i=1}^m \sum_{l=0}^{L-1} X_{i,l}(g, f, g_f^*(l)) = \mathcal{L}_l(g; f) - \mathcal{L}_l(g_f^*(l); f).$$

Now, we have that

$$\mathrm{Var}[X(g, f, g_f^*(l))] \le \mathbb{E}[X(g, f, g_f^*(l))^2]$$

$$= \mathbb{E}[((g(s, a, \theta, l+1) - r - V_f(s', \theta, l))^2 - ((g_f^*(l))(s, a, \theta, l+1) - r - V_f(s', \theta, l))^2)^2]$$

$$= \mathbb{E}[(g(s, a, \theta, l+1) - (g_f^*(l))(s, a, \theta, l+1))^2 (g(s, a, \theta, l+1) + (g_f^*(l))(s, a, \theta, l+1) - 2r - 2V_f(s', \theta, l))^2].$$

Now, since all rewards are in $[0, 1]$,

$$g(s, a, \theta, l+1) + (g_f^*(l))(s, a, \theta, l+1) - 2r - 2V_f(s', \theta, l) \le 2l$$

and so the above is at most

$$4l^2 \mathbb{E}[(g(s, a, \theta, l+1) - (g_f^*(l))(s, a, \theta, l+1))^2] = 4l^2 \|g(\cdot, \cdot, \cdot, l+1) - (g_f^*(l))(\cdot, \cdot, \cdot, l+1)\|_{2, \mu \times \mathcal{P}_\Theta}^2$$

$$\le 8l^2 \left(\|g(\cdot, \cdot, \cdot, l+1) - (\mathcal{T}_l(\cdot))f(\cdot, \cdot, \cdot, l)\|_{2, \mu \times \mathcal{P}_\Theta}^2 + \|(\mathcal{T}_l(\cdot))f(\cdot, \cdot, \cdot, l) - (g_f^*(l))(\cdot, \cdot, \cdot, l+1)\|_{2, \mu \times \mathcal{P}_\Theta}^2 \right)$$

$$= 8l^2 \left(\|g(\cdot, \cdot, \cdot, l+1) - (\mathcal{T}_l(\cdot))f(\cdot, \cdot, \cdot, l)\|_{2, \mu \times \mathcal{P}_\Theta}^2 - \|(\mathcal{T}_l(\cdot))f(\cdot, \cdot, \cdot, l) - (g_f^*(l))(\cdot, \cdot, \cdot, l+1)\|_{2, \mu \times \mathcal{P}_\Theta}^2 \right.$$

$$\left. + 2\|(\mathcal{T}_l(\cdot))f(\cdot, \cdot, \cdot, l) - (g_f^*(l))(\cdot, \cdot, \cdot, l+1)\|_{2, \mu \times \mathcal{P}_\Theta}^2 \right)$$

$$= 8l^2 \left(\mathbb{E}[X(g, f, g_f^*(l))] + 2\|(\mathcal{T}_l(\cdot))f(\cdot, \cdot, \cdot, l) - (g_f^*(l))(\cdot, \cdot, \cdot, l+1)\|_{2, \mu \times \mathcal{P}_\Theta}^2 \right)$$

$$\le 8l^2 \left(\mathbb{E}[X(g, f, g_f^*(l))] + 2\epsilon_{\mathcal{F}, \mathcal{G}} \right).$$

Now, notice that the random variables $X_{1,l}, X_{2,l}, \ldots, X_{m,l}$ are i.i.d for fixed l. Thus, as in Chen and Jiang [2019], we will use one-sided Bernstein's inequality along with a union-bound over $f \in \mathcal{F}, g \in \mathcal{G}$ and finally over $l = 1, \ldots, L$. In particular, as Chen and Jiang [2019] note, union-bounding over the first two sets says that, with probability at least $1 - \delta$, every $f \in \mathcal{F}, g \in \mathcal{G}$, and fixed l satisfy

$$\mathbb{E}[X(g, f, g_f^*(l))] - \sum_{i=1}^m X_{i,l}(g, f, g_f^*(l)) \le \sqrt{\frac{2\mathrm{Var}[X(g, f, g_f^*(l))] \ln\left(\frac{|\mathcal{F}||\mathcal{G}|}{\delta}\right)}{m}} + \frac{4L^2 \ln\left(\frac{|\mathcal{F}||\mathcal{G}|}{\delta}\right)}{3m}$$

$$\le \sqrt{\frac{16l^2 \left(\mathbb{E}[X(g, f, g_f^*(l))] + 2\epsilon_{\mathcal{F}, \mathcal{G}} \right) \ln\left(\frac{|\mathcal{F}||\mathcal{G}|}{\delta}\right)}{m}} + \frac{4L^2 \ln\left(\frac{|\mathcal{F}||\mathcal{G}|}{\delta}\right)}{3m}.$$

Now, notice that

$$\frac{1}{m} \sum_{i=1}^m X_{i,l}(\widehat{\mathcal{T}}_l^{\mathcal{G}} f, f, g_f^*(l)) \le \frac{1}{m} \sum_{i=1}^m X_{i,l}(g_f^*(l), f, g_f^*(l)) = 0,$$

and so we really have that, with probability at least $1 - \delta$, every $f \in \mathcal{F}$ satisfies

$$\mathbb{E}[X(\widehat{\mathcal{T}}_l^{\mathcal{G}} f, f, g_f^*(l))] \le \sqrt{\frac{16l^2 \left(\mathbb{E}[X(\widehat{\mathcal{T}}_l^{\mathcal{G}} f, f, g_f^*(l))] + 2\epsilon_{\mathcal{F}, \mathcal{G}} \right) \ln\left(\frac{|\mathcal{F}||\mathcal{G}|}{\delta}\right)}{m}} + \frac{4L^2 \ln\left(\frac{|\mathcal{F}||\mathcal{G}|}{\delta}\right)}{3m},$$

which, by algebraic manipulations (see Lemma 16 of Chen and Jiang [2019]) implies

$$\mathbb{E}[X(\widehat{\mathcal{T}}_l^{\mathcal{G}} f, f, g_f^*(l))] \le \frac{56L^2 \ln\left(\frac{|\mathcal{F}||\mathcal{G}|}{\delta}\right)}{3m} + \sqrt{\frac{32L^2 \ln\left(\frac{|\mathcal{F}||\mathcal{G}|}{\delta}\right)}{m} \epsilon_{\mathcal{F}, \mathcal{G}}},$$

for the fixed l. Having union-bounded with respect to l will yield the desired upper bound of

$$\frac{56L^2 \ln\left(\frac{L|\mathcal{F}||\mathcal{G}|}{\delta}\right)}{3m} + \sqrt{\frac{32L^2 \ln\left(\frac{L|\mathcal{F}||\mathcal{G}|}{\delta}\right)}{m} \epsilon_{\mathcal{F}, \mathcal{G}}}.$$

$\square$

Finally, we derive a sample-complexity bound:

Theorem A.2. *We can bound the error of the hypothesis returned by FQI, with probability at least $1 - \delta$ as*

$$v_{\pi^*}^L - v_{\pi_f}^L \leq L(L+1) \sqrt{C \left(\frac{56L^2 \ln\left(\frac{L|\mathcal{F}|^2}{\delta}\right)}{3m} + \sqrt{\frac{32L^2 \ln\left(\frac{L|\mathcal{F}|^2}{\delta}\right)}{m}} \epsilon_{\mathcal{F},\mathcal{F}} + \epsilon_{\mathcal{F},\mathcal{F}} \right)}.$$

Proof. From Lemma A.1.5 we have

$$||f_{l+1}(\cdot,\cdot,\cdot,l+1) - Q^*(\cdot,\cdot,\cdot,l+1)||_{2,\nu \times \mathcal{P}_\Theta}$$

$$\leq \sqrt{C}||f_{l+1}(\cdot,\cdot,\cdot,l+1) - (\mathcal{T}_l(\cdot))f_l(\cdot,\cdot,\cdot,l)||_{2,\mu \times \mathcal{P}_\Theta} + ||f_l(\cdot,\cdot,\cdot,l) - Q^*(\cdot,\cdot,\cdot,l)||_{2,P(\nu) \times \pi_{f_l}, Q^*(\cdot,\cdot,l) \times \mathcal{P}_\Theta},$$

for all l. We bound the first term:

$$||f_{l+1}(\cdot,\cdot,\cdot,l+1) - (\mathcal{T}_l(\cdot))f_l(\cdot,\cdot,\cdot,l)||_{2,\mu \times \mathcal{P}_\Theta}^2 = \mathbb{E}_\mathcal{D}[\mathcal{L}_l(f_{l+1}; f_l)] - \mathbb{E}_\mathcal{D}[\mathcal{L}_l((\mathcal{T}_l(\cdot))f_l; f_l)]$$

$$= \mathbb{E}_\mathcal{D}[\mathcal{L}_l(f_{l+1}; f_l) - \mathcal{L}_l(g_f^*(l); f)] + \mathbb{E}_\mathcal{D}[\mathcal{L}_l(g_f^*(l); f) - \mathcal{L}_l((\mathcal{T}_l(\cdot))f_l; f_l)],$$

which, by the previous lemma and taking $\mathcal{G} = \mathcal{F}$ is, with probability at least $1 - \delta$, at most

$$\frac{56L^2 \ln\left(\frac{L|\mathcal{F}|^2}{\delta}\right)}{3m} + \sqrt{\frac{32L^2 \ln\left(\frac{L|\mathcal{F}|^2}{\delta}\right)}{m}} \epsilon_{\mathcal{F},\mathcal{F}} + \epsilon_{\mathcal{F},\mathcal{F}},$$

so, by repeating the above recursion, we have that

$$||f_{l+1}(\cdot,\cdot,\cdot,l+1) - Q^*(\cdot,\cdot,\cdot,l+1)||_{2,\nu \times \mathcal{P}_\Theta}$$

$$\leq (L+1) \sqrt{C \left(\frac{56L^2 \ln\left(\frac{L|\mathcal{F}|^2}{\delta}\right)}{3m} + \sqrt{\frac{32L^2 \ln\left(\frac{L|\mathcal{F}|^2}{\delta}\right)}{m}} \epsilon_{\mathcal{F},\mathcal{F}} + \epsilon_{\mathcal{F},\mathcal{F}} \right)},$$

for all l, with probability at least $1 - \delta$, which, by Lemma A.1.3 implies

$$v_{\pi^*}^L - v_{\pi_f}^L \leq L(L+1) \sqrt{C \left(\frac{56L^2 \ln\left(\frac{L|\mathcal{F}|^2}{\delta}\right)}{3m} + \sqrt{\frac{32L^2 \ln\left(\frac{L|\mathcal{F}|^2}{\delta}\right)}{m}} \epsilon_{\mathcal{F},\mathcal{F}} + \epsilon_{\mathcal{F},\mathcal{F}} \right)},$$

with probability at least $1 - \delta$. $\qquad \square$

The corollary regarding model-based learning immediately follows by constructing $\mathcal{F}$ from $\mathcal{H}$.