

PROVABLE BENEFITS OF REPRESENTATIONAL TRANSFER IN REINFORCEMENT LEARNING

Anonymous authors

Paper under double-blind review

ABSTRACT

We study the problem of representational transfer in RL, where an agent first pretrains in a number of source tasks to discover a shared representation, which is subsequently used to learn a good policy in a target task. We propose a new notion of task relatedness between source and target tasks, and develop a novel approach for representational transfer under this assumption. Concretely, we show that given a generative access to source tasks, we can discover a representation, using which subsequent linear RL techniques quickly converge to a near-optimal policy, with only online access to the target task. The sample complexity is close to knowing the ground truth features in the target task, and comparable to prior representation learning results in the source tasks. We complement our positive results with lower bounds without generative access, and validate our findings with empirical evaluation on rich observation MDPs that require deep exploration.

1 INTRODUCTION

Leveraging historical experiences acquired in learning past skills to accelerate the learning of a new skill is a hallmark of intelligent behavior. In this paper, we study this question in the context of reinforcement learning (RL). Specifically, we consider a setting where the learner is exposed to multiple tasks and ask the following question:

Can we accelerate RL by sharing representations across multiple related tasks?

There is rich empirical literature which studies multiple approaches to this question and various paradigms for instantiating it. For instance, in a multi-task learning scenario, the learner has simultaneous access to different tasks and tries to improve the sample complexity by sharing data across them (Caruana, 1997). Other works study a transfer learning setting, where the learner has access to multiple source tasks during a *pre-training* phase, followed by a target task (Pan and Yang, 2009). The goal is to learn features and/or a policy which can be quickly adapted to succeed in the target task. More generally, the paradigms of meta-learning (Finn et al., 2017), lifelong learning (Parisi et al., 2019) and curriculum learning (Bengio et al., 2009) also consider related questions.

On the theoretical side, questions of representation learning have received an increased recent emphasis owing to their practical significance, both in supervised learning and RL settings. In RL, a limited form of transfer learning across multiple downstream reward functions is enabled by several recent reward-free representation learning approaches (Jin et al., 2020a; Zhang et al., 2020; Wang et al., 2020; Du et al., 2019; Misra et al., 2020; Agarwal et al., 2020; Modi et al., 2021). Inspired by recent treatments of representation transfer in supervised (Maurer et al., 2016; Du et al., 2020) and imitation learning (Arora et al., 2020), some works also study more general task collections in bandits (Hu et al., 2021; Yang et al., 2020, 2022) and RL (Hu et al., 2021; Lu et al., 2021). Almost all these works study settings where the representation is *frozen* after pre-training in the source tasks, and a linear policy or optimal value function approximation is trained in the target task using these learned features. This setting, which we call *representational transfer*, is the main focus of our paper.

A crucial question in formalizing representational transfer settings is the notion of similarity between source and target tasks. Prior works in supervised learning make the stringent assumption that the covariates x follow the same underlying distribution in all the tasks, and only the conditional $P(y|x)$ can vary across tasks (Du et al., 2020). This assumption does not nicely generalize to RL settings, where state distributions are typically policy dependent, and prior attempts to extend this assumption to RL (Lu et al., 2021) result in strong assumptions on the learning setup. Other works (Hu et al.,

2021; Yang et al., 2020, 2022) focus on linear representations only, which limits the expressivity of the feature maps, and does not adequately represent the empirical literature in the field.

Our contributions. In this context, our work makes the following contributions:

- We propose a new linear span assumption of task relatedness for representational transfer, where the target task dynamics can be expressed as a (state-dependent) linear span of source task dynamics, in addition to the dynamics being low-rank under a shared representation. We give examples captured by this assumption, and it generalizes all prior settings for representational transfer in RL. We do not make any linearity assumptions on our feature maps.
- When we have *generative access* to source tasks, we provide a novel algorithm REPTRANSFER that successfully pretrains a representation for downstream *online learning in any target task* (i.e., no generative access in target task) satisfying the linear span assumption, when the source tasks satisfy a common latent reachability assumption. The regret bound of learning in the target task is close to that of learning in a linear MDP equipped with the ground truth features, the strongest possible yardstick in our setup. The additional terms in our regret largely arise out of the distributional mismatch between source and target tasks, which is expected. We complement the theory with an empirical validation of REPTRANSFER on the challenging rich observation combination lock benchmarks (Misra et al., 2020), confirming our theoretical findings.
- Without generative access to source tasks, we show the statistical hardness of representational transfer under the linear span assumption, and confirm this hardness in our empirical evaluation. We show that an additional assumption that every *observed state* is reachable in every source task is sufficient for allowing fully online learning in source tasks.

The new task relatedness assumption, reward-free learning result for low-rank MDPs and our analysis of LSVI-UCB under average case misspecification may be of independent interest.

2 RELATED WORK

In this section, we focus on survey related works that obtained concrete PAC or regret guarantees, and we defer a discussion of the empirical literature to the appendix.

Multi-task and Transfer Learning in Supervised Learning. The theoretical benefit of representation learning are well studied under conditions such as the i.i.d. task assumption (Maurer et al., 2016) and the diversity assumption (Du et al., 2020; Tripuraneni et al., 2020). Many works below successfully adopt the frameworks and assumptions to sequential decision making problems.

Multi-task and Transfer Learning in Bandit and small-size MDPs. Several recent works study multi-task linear bandits with linear representations ($\phi(s) = A s$ with unknown A) (Hu et al., 2021; Yang et al., 2020, 2022). The techniques developed in these works crucially rely on the linear structure and can not be applied to nonlinear function classes. Lazaric et al. (2013) study spectral techniques for online sequential transfer learning. Brunskill and Li (2013) study multi-task RL under a fixed distribution over finitely many MDPs, while Brunskill and Li (2014) consider transfer in semi-MDPs by learning options. Lecarpentier et al. (2021) consider *lifelong learning in Lipschitz MDP*. All these works consider small size tabular models while we focus on large-scale MDPs.

Multi-task and Transfer Learning in RL via representation learning. Beyond tabular MDPs, Arora et al. (2020) and D’Eramo et al. (2019) show benefits of representation learning in imitation learning and planning, but do not address exploration. Lu et al. (2021) study transfer learning in low-rank MDPs with general nonlinear representations, but make a generative model assumption on both the source tasks and the target task, along with other distributional and structural assumptions. We do not require generative access to the target task and make much weaker structural assumptions on the source-target relatedness. Recently and independently, Cheng et al. (2022) also studied transfer learning in low-rank MDPs in the online learning setting, identical to the setting we study in Section 5. However, their analysis relies on an additional assumption that bounds the point-wise TV error with the population TV error, which we show is in fact not necessary (details in Appendix C).

Efficient Representation Learning in RL. Even in the single task setting, efficient representation learning is an active area witnessing recent advances with exploration (Agarwal et al., 2020; Modi et al., 2021; Uehara et al., 2021; Zhang et al., 2022) or without (Ren et al., 2021). Other papers study feature selection (e.g. Farahmand and Szepesvári, 2011; Jiang et al., 2015; Pacchiano et al., 2020; Cutkosky et al., 2021; Lee et al., 2021; Zhang et al., 2021) or sparse models (Hao et al., 2021a,b).

3 PRELIMINARIES

Notations: We denote total variation distance of P_1 and P_2 by $\|P_1 - P_2\|_{TV}$. Given a vector a , we define $\|a\|_B = \sqrt{a^\top B a}$. c_0, c_1, \dots are universal constants. We use $a \leq Cb$ to mean $a \leq Cb$ for some universal constant C . Also, $[K] = \{1, \dots, K\}$ and $\lambda_{\min}(A)$ is the smallest eigenvalue of matrix A . Please see [Table 1](#) for a full list of notations and [Table 2](#) for a list of algorithms.

Low-rank Markov Decision Processes: We consider a finite-horizon episodic Markov Decision Process $\mathcal{M} = \langle H, \mathcal{S}, \mathcal{A}, \{P_h^?\}_{0:H-1}, \{r_h\}_{0:H-1}, d_0 \rangle$, specified by the episode length H , state space \mathcal{S} , discrete action space \mathcal{A} of size A , transition models $P_h^? : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$, reward functions $r_h : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$, and an initial distribution $d_0 \in \Delta(\mathcal{S})$. Occasionally, we use $P^?$ to denote $\{P_h^?\}_{0:H-1}$ and r to denote $\{r_h\}_{0:H-1}$. Starting from an initial state $s_0 \sim d_0$, an agent chooses actions a_h at step h upon observing the state s_h , receives a reward $r_h(s_h, a_h)$, and transitions to $s_{h+1} \sim P_h^?(s_h, a_h)$. We assume that r and d_0 are known. Given a policy $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$ and transition kernel P , we use the notation $\mathbb{E}_{\pi, P}[\cdot]$ to denote expectation under the distribution of trajectories when π is executed in P . The value function $V_{P,r,h}(s) = \mathbb{E}_{\pi, P} \left[\sum_{t=h}^H r(s_t, a_t) \mid s_h = s \right]$ gives the expected reward-to-go of π under the MDP with transitions P , rewards r , starting at state s in step h . Similarly, we define the Q function $Q_{P,r,h}(s, a) := r_h(s, a) + \mathbb{E}_{s' \sim P_h(s, a)} V_{P,r,h+1}(s')$. The expected total reward of a policy π under transitions P and rewards r is denoted as $V_{P,r} := \mathbb{E}_{s_0 \sim d_0} V_{P,r,0}(s_0)$. We define the state-action occupancy distribution $d_{P,h}(s, a)$ as the probability of π visiting (s, a) at time step h under π and P . $d_{P,h}(s)$ is the marginal state visitation, which is equal to $\sum_{a \in \mathcal{A}} d_{P,h}(s, a)$.

We study low-rank MDPs defined as follows ([Jiang et al., 2017](#); [Agarwal et al., 2020](#)). The conditions on the upper bounds of the norm of $\phi^?, \mu^?$ are just for normalization.

Definition 3.1 (Low-rank MDP) *A transition model $P_h^? : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ admits a low rank decomposition with rank $d \in \mathbb{N}$ if there exist two unknown embedding functions $\phi_h^? : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}^d$, $\mu_h^? : \mathcal{S} \mapsto \mathbb{R}^d$ such that $\forall s, s^0 \in \mathcal{S}, a \in \mathcal{A} : P_h^?(s^0 | s, a) = \mu_h^?(s^0)^\top \phi_h^?(s, a)$, where $\|\phi_h^?(s, a)\|_2 \leq 1$ for all (s, a) and for any function $g : \mathcal{S} \rightarrow [0, 1]$, $\int g(s) d\mu_h^?(s) \leq \sqrt{d}$. An MDP is a low rank MDP if $P_h^?$ admits such a low rank decomposition for all $h = 0, 1, \dots, H-1$.*

Low-rank MDPs capture the latent variable model ([Agarwal et al., 2020](#)) where $\phi^?(s, a)$ is a distribution over a discrete latent state space \mathcal{Z} , and the block-MDP model ([Du et al., 2019](#)) where $\phi^?(s, a)$ is a one-hot encoding vector. Note that the linear MDP model ([Yang and Wang, 2020](#); [Jin et al., 2020b](#)) assumes a *known* $\phi^?$, which significantly simplifies the algorithm design.

Transfer Learning: In contrast to the classic single-task learning setting, in this paper, we explore the setting of *transfer learning*, where learning consists of two phases: (1) the **pre-training phase** where the agent interacts with $K-1$ source tasks with dynamics $P_k^?$, and (2) the **deployment phase** where the agent is deployed into the K -th target task and no longer has access to the source tasks. The performance is measured mainly by the regret incurred in the target task upon deployment, while we also desire small sample complexity in the source tasks. We denote $d_{k,h} := d_{P_k^?,h}$.

In order for the pre-training phase to help with learning in the target task, we must make assumptions on the connections between tasks. In this work, we make the following fundamental structural assumption on all the tasks at hand, namely that they share the same underlying representation $\phi^?$.

Assumption 3.1 (Common representation) *We assume that all tasks $P_k^?$ are low-rank MDPs with a shared representation $\phi_h^?(s, a)$ but distinct $\mu_{k,h}^?(s^0)$, that is $P_{k,h}^?(s^0 | s, a) = \phi_h^?(s, a)^\top \mu_{k,h}^?(s^0)$.*

Assumption 3.2 (Realizability) *For any source task $k \in [K-1]$ and any $h \in [H]$, we assume that the agent has access to realizable function classes Φ_h and $\Upsilon_{k,h}$, such that $\phi_h^? \in \Phi_h$ and $\mu_{k,h}^? \in \Upsilon_{k,h}$. For normalization, we assume that for all k, h , all $\phi \in \Phi_h$ satisfy $\|\phi(s, a)\|_2 \leq 1$, and for all $\mu \in \Upsilon_{k,h}$ and any function $g : \mathcal{S} \rightarrow [0, 1]$, $\int g(s) d\mu_h^?(s) \leq \sqrt{d}$.*

[Assumption 3.2](#) is a standard realizability condition in the function class which is made in almost all prior works on RL with linear and nonlinear function approximation.

Under [Assumption 3.1](#), we expect the agent to learn a good representation during the pre-training phase by interacting with the source tasks, and then learn a good policy for the target task using this representation in the deployment phase. In order to realize this intuition though, we need two additional assumptions which we describe next.

Assumptions for representational transfer. In addition to standard assumptions in the source tasks, we make the following structural and relatedness assumptions on the source and target tasks.

Assumption 3.3 (Feature reachability in the source tasks) We assume that $\gamma := \min_{k \in [K-1], h \in [H]} \max_{\mathbf{s}^0} \min_{\mathbf{a}_h} \mathbb{E}_{\mathbf{P}_{k,h}^?} [\gamma_h^?(s_h; \mathbf{a}_h)]$ is strictly positive.

Assumption 3.3 intuitively requires that no subspace is unreachable in the source tasks, as otherwise it's impossible to guarantee the quality of the learned representation in that subspace which may contain high rewards in the target task. Note that no reachability is required in the target task.

The next assumption quantifies the relatedness of the target task and the source tasks.

Assumption 3.4 (Relatedness: Point-wise Linear span) For any $h \in [H]$ and $\mathbf{s}^0 \in \mathcal{S}_{h+1}$, there is a vector $\mathbf{p}_{k,h}(\mathbf{s}^0) \in \mathbb{R}^{K-1}$ such that $\mathbf{P}_{k,h}^?(s^0) = \sum_{k=1}^{K-1} \mathbf{p}_{k,h}(\mathbf{s}^0) \mathbf{P}_{k,h}^?(s^0)$, with $\max_{\mathbf{s}^0} \sum_{k=1}^{K-1} \mathbf{p}_{k,h}(\mathbf{s}^0) = 1$ and $\mathbf{p}_{k,h}(\mathbf{s}^0) \geq 0$.

Assumption 3.4 ensures that if $(\mathbf{s}; \mathbf{a})$ is reachable from a $(\mathbf{s}; \mathbf{a})$ pair in the target task, then it must be reachable from the same pair in at least one of the source tasks. This intuitively is also necessary for transfer learning, as could be a high rewarding state in the target. A special case is the convex combination, i.e., for any $\mathbf{s}^0 \in \mathcal{S}_h$, $\mathbf{P}_{k,h}^?(s^0) = \mathbf{p}_k$ with $\mathbf{p}_k \geq 0$; $\sum_k \mathbf{p}_k = 1$, which implies $\max_{\mathbf{s}^0} \sum_k \mathbf{p}_k = 1$. On the other hand, if the target task largely focuses on observations quite rare under the source tasks, they can grow large, and this is unavoidable in a transfer learning setting.

We remark that unlike prior work (Du et al., 2020; Tripuraneni et al., 2020; Lu et al., 2021), we do not make any assumption on the data generating distribution in either the source or the target tasks. Instead, our approach is end-to-end, i.e., we collect our own data from scratch for representation learning by doing strategic exploration in the source tasks (see Section C for a detailed discussion). In fact, in Theorem D.1 we show that the above assumptions do not permit successful transfer in the supervised learning setting, thus establishing an interesting separation between supervised and reinforcement learning.

We conclude this section with a couple of examples where our assumptions are satisfied.

Example 3.1 (Mixture of source tasks) Perhaps the simplest example of our assumptions is where $\mathbf{P}_k^?(s^0; \mathbf{a}) = \sum_{k=1}^{K-1} \mathbf{p}_k \mathbf{P}_k^?(s^0; \mathbf{a})$ with the coefficients independent of \mathbf{s}^0 , $\mathbf{p}_k \geq 0$ and $\sum_{k=1}^{K-1} \mathbf{p}_k = 1$. Such mixtures of base models have been considered in several prior works (Modi et al., 2020; Ayoub et al., 2020). While prior works study the case of arbitrary known base models, we instead allow structured, unknown base models with a shared representation.

Example 3.2 (Block MDPs with shared latent dynamics) In this example, each MDP $\mathbf{P}_k^?$ is a Block MDP (Du et al., 2019) with a shared latent space and a shared decoder $\gamma: \mathcal{S} \rightarrow \mathcal{Z}$. In a block MDP, given state action pair $(\mathbf{s}; \mathbf{a})$, the decoder γ maps to a latent state \mathbf{z} , the next latent state is sampled from the latent transition $\mathbf{P}(\mathbf{z}; \mathbf{a})$, and the next state is generated from an emission distribution $\alpha(\mathbf{z}^0)$. Recall that $\alpha(\mathbf{z}^0) > 0$ at only one $\mathbf{z}^0 \in \mathcal{Z}$ for any $\mathbf{s}^0 \in \mathcal{S}$ for a block MDP. We assume that the latent transition model $\mathbf{P}(\mathbf{z}; \mathbf{a})$ is shared across all the tasks, but the emission process differs across the MDPs. For instance, in a typical navigation example used to motivate Block MDPs, the latent dynamics might correspond to navigating in a shared 2-D map, while emission distributions capture different wall colors or lighting conditions across multiple rooms. Then Assumption 3.3 requires that the agent can visit the entire 2-D map, while Assumption 3.4 requires that the color/lighting conditions of the target task resemble that of at least one source task. The coefficients \mathbf{p}_k for any \mathbf{s}^0 are non-zero on the source tasks which can generate that observation.

4 TRANSFER LEARNING WITH GENERATIVE ACCESS TO SOURCE TASKS

We first study a setting where we assume generative model access to the source tasks, while having only online access to the target task.

Assumption 4.1 (Generative access to the source tasks) We assume that we have access to generative models for the $K-1$ source tasks. Specifically, for any $\mathbf{P}_k^?$ with $k \in [K-1]$, we can query any $(\mathbf{s}_h; \mathbf{a}_h)$ pair, and the generative model will return a next state sample $\mathbf{P}_{k,h}^?(s_h; \mathbf{a}_h)$.

Having generative model access is not unrealistic, especially in applications where a high-quality simulation environment is available. Generative access also does not trivialize the challenge of efficient exploration in the source tasks, since there are a potentially infinite number of states and the ground truth representation is unknown. Prior works using generative access typically

Algorithm 1 Transfer learning with generative access (REPTRANSFER)

PRE-TRAINING PHASE

Input: exploratory policies $\{g_{k=1}^K\}$, size of cross-sampled dataset, failure probability δ .

1: for task pairs i, j , i.e. for all $i, j \in [K-1]$ s.t. $i \neq j$ do . cross sampling procedure

2: For each $h \in [H-1]$, sample dataset $\mathcal{D}_{ij;h}$ containing n i.i.d. $(s; a; s^0)$ tuples sampled as:

$$(s; a) \sim d_{i;h}^i; s \sim P_{j;h}^j(j; s; a); a \sim U(A); s^0 \sim P_{i;h}^i(j; s; a); \quad (1)$$

3: $\forall h \in [H-1]$, learn $b_h = \text{Multi-task REPLEARN}(\{\mathcal{D}_{kj;h}\}_{k \in [K-1]})$. (Algorithm 3)

DEPLOYMENT PHASE Additional Input: number of deployment episodes \bar{d}

1: Set $\epsilon = H^{-\frac{1}{d}} + dH^{-\frac{1}{d}} \log(dHT) \epsilon$.

2: Run LSVI-UCB $\{b_h\}_{h=0}^{H-1}; r = r_K; T; \epsilon$ in the target task $\mathcal{P}_K^?$ (Algorithm 6).

require either a known representation (so that one can perform D-optimal design to construct an exploratory state-action distribution (Agarwal et al., 2019)), or directly assume access to a diverse state-action distribution which provides coverage and from which one can sample. Neither such a diverse sampling distribution is given in our case. We also note that in Section 5 we will show that without any additional assumptions, generative access in source tasks is necessary.

Algorithm overview. Given the above setup, now we present our algorithm REPTRANSFER, detailed in Algorithm 1. REPTRANSFER takes a representation learning approach to the transfer learning problem and operates in two phases. During the pre-training phase, REPTRANSFER performs reward-free exploration in each of the source tasks to learn task-specific policies which satisfies for all $h = 0; 1; \dots; H-1$, that $E_{k;P_k^?} \int b_h(s; a) \int_{h^?}^? (s; a) > \min_{h^?} I$ for some $\min > 0$. We call such policies \min -exploratory and will use this definition in Lemma 4.1 and Theorem 4.1. In Section 4 we present one particular algorithm that finds such policies with a specific \min , but our main analysis below is modular to the choice of the reward-free exploration algorithm.

Given the exploratory policies $\{g_{k \in [K-1]}\}$, REPTRANSFER collects a joint dataset across the source tasks using these policies and cross-sampling across pairs of tasks and then learns a single representation b . In the deployment phase, REPTRANSFER runs optimistic least squares value iteration (Algorithm 6) using the learned representation in the target task.

The cross sampling procedure and learning a shared representation. We now describe the cross sampling procedure in detail. Given an exploratory policy for each source task, the next step is to sample fresh data under cross-sampling procedure using the generative model. Consider a particular $(k; j)$ pair of source environments. For each $h \in [H]$, we first sample $s_{h-1}; a_{h-1} \sim d_{k;h-1}^k$. Then, in the simulator of task j , we reset to $(s_{h-1}; a_{h-1})$ and perform a transition step s_h , i.e., $s_h \sim P_{j;h}^j(s_{h-1}; a_{h-1})$. Then we reset the simulator for task k to states s_h , sample a uniformly random action a_h and then perform another transition step s_{h+1} , i.e., $s_{h+1} \sim P_k^?(s_h; a_h)$. Such a procedure is only possible under the generative model setting. **Intuitively, cross-sampling ensures that our training data contains all possible states that can be encountered in the target task, and failure modes without this can be found in the discussion following Theorem 5.1.**

Given datasets, each of size n collected using the cross-sampling procedure, we perform a Maximum Likelihood Estimation (MLE)-based representation learning procedure (Algorithm 3) jointly on all source tasks, with the goal of finding a representation function that can predict the transition probability well across all source tasks. An existing MLE generalization analysis (Agarwal et al., 2020) guarantee that Multi-task REPLEARN (Algorithm 3) achieves the following total variation guarantee with probability at least $1 - \delta$:

$$\sum_{k=1}^{K-1} E_{k;h} \int b_h(s; a) > b_{k;h}(\cdot) \int_{h^?}^? (s; a) > \int_{k;h}^? (\cdot) \quad TV \quad N := O((\log(j_{i=j}) + K \log(j_{i \neq j})) = N): \quad (2)$$

¹Given any dataset $\{s; a; r; s^0\}$ feature ϕ , and reward r , LSVI learns a Q function backward, i.e., at step t via $W_h = \arg \min_w \sum_{s; a; s^0} (W^>(s; a) - \hat{V}_{h+1}(s^0))^2 + kwk^2$ and set $\hat{V}_h(s) = \max_a (r(s; a) + W_h^>(s; a)); \delta s$. UCB, short for Upper Confidence Bound, refers to an exploration bonus added to basic LSVI.

Algorithm 2 REWARDFREE

- 1: Input: MDP $P^?$ with online access, num. LSVI-UCB episodes $N_{\text{LSVI-UCB}}$, num. model-learning episodes $N_{\text{REWARDFREE}}$, failure probability δ .
 - 2: Learn model $\hat{P}_h = (b_h; b_h)_{h=0}^{H-1}$ by running REWARDFREE REP-UCB (Algorithm 4) in $P^?$ for $N_{\text{REWARDFREE}}$ episodes.
 - 3: Set $\epsilon = dH^{-1} \log(dHN_{\text{LSVI-UCB}})$.
 - 4: Return $\pi = \text{LSVI-UCB}(f, b_h)_{h=0}^{H-1}; r = 0; N_{\text{LSVI-UCB}}; \text{UNIFORMACTIONS} = \text{TRUE}$ by simulating in the learned model \hat{P} (Algorithm 6). Note this step requires no samples from $P^?$
-

where $\mathcal{D}_{k;h}$ denotes the generating distribution of inputted data sets, each with size N . For example, in REPTTRANSFER we have $\mathcal{D}_{k;h} = [\bigcup_{j=1}^K \mathcal{D}_{kj;h}]$ and $N = n(K-1)$.

Representational transfer. Next we transfer this guarantee to the target task via Assumption 3.4, a key insight of our work. Specifically, we show that under cross sampling and Assumption 3.4, \hat{P}_h also linearly approximates the true transition in the target task well, under the occupancy measure of policy. Remarkably, this holds before the agent has ever interacted with the target task.

Lemma 4.1 Suppose Assumption 3.4 and that for all source tasks $k \in [K-1]$ we have a \min -exploratory policy π_k . Then, for any $\epsilon \in (0, 1)$, learning features \hat{P}_h using the cross-sampling procedure of Algorithm 1 satisfies ϵ w.p. $1 - \delta$. Furthermore, for any $h = 0, 1, \dots, H-1$, there exist $e_h : S \rightarrow \mathbb{R}^d$ such that for any function $g : S \rightarrow [0, 1]$, $\|g(s) - e_h(s)\|_2 \leq \epsilon$ and

$$\sup_{P \in \mathcal{P}} \mathbb{E}_{\pi_k} \|b_h(s_h; a_h) - e_h(s_h)\|_2 \leq \epsilon_{k;h} \quad \epsilon_{k;h} := \frac{P}{jA_j} \frac{3}{\max} K \quad n = \min :$$

Lemma 4.1 implies that \hat{P}_h is a feature such that \hat{P}_h is an approximately linear MDP in \hat{P}_h . Learning in approximately linear MDPs has been studied in Jin et al. (2020b), but under a much stronger error bound rather than the average misspecification here. Our result for downstream task learning under such an average model misspecification case presents the strongest result for learning in an approximately linear MDP, and the result might be of independent interest.

Regret bound for REPTTRANSFER. Putting things together, we obtain the following regret bound for online learning in the target task using our learned representation.

Theorem 4.1 (Regret under generative source access) Suppose Assumptions 3.1-3.4 and 4.1, and suppose the input policies π_k are \min -exploratory. Then, for any $\epsilon \in (0, 1)$, w.p. $1 - \delta$, REPTTRANSFER when deployed in the target task has regret at most $H^2 d^{1.5} \frac{P}{T} \log(1/\epsilon)$, with at most Kn generative accesses per source task, with $O\left(\frac{1}{\min} A \frac{3}{\max} KT \log \frac{j}{j} + K \log j\right)$.

Remarkably, Theorem 4.1 shows that with the pre-trained features, we achieve the same regret bound on the target task to the setting of linear MDP with known transition (Jin et al., 2020b), up to the additional $\frac{P}{T}$ factor. The scaling factor only depends on itself and captures the hardness of transfer learning. For special cases such as convex combination, i.e., state-independent and $\epsilon \in (0, 1/K)$, then $\frac{P}{T} = 1$. In the worst-case, some dependence on the scales seems unavoidable as we can have a state s^0 such that $\pi_k(s^0) = 1$ and $\pi_1(s^0) = 1$ with $\pi_1(s^0) = 1$. This corresponds to a rarely observed state for the source task encountered often in the target, and our estimates of transitions involving this state can be highly unreliable if it is not seen in any other source, roughly scaling the error between target and source tasks as $\frac{1}{\pi_1(s^0)}$. Obtaining formal lower bounds that capture a matching dependence on structural properties is an interesting question for future research.

Reward-free exploration in the source environments. So far we have assumed that we have a reward-free exploration black box algorithm that can provide an exploratory policy for each source task. In this section, we give a detailed algorithm that achieves this goal.

Recall that our transfer learning algorithm, for source task k relies on an exploratory policy π_k such that the empirical covariance with respect to π_k is lower bounded. This ensures good exploration in the underlying ground truth feature space $P_k^?$. Algorithm 2 achieves this goal by first invoking the REWARDFREE REP-UCB algorithm (Algorithm 4) to learn an estimated linear MDP model \hat{P}_k for $P_k^?$. **The REWARDFREE REP-UCB algorithm is a reward-free generalization of the recent**

REP-UCB algorithm (Uehara et al., 2021) that perform reward-free exploration in low-rank MDPs and can be of independent interest. Subsequently, we perform reward-free exploration (e.g., using LSVI-UCB with zero reward) within each \mathcal{P}_k , which involves no further environment interactions. Lemma 4.2 (Reward-free Exploration) Fix any source task $k \in [K - 1]$. Suppose Assumptions 3.2 and 3.3. Then, for any $\beta \in (0, 1)$, w.p. $1 - \beta$, REWARDFREE (Algorithm 2) with $N_{\text{LSVI-UCB}} = \lceil e A^3 d^6 H^8 \beta^{-2} \rceil$ and $N_{\text{REWARDFREE}} = \lceil e A^3 d^4 H^6 \log \frac{1}{\beta} \rceil N_{\text{LSVI-UCB}}^2$ returns a β -exploratory policy π_k where $\beta = e A^3 d^5 H^7 \beta^{-2}$. The sample complexity here is $N_{\text{REWARDFREE}}$ episodes in the source task.

To the best of our knowledge, Lemma 4.2 is the first result that finds a full-rank policy cover in the low-rank MDP setting, and might be of independent interest. Wagenmaker et al. (2022) recently obtained a related guarantee in the linear MDP setting with known features. Incorporating Lemma 4.2 and Theorem 4.1, we get our final sample complexity bound.

Theorem 4.2 (Regret for REPTTRANSFER with REWARDFREE subroutine) Suppose Assumptions 3.1-3.4 and 4.1 hold, and let $\beta \in (0, 1)$. Let $\{g_{k=1}^K\}$ be exploratory policies learned from running REWARDFREE on each source task with $N_{\text{LSVI-UCB}}$ and $N_{\text{REWARDFREE}}$ set as in Lemma 4.2. Then, w.p. $1 - \beta$, running REPTTRANSFER with $\{g_{k=1}^K\}$ has regret in the target task of $\lceil e H^2 d^{1.5} \beta^{-1} T \log(1/\beta) \rceil$, with at most $\lceil e A^4 \beta_{\max}^3 d^5 H^7 K^2 T^{-2} (\log(j) + K \log j) \rceil$ generative accesses per source task.

5 TRANSFER LEARNING WITH ONLINE ACCESS TO SOURCE TASKS

In the previous section, we show that efficient transfer learning is possible under very weak structural assumptions, with generative access to the source tasks. One natural question is whether transfer learning is possible with only online access to the source tasks. Somewhat surprisingly, we show that this is impossible without significantly stronger assumption.

Theorem 5.1 (Impossibility Result) Let \mathcal{M}_K be a K -task multi-set that satisfies (1) all tasks are Block MDPs; (2) all tasks satisfy Assumption 3.3 and Assumption 3.4; (3) the latent dynamics are exactly the same for all source and target tasks. For any pre-training algorithm which outputs a feature $\hat{\phi}$ by interacting with the source tasks $\mathcal{S} \subseteq [K - 1]$, there exists $\mathcal{P}_K \subseteq \mathcal{M}_K$, such that with probability at least $1 - \beta$, \mathcal{A} will output a feature $\hat{\phi}$, such that for any policy taking the functional form of $\pi(s) = \text{argmax}_a \sum_{z \in \mathcal{Z}} \hat{\phi}(s; a) g_{z,a}$, we have $\mathbb{E} \sum_{k \in \mathcal{S}} V_k - V_K \geq 1 - \beta$.

The above theorem implies that a representation learned only from online access to source tasks does not enable learning in downstream tasks if the downstream task algorithm is restricted to use the representation as the only information of the state-action pairs (e.g., running UCB with $\hat{\phi}$).

We briefly explain the intuition behind the above lower bound. For a Block MDP, for any (s, a) , we can model the ground-truth $g_{z,a}$ as a one-hot encoding $\mathbb{1}_{z=(z,a)}$ corresponding to the latent state-action pair $(z; a)$ with $z = \phi(s)$ being the encoded latent state. The key observation here is that any permutation of ϕ will also be a perfect feature in terms of characterizing the Block MDPs, since it corresponds to simply permuting the indices of the latent states. Therefore, without cross referencing, the agent could potentially learn different permutations in different source tasks, which would collapse in the target task. A precise constructive proof of Theorem 5.1 can be found in Section D.

Part of the reason that the above example fails is that each source task has its own observed subset of raw states, which permits such permutation to happen. In what follows, we show that, indeed, under an additional assumption on the reachability of raw states, a slight variant of the same algorithm (Algorithm 5 in Section G) can achieve the same regret with only online access to the source tasks.

Assumption 5.1 (Reachability in the raw state space) For all source tasks $k \in [K - 1]$, any policy π and $h \in \{0, 1, \dots, H - 1\}$, we have $\inf_{s \in \mathcal{S}; a \in \mathcal{A}} \mathbb{P}_{\pi}(\phi_h(s; a) = \text{raw}_k) \geq \beta$.

Assumption 5.1 implies that for each source task, any policy that achieves a full-rank covariance matrix also achieves global coverage across the raw state-action space. In addition, in order to apply importance sampling (IS) to transfer the TV error from source task to target task, we need to assume that the target task distribution has bounded density. This is true, for example, when \mathcal{S} is discrete.

	Source	O-REPTRANSFER	G-REPTRANSFER	Oracle	Target
Comblock	1	8006 (294.7)	7790 (267.6)	7048 (164.8)	181450 (147600.2)
Comblock PO	1	1	7764 (145.7)	7336 (270.7)	85000 (14469.8)

(a)

(b)

Figure 1:(a): A visualization of the rich observation comblock environment. Latent states (white and black) emit continuous high-dimensional observations. The reward is sparse (only in white states at H). Each white state has 10 actions where one of them is a good action that leads to the next two white states while the other 9 lead to black states (good action for each white state is different). Once the agent is in black states, it stays stuck in black until the end of the episode. Thus, a random exploration strategy has an exponentially small probability of hitting the goal (i.e., $\frac{1}{(dH)^H}$). (b) Top: Number of episodes required to solve the target environment under the setting from Sec. 6. Bottom: Number of episodes required to solve the target environment under the setting from Sec. 6. An algorithm solves the target task if it can achieve the optimal return (i.e., 1) for 5 consecutive iterations with 50 evaluation runs each. We include the mean and standard deviation (in the brackets) for 5 random seeds. \dashv denotes that an algorithm can not solve the target task within a fixed sample budget.

Assumption 5.2 (Bounded density). For all $(\delta; h; s; a)$, we have $d_{K;h}(s; a) \leq 1$.

Theorem 5.2 (Regret with online access). Suppose Assumptions 3.1-3.4, 5.1, 5.2 hold. W.p. $1 - \delta$, Algorithm 5 with appropriate parameters achieves a regret in the target $\leq d^{1.5} H^2 \frac{p}{T} \log(1/\delta)$, with at most $\text{poly}(A; \max\{d; H; K; T; \frac{1}{\delta}\}; \frac{1}{\text{raw}}; \log(\prod_{j=1}^J j))$ online queries in the source tasks.

Assumption 5.1 is satisfied in a Block MDP, when, for example, the emission function $(s; z)$ satisfies that $\exists s; z; s.t., \exists (s; z) \in \mathcal{C}$. That is, for any source task, any state in the state space can be generated by at least one latent state.

6 EXPERIMENTS

We investigate empirically the benefit of transfer learning under the Block MDP setting, on the challenging Rich Observation Combination Lock (comblock) benchmark. In this environment, one must recover the correct feature from rich observations and perform strategic exploration or otherwise pay for exponential sample complexity. We include a visual overview in Fig. 1(a). This environment is uniquely challenging since it requires strategic exploration and latent state discovery at the same time, which results in failures of common deep RL methods (Misra et al., 2020), and theoretical RL approaches based on linear function approximation (and kernels) (Zhang et al., 2022).

In particular, in this section we study the following questions: i) what are the benefits of representational transfer using multiple source tasks, and ii) whether generative access to source tasks is needed. We design two sets of experiments with various source and target environment configurations. We defer the details of the experiments (the design of vanilla comblock, hyperparameters, etc.) in Appendix J and include the essential design information in the following two sections.

Baselines. We denote \mathcal{C} as the smallest sample complexity of ϵ -VI-UCB using learned features from any of the source tasks; \mathcal{G} -REPTRANSFER as REPTRANSFER with only online access to the source tasks; \mathcal{G} -REPTRANSFER as REPTRANSFER with generative access to the source tasks; Oracle as learning in the target task with ground truth features; Target as running BTRIEE (Zhang et al., 2022) — the SOTA Block MDP algorithm, in the target task with no pretraining.

Comblock without partitioned observations. In this section we start off with an easier setting: we use 5 source tasks, each with horizon $t_{\text{hor}} = 25$, 3 latent states per level and 10 actions. The latent transition dynamics are different for each source task, but notice that for comblock, the reachability assumption (c.r. Assumption 3.3) is always satisfied. The emission distribution of all the source and target tasks is identical, so Assumption 5.1 holds here. For the construction of the target task, for each timestep t we choose the latent transition dynamics from one of the sources uniformly at random (thus Assumption 3.4 is satisfied). We record the number of episodes in the

(a) (b)

Figure 2: (a): Visualization of the decoder source (top) and REPTTRANSFER (bottom). (b): Visualization of the decoder G-REPTTRANSFER (top) and GREPTTRANSFER (bottom). For each baseline, the i -th column in the i -th image denotes the averaged decoded states from the 30 observations generated by latent state, for $i \in \{0, 1, 2\}$ and $h \in \{1, 2\}$ [25], from the corresponding target environment. The optimal decoder should recover the latent states up to a permutation. In Fig. 2(a) (top), note that the learned features in source task fail to solve the target because of the collapse at timestep 5: both observations from state 0 and 1 are mapped to state 0. Note in the source task where this feature is trained, such collapse can happen when states have identical latent transition (for detailed discussion we refer to Misra et al. (2020)). In Fig. 2(b) (top), REPTTRANSFER with only online access learns an incorrect decoder when the source tasks' observation spaces are disjoint. This is because the learned feature can decode each source task with a different permutation.

target environment that each method takes to solve the target environment in Table. 1(b). We first observe that REPTTRANSFER with either online or generative access can solve the target task (since Assumption 5.1 holds). Second, we observe that directly applying the learned feature from a single source task does not suffice to solve the target environment. This is because the representation learned from a single source task may collapse two latent states into a single one during encoding (e.g., if two latent states at the same time step have exactly identical latent transitions). See the visualization of comparisons between the learned decoders from a single source task and REPTTRANSFER in Fig. 2(a). Third, the result shows that REPTTRANSFER saves order of magnitude of target samples compared with training in the target environment from scratch using the SOTA Block MDPs algorithm BRIEE. This set of results verifies the empirical benefits of representation learning from multiple tasks, i.e., resolves ambiguity and speeds up downstream task learning.

Comblock with partitioned observation space (Comblock-PO) In this section, following the intuition of our lower bound (Theorem 5.1), we construct a setting where the supports of the emission distributions from each task are completely disjoint, while the emission distribution in the target task is a mixture of all source emissions and the latent dynamics are identical across tasks. Hence Assumption 3.4 holds while Assumption 5.1 fails. So we expect that an algorithm without generative access to source tasks will fail based on Theorem 5.1. Specifically, under disjoint emission supports, a representation can decode the latent state correctly for each source, but permute latent state labels across sources, causing decoding errors on the target. We record the number of target episodes for each method to solve the target task in Table. 1(b). We observe that indeed the online version fails while the generative version still succeeds. We show the visualizations of REPTTRANSFER and G-REPTTRANSFER in Fig. 2(b), and we note that the empirical results exactly match our theoretical results. This ablation verifies that source generative model access is needed.

7 CONCLUSION

We study representational transfer among low rank MDPs which share the same unknown representation. Under a reasonably flexible linear span task relatedness assumption, we propose an algorithm that provably transfers the representation learned from source tasks to the target task. The regret in target task matches the bound obtained with oracle access to the true representation, using only polynomial number of samples from source tasks. Our approach relies on the generative model access in source tasks, which we prove is not avoidable in the worst case under the linear span assumption. To complement the lower bound, we propose a stronger assumption on the conditions of the reachability in raw states, under which online access to source tasks suffices for provably efficient representation transfer. Finding modalities other than generative access which avoid the lower bound, and a more extensive empirical evaluation beyond the proof-of-concept experiments here are important directions for future research.

Reproducibility Statement: For theory, we include detailed and complete proofs for all of our claims in the appendices. For empirics, we include anonymized code in the supplement. Please refer to Appendix J for a comprehensive experimental setup and list of hyperparameters. All experiments were run on CPUs and no external datasets were used.

REFERENCES

- Alekh Agarwal, Nan Jiang, Sham M Kakade, and Wen Sun. Reinforcement learning: Theory and algorithms. CS Dept., UW Seattle, Seattle, WA, USA, Tech, Rep.
- Alekh Agarwal, Sham Kakade, Akshay Krishnamurthy, and Wen Sun. Flambe: Structural complexity and representation learning of low rank mdps. *Advances in Neural Information Processing Systems*, volume 33, pages 20095–20107, 2020.
- Sanjeev Arora, Simon Du, Sham Kakade, Yuping Luo, and Nikunj Saunshi. Provable representation learning for imitation learning via bi-level optimization. *International Conference on Machine Learning* pages 367–376. PMLR, 2020.
- Alex Ayoub, Zeyu Jia, Csaba Szepesvari, Mengdi Wang, and Lin Yang. Model-based reinforcement learning with value-targeted regression. *International Conference on Machine Learning* pages 463–474. PMLR, 2020.
- André Barreto, Will Dabney, Rémi Munos, Jonathan J Hunt, Tom Schaul, Hado P van Hasselt, and David Silver. Successor features for transfer in reinforcement learning. *Advances in neural information processing systems*, 30, 2017.
- Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning* pages 41–48, 2009.
- Léonard Blier, Corentin Tallec, and Yann Ollivier. Learning successor states and goal-dependent values: A mathematical viewpoint. *arXiv preprint arXiv:2101.07123*, 2021.
- Emma Brunskill and Lihong Li. Sample complexity of multi-task reinforcement learning. *arXiv preprint arXiv:1309.6821*, 2013.
- Emma Brunskill and Lihong Li. Pac-inspired option discovery in lifelong reinforcement learning. In *International conference on machine learning* pages 316–324. PMLR, 2014.
- Rich Caruana. Multitask learning. *Machine learning* 28(1):41–75, 1997.
- Yuan Cheng, Songtao Feng, Jing Yang, Hong Zhang, and Yingbin Liang. Provable benefit of multitask representation learning in reinforcement learning. *arXiv preprint arXiv:2206.05900*, 2022.
- Ashok Cutkosky, Christoph Dann, Abhimanyu Das, Claudio Gentile, Aldo Pacchiano, and Manish Purohit. Dynamic balancing for model selection in bandits and reinforcement learning. *International Conference on Machine Learning* pages 2276–2285. PMLR, 2021.
- Peter Dayan. Improving generalization for temporal difference learning: The successor representation. *Neural computation* 5(4):613–624, 1993.
- Carlo D'Eramo, Davide Tateo, Andrea Bonarini, Marcello Restelli, and Jan Peters. Sharing knowledge in multi-task deep reinforcement learning. *International Conference on Learning Representations*, 2019.
- Simon Du, Akshay Krishnamurthy, Nan Jiang, Alekh Agarwal, Miroslav Dudík, and John Langford. Provably efficient rl with rich observations via latent state decoding. *International Conference on Machine Learning* pages 1665–1674. PMLR, 2019.
- Simon S Du, Wei Hu, Sham M Kakade, Jason D Lee, and Qi Lei. Few-shot learning via learning the representation, provably. *arXiv preprint arXiv:2002.09434*, 2020.
- Amir-massoud Farahmand and Csaba Szepesvári. Model selection in reinforcement learning. *Machine learning* 85(3):299–332, 2011.

- Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. *International conference on machine learning*, pages 1126–1135. PMLR, 2017.
- Botao Hao, Yaqi Duan, Tor Lattimore, Csaba Szepesvári, and Mengdi Wang. Sparse feature selection makes batch reinforcement learning more sample efficient. *International Conference on Machine Learning*, pages 4063–4073. PMLR, 2021a.
- Botao Hao, Tor Lattimore, Csaba Szepesvári, and Mengdi Wang. Online sparse reinforcement learning. *International Conference on Artificial Intelligence and Statistics*, pages 316–324. PMLR, 2021b.
- Jiachen Hu, Xiaoyu Chen, Chi Jin, Lihong Li, and Liwei Wang. Near-optimal representation learning for linear bandits and linear rl. *International Conference on Machine Learning*, pages 4349–4358. PMLR, 2021.
- Nan Jiang, Alex Kulesza, and Satinder Singh. Abstraction selection in model-based reinforcement learning. *International Conference on Machine Learning*, pages 179–188. PMLR, 2015.
- Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, John Langford, and Robert E Schapire. Contextual decision processes with low bellman rank are pac-learnable. *International Conference on Machine Learning*, pages 1704–1713. PMLR, 2017.
- Chi Jin, Akshay Krishnamurthy, Max Simchowitz, and Tiancheng Yu. Reward-free exploration for reinforcement learning. *International Conference on Machine Learning*, pages 4870–4879. PMLR, 2020a.
- Chi Jin, Zhuoran Yang, Zhaoran Wang, and Michael I Jordan. Provably efficient reinforcement learning with linear function approximation. *Conference on Learning Theory*, pages 2137–2143. PMLR, 2020b.
- Alessandro Lazaric, Emma Brunskill, et al. Sequential transfer in multi-armed bandit with finite set of models. *Advances in Neural Information Processing Systems*, 26, 2013.
- Erwan Lecarpentier, David Abel, Kavosh Asadi, Yuu Jinnai, Emmanuel Rachelson, and Michael L Littman. Lipschitz lifelong reinforcement learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 8270–8278, 2021.
- Jonathan Lee, Aldo Pacchiano, Vidya Muthukumar, Weihao Kong, and Emma Brunskill. Online model selection for reinforcement learning with function approximation. *International Conference on Artificial Intelligence and Statistics*, pages 3340–3348. PMLR, 2021.
- Rui Lu, Gao Huang, and Simon S Du. On the power of multitask representation learning in linear mdp. *arXiv preprint arXiv:2106.08053*, 2021.
- Thodoris Lykouris, Max Simchowitz, Alex Slivkins, and Wen Sun. Corruption-robust exploration in episodic reinforcement learning. *Conference on Learning Theory*, pages 3242–3245. PMLR, 2021.
- Andreas Maurer, Massimiliano Pontil, and Bernardino Romera-Paredes. The benefit of multitask representation learning. *Journal of Machine Learning Research*, 17(81):1–32, 2016.
- Dipendra Misra, Mikael Henaff, Akshay Krishnamurthy, and John Langford. Kinematic state abstraction and provably efficient rich-observation reinforcement learning. *International conference on machine learning*, pages 6961–6971. PMLR, 2020.
- Aditya Modi, Nan Jiang, Ambuj Tewari, and Satinder Singh. Sample complexity of reinforcement learning using linearly combined model ensembles. *International Conference on Artificial Intelligence and Statistics*, pages 2010–2020. PMLR, 2020.
- Aditya Modi, Jinglin Chen, Akshay Krishnamurthy, Nan Jiang, and Alekh Agarwal. Model-free representation learning and exploration in low-rank mdp. *arXiv preprint arXiv:2102.07035*, 2021.
- Aldo Pacchiano, Christoph Dann, Claudio Gentile, and Peter Bartlett. Regret bound balancing and elimination for model selection in bandits and rl. *arXiv preprint arXiv:2012.13045*, 2020.

- Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering* 22(10):1345–1359, 2009.
- German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural Networks* 113:54–71, 2019.
- Emilio Parisotto, Jimmy Lei Ba, and Ruslan Salakhutdinov. Actor-mimic: Deep multitask and transfer reinforcement learning. *arXiv preprint arXiv:1511.06342*, 2015.
- Tongzheng Ren, Tianjun Zhang, Csaba Szepesvári, and Bo Dai. A free lunch from the noise: Provable and practical exploration for representation learning. *arXiv preprint arXiv:2111.11485*, 2021.
- Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *arXiv preprint arXiv:1606.04671*, 2016.
- Ahmed Touati and Yann Ollivier. Learning one representation to optimize all rewards. *Advances in Neural Information Processing Systems* 34:13–23, 2021.
- Nilesh Tripuraneni, Michael Jordan, and Chi Jin. On the theory of transfer learning: The importance of task diversity. *Advances in Neural Information Processing Systems* 33:7852–7862, 2020.
- Masatoshi Uehara, Xuezhou Zhang, and Wen Sun. Representation learning for online and of ine rl in low-rank mdps. In *International Conference on Learning Representations*, 2021.
- Andrew Wagenmaker, Yifang Chen, Max Simchowitz, Simon S Du, and Kevin Jamieson. Reward-free rl is no harder than reward-aware rl in linear markov decision processes. *arXiv preprint arXiv:2201.11206*, 2022.
- Ruosong Wang, Simon S Du, Lin Yang, and Russ R Salakhutdinov. On reward-free reinforcement learning with linear function approximation. *Advances in neural information processing systems* 33:17816–17826, 2020.
- Jiaqi Yang, Wei Hu, Jason D Lee, and Simon Shaolei Du. Impact of representation learning in linear bandits. In *International Conference on Learning Representations*, 2020.
- Jiaqi Yang, Qi Lei, Jason D Lee, and Simon S Du. Nearly minimax algorithms for linear bandits with shared representation. *arXiv preprint arXiv:2203.15664*, 2022.
- Lin Yang and Mengdi Wang. Reinforcement learning in feature space: Matrix bandit, kernels, and regret bound. In *Proceedings of the 37th International Conference on Machine Learning*, pages 10746–10756, 2020.
- Andrea Zanette, Ching-An Cheng, and Alekh Agarwal. Cautiously optimistic policy optimization and exploration with linear function approximation. *COLT*, 2021.
- Weitong Zhang, Jiafan He, Dongruo Zhou, Amy Zhang, and Quanquan Gu. Provably efficient representation learning in low-rank markov decision processes. *arXiv preprint arXiv:2106.11935*, 2021.
- Xuezhou Zhang, Yuzhe Ma, and Adish Singla. Task-agnostic exploration in reinforcement learning. *Advances in Neural Information Processing Systems* 33:11734–11743, 2020.
- Xuezhou Zhang, Yuda Song, Masatoshi Uehara, Mengdi Wang, Wen Sun, and Alekh Agarwal. Efficient reinforcement learning in block mdps: A model-free representation learning approach. *arXiv preprint arXiv:2202.00063*, 2022.
- Zhuangdi Zhu, Kaixiang Lin, and Jiayu Zhou. Transfer learning in deep reinforcement learning: A survey. *arXiv preprint arXiv:2009.07888*, 2020.

Appendices

A NOTATIONS

Table 1: List of Notations

$S; A; A$	State and action spaces, and $ A $.
(S)	The set of distributions supported by S .
$\lambda_{\min}(A)$	Smallest eigenvalue of matrix A .
e_j	One-hot encoding of j , i.e. 0 at each index except the one corresponding to j .
$(x)_y$	Length of vector implied from context.
$\min_x y$	$\min_x x; y$.
H	Episode length of MDPs, a.k.a. time horizon. We index steps as $0; 1; \dots; H - 1$.
K	There are a total of K tasks. The $k = 1; \dots; K - 1$ are source tasks, and the K -th task is the target task.
d	dimension of the low-rank MDP, i.e. dimension of \mathcal{F} .
$P_{k;h}^*$	Ground truth transition at time h for task k .
r_h	Reward function of the target task (i.e. task K).
$E_{\mathcal{P}}[\cdot]$	Expectation under the distribution of trajectories when executed in \mathcal{P} . We sometimes omit \mathcal{P} when the MDP is clear from context.
$d_{\mathcal{P};h}$	Occupancy distribution of \mathcal{P} under transition P at time h .
$d_{k;h}$	Occupancy distribution for the k -th task, i.e. $d_{\mathcal{P}_k;h}$.
$\phi_h^*(s; a)$	Embedding function $f(s; a)$ at time h .
\mathcal{F}_h	Realizable function class for ϕ_h^* .
$j \cdot j$	Defined as $\max_h j \cdot j_h$.
$\phi_{k;h}^*(s^0)$	Emission embedding function $f(s^0)$ at time h for environment k .
$\mathcal{F}_{k;h}$	Realizable function class for $\phi_{k;h}^*$.
$j \cdot j$	Defined as $\max_{k;h} j \cdot j_{k;h}$.
\max_i	Feature reachability in the source task (Assumption 3.3).
\max_i	Constants defined in point-wise linear span assumption (Assumption 3.4).
raw	Raw states reachability parameter (Assumption 5.1).

Table 2: List of Algorithms

REPTRANSFER (Algorithm 1)	Transfer learning with generative access.
REWARDFREE (Algorithm 2)	Reward Free exploration in low-rank MDP.
Multi-task REPLEARN (Algorithm 3)	Multi-task Maximum Likelihood (MLE) Representation Learning.
REWARDFREE REP-UCB (Algorithm 4)	Modified REP-UCB for reward-free model learning.
REPTRANSFER for online access (Algorithm 5)	Simplified version of REPTRANSFER if we have reachability under raw states (see Theorem 5.2).
LSVI-UCB (Algorithm 6)	Optimistic Least Squares Value Iteration (with bonus).

B ADDITIONAL RELATED WORKS

The idea of learning transferable representation has been extensively explored in the empirical literature. Here we don't intend to provide a comprehensive survey of all existing works on this topic. Instead, we discuss a few representative approach that may be of interest.

Towards transfer learning across different environments, progressive neural network (Rusu et al., 2016) is among the first neural-based attempt to learning a transferable representation for a sequence of downstream tasks that tries to overcome the challenge of catastrophic forgetting. It maintains the learned neural models for all previous tasks and introduce additional connections between the network of the current tasks to those of prior tasks to allow information reuse. However, a drawback common to such an approach is that the network size grows linearly with the number of tasks. Other approaches include directly learning a multi-task policy that can perform well on a set of source tasks, with the hope that it will generalize to future tasks (Parisotto et al., 2015). Such an approach requires the tasks to be similar in their optimal policy, which is a much stronger assumption than ours.

Slightly off-topic are the works about "transfer learning" inside the same environment but across different reward functions, which is more restricted than the setting considered in this paper. Several prior works design representation learning algorithms that aim to learn a representation that generalize across multiple reward function/goals (Dayan, 1993; Barreto et al., 2017; Touati and Ollivier, 2021; Blier et al., 2021). These are related to REWARDFREE REP-UCB we developed in Section E.

The key difference is that we concern representation learning along with efficient exploration to derive an end-to-end polynomial sample complexity bound. These prior works do not consider exploration and do not come with provable sample complexity bounds. We refer interested readers to a recent survey (Zhu et al., 2020) for a comprehensive discussion of other empirical approaches.

C COMPARISONS TO CLOSELY RELATED WORKS

Recall that in addition to the commonly made shared representation assumption (Assumption 3.1), we made two additional structural assumptions: reachability (Assumption 3.3) and linear span on \mathcal{X} (Assumption 3.4). The reachability assumption is commonly made in prior works even in single agent RL, e.g. (Modi et al., 2021; Misra et al., 2020). It ensures that there is no redundant dimensions in the ground-truth representation, which is a reasonable requirement. The linear span assumption is closely related to the diversity assumptions made in prior works of transfer learning in both supervised learning (Tripuraneni et al., 2020) and reinforcement learning (Lu et al., 2021).

Lu et al. (2021): In the prior work of Lu et al. (2021), which also studies transfer learning in low-rank MDPs with nonlinear function approximations, they need to make the following assumptions:

1. shared representation (identical to our Assumption 3.1).
2. task diversity (similar to our Assumption 3.4).
3. generative model access to both the source and the target tasks. In contrast, we only require generative model access to the source tasks and allow online learning in the target task.
4. a somewhat strong coverage assumption saying that the data covariance matrix (under the generative data distribution) between arbitrary pairs of features^{0,2} must be full rank. In contrast, our analysis only requires coverage in the true features^{0,2} in the source tasks.
5. the existence of an ideal distribution q on which the learned representation can extrapolate. We do not require an assumption of a similar nature. Instead, we show that the data collected from our strategic reward-free exploration phase suffices for successful transfer.
6. the uniqueness for each in the sense of linear-transform equivalence. Two representation functions ϕ and ϕ_0 can yield similar estimation result if and only if they differ by just an invertible linear transformation. In contrast, we do not make any additional structural assumptions on the function class beyond realizability.

In summary, our work present a theoretical framework that permits successful representation transfer based on significantly weaker assumptions. We believe that this is a solid step towards understanding transfer learning in RL.

Cheng et al. (2022): A concurrent work of Cheng et al. (2022) also studies the exact same problem as ours. Both works study the setting where the agent performs reward-free exploration in the source tasks for representation learning and used the learned representation for the target task. Both works achieves similar sample complexity in source and target tasks, albeit using very different algorithms. However, in addition to the assumptions that we made Cheng et al. (2022) has to make an additional assumption.

Assumption C.1 For any two different models in the model class \mathcal{H} , say $P^1(s^0; \mathbf{s}; \mathbf{a}) = h^1(\mathbf{s}; \mathbf{a}; \mathbf{s}^0)$ and $P^2(s^0; \mathbf{s}; \mathbf{a}) = h^2(\mathbf{s}; \mathbf{a}; \mathbf{s}^0)$, there exists a constant C_R such that for all $(\mathbf{s}; \mathbf{a}) \in \mathcal{S} \times \mathcal{A}$ and $h \in \mathcal{H}$,

$$\|P^1(\mathbf{j}; \mathbf{s}; \mathbf{a}) - P^2(\mathbf{j}; \mathbf{s}; \mathbf{a})\|_{TV} \leq C_R E_{(\mathbf{s}; \mathbf{a}) \sim U(\mathcal{S}; \mathcal{A})} \|P^1(\mathbf{j}; \mathbf{s}; \mathbf{a}) - P^2(\mathbf{j}; \mathbf{s}; \mathbf{a})\|_{TV}$$

This assumption ensures that the point-wise TV error is bounded, as long as the population-level TV error is bounded. This assumption is used to transfer the MLE error from the source tasks to the target task. This type of assumption is strong in the sense that we typically expect scale with $|\mathcal{S}|$. In contrast, our analysis (Lemma G.1) shows that this assumption is in fact not necessary assuming online access only to source tasks. The generative access to source task studied here, which enables transfer under weaker reachability assumptions is not studied in their work.

It is worth noting that Cheng et al. (2022) also study of RL in the target task which we do not cover, while we mainly focus on the setting of generative models in the source tasks and demonstrating a more complete picture by proving generative model access in source tasks is needed without additional assumptions. Comparing to (Cheng et al., 2022), we also further implement and perform experimental evaluations of our algorithm.

D IMPOSSIBILITY RESULTS

Here, we present an interesting result showing that the above assumptions we make are so weak that they do not even permit efficient transfer in supervised learning:

Theorem D.1 (Counter-example in supervised learning) Assume that we want to perform conditional density estimation, where $P_k(y|x) = \frac{P_k(x,y)}{P_k(x)}$. Under Assumption 3.1 (shared representation) and Assumption 3.4 (linear span), and assume that in each source task, one have access to a data generating distribution $P_k(x)$ such that $\min_k E_k[\frac{P_k(x)}{P_k(x)}]$ (reachability). No algorithm can consistently achieve $E_k[\|P_k(y|x) - \hat{P}_k(y|x)\|_{TV}] \leq \epsilon$ on the target task using the feature learned from the source tasks with probability more than $1-\epsilon$.

Proof of Theorem D.1 Consider the following example $\mathcal{X} = \mathbb{R}^2$ and we have the following 3 sets.

$$S_1 = B_{1/2}((1; 1))$$

$$S_2 = B_{1/2}((2; 2))$$

$$S_3 = B_{1/2}((0; 1))$$

where $B_a((x; y))$ stands for the ball with radius a centered at $(x; y)$. These will be the support of 3 tasks: task 1 and 2 are two source tasks, task 3 is the target task. Let's assume that P_k are uniform distribution on S_k .

Suppose that the feature class only contains two functions:

$$\begin{aligned} f_1 &: f(x_1, x_2) = 0 \text{ if } x_1 \leq x_2 \text{ and } x_2 \leq x_1 \text{ and } (1; 0), \text{ and } (0; 1) \text{ otherwise} \\ f_2 &: f(x_1, x_2) = 0 \text{ if } x_1 \leq x_2 \text{ and } x_2 \leq x_1 \text{ and } (0; 1), \text{ and } (1; 0) \text{ otherwise} \end{aligned}$$

²with a the minor difference that they directly assume the state space is compact with bounded measure

That is, the feature maps from \mathbb{R}^2 to the set of binary encoding of dimension 2, $\{(\pm 1; 0); (0; 1)\}$. We further assume that $f_k^? = (p_1(y); p_2(y))$ for some distributions $p_1; p_2$, which is identical for all task k , where $k \in [K]$, $p_1 + p_2 = 1$. We also assume that $f_k^?$ is known to the learner a priori, i.e. $f_k^? = (p_1(y); p_2(y))$ for all $k \in [K]$, so all the learner needs to do is to pick the correct out of two candidates.

Given the above setup, it's easy to verify that both [Assumption 3.1](#) and [Assumption 3.4](#) are satisfied, because the decision boundary of both f_1 and f_2 passes through the support of the source tasks, and all $f_k^?$'s are identical. However, f_1 and f_2 are equivalent in \mathcal{S}_1 and \mathcal{S}_2 in terms of their representation power, therefore no algorithm can always pick the correct feature function with probability more than $1/2$, regardless of the number of samples. Suppose the true feature and the algorithm incorrectly chooses f_2 . Then, for $x \in \mathcal{S}_3$ $f(x) = 0$ which has probability mass $1/2$, $P_3^?(y|x) = p_2$ whereas $P_3^?(y|x) = p_1$. Thus, the expected total variation distance between $P_3^?$ and $P_k^?$ is $1/2$. \square

The above construction shows that our assumption are not sufficient to permit reliable representation transfer, even in the supervised learning setting. Yet, surprisingly, these assumptions are sufficient in the RL setting, implying somehow that transfer learning in RL is easier than transfer learning in SL. To understand this phenomenon, observe that in RL, the marginal distribution $\pi(s)$ is not independent from the conditional density $P(y|s; a)$ we desire to estimate. In particular, if one collects data in the source tasks in an online fashion via running a policy $\pi(s; a)$ is structurally restricted to be an occupancy distribution generated by the ground-truth transition $P(y|s; a)$. Such a connection can only exist in Markov chains, and our analysis elegantly utilizes this additional structure to establish the soundness of the learned representation. Also note, crucially, that we never learn a representation to capture $\pi(s)$ which would suffer from similar issues as the supervised learning setting, but is not necessary for sample-efficient RL.

Next, we prove the impossibility result in [Theorem 5.1](#), restated below as [Theorem D.3](#). This result shows that one can not achieve online learning in the source tasks without significantly stronger assumptions such as [Assumption 5.1](#). Before that, we provide a preliminary version, showing that the learned \hat{f} is not sufficient to fit the transition model in the target task, which motivates the construction in [Theorem D.3](#).

Theorem D.2 (Impossibility Result: Model Learning) Let M_K be the set of K -task multi-set that satisfies

1. all tasks are Block MDPs;
2. all tasks satisfy [Assumption 3.3](#) and [Assumption 3.4](#);
3. the latent dynamics are exactly the same for all source and target tasks.

For any pre-training algorithm A , there exists $P_k^? \forall k=1:K \in M_K$ and an occupancy distribution π_k on the target task, such that with probability at least $1/2$, A will output a feature \hat{f} and for any

$$E_k \pi_k^{\hat{f}}(s; a) > \frac{1}{2} \sum_{j \in \mathcal{S}} P_k^?(j; a) \pi_k(j) \quad \forall k=1:K$$

Proof of [Theorem D.2](#) Consider a tabular MDP with 2 latent states z_1, z_2 and an observation state space $\mathcal{S} = \{R_1, R_2, B_1, B_2\}$, where in task 1 one can only observe R_1, R_2 and in task 2 one can only observe B_1, B_2 . Correspondingly, $\pi_1(s; z)$ is only supported on $\{R_1, R_2\}$ (i.e., $\pi_1(R_i | z_i) = 1$) and similar for task 2. Let the latent state transition be such that $P(z_1 | z_1; a) = 1$ and $P(z_2 | z_2; a) = 1$, i.e. only self-transition regardless of the actions.

Now, consider a 2-element feature class $f = \{f_1; f_2\}$ such that

$$\begin{aligned} f_1 &= (R_1: 1; R_2: 2; B_1: 1; B_2: 2) \\ f_2 &= (R_1: 1; R_2: 2; B_1: 2; B_2: 1) \end{aligned}$$

Denote $\pi_i(s; a) = \pi_i(s; a)$ for $i \in [1; 2]$. Consider for each task k , a 2-element class in the form of $f_k = (o_k(s|z_1); o_k(s|z_2)); (o_k(s|z_2); o_k(s|z_1))$.

Notice that π_1 and π_2 are merely permutations of one another and so given any single task data, the two hypothesis will not be distinguishable by any means. Therefore, for any algorithm, there is at least probability $\frac{1}{2}$ that it will choose the wrong hypothesis if the ground truth is sampled between π_1 and π_2 uniformly at random. Suppose π_1 is the correct hypothesis and π_2 is the one that the algorithm picks (i.e., $\hat{\pi} = \pi_2$). Let task 3 be such that any state emits R_2 and $B_1 = B_2$ each with probability $\frac{1}{2}$ (i.e., $o_3(R_i|z_i) = o_3(B_i|z_i) = 0.5$). This construction satisfies [Assumption 3.3](#) and [Assumption 3.4](#).

Then, within task 3, one would encounter observations from B_1 and B_2 which should be mapped to latent states z_1 and z_2 respectively by the true decoder f , but are instead both mapped to latent state z_1 by the learned decoder \hat{f} , and thus z_1 and z_2 become indistinguishable. Suppose $\kappa(z_1) = \kappa(z_2) = \frac{1}{2}$, then

$$\begin{aligned} & E_{\kappa} [k^{\wedge}(s; a)^{\>}(\cdot) - P^?(j; s; a)_{k_{TV}}] \\ &= \frac{1}{4} k^{\wedge}(R_1)^{\>}(\cdot) - (R_1)^{?>}(\cdot)_{k_{TV}} + \frac{1}{4} k^{\wedge}(B_1)^{\>}(\cdot) - (B_1)^{?>}(\cdot)_{k_{TV}} \\ &\quad + \frac{1}{4} k^{\wedge}(R_2)^{\>}(\cdot) - (R_2)^{?>}(\cdot)_{k_{TV}} + \frac{1}{4} k^{\wedge}(B_2)^{\>}(\cdot) - (B_2)^{?>}(\cdot)_{k_{TV}} \\ &= k_{o_1} - o_1^2_{k_{TV}} = 4 + k_{o_2} - o_2^2_{k_{TV}} = 4 + k_{o_2} - o_2^2_{k_{TV}} = 4 + k_{o_1} - o_1^2_{k_{TV}} = 4 \\ &\quad - \frac{1}{4} k_{o_1}^2 - o_2^2_{k_{TV}} + \frac{1}{4} k_{o_1}^2 - o_2^2_{k_{TV}} \\ &= \frac{1}{2}; \end{aligned}$$

where the last second inequality uses triangle inequality, and the last equality comes from the fact that $o_3(z_1)$ and $o_3(z_2)$ have disjoint support which implies that $k_{o_1}^2 + k_{o_2}^2_{k_{TV}} = 1$. \square

Now, we are ready to restate and prove [Theorem 5.1](#).

Theorem D.3 (Impossibility Result: Optimal Policy Identification) Let M_{κ} be the set of K -task multi-set that satisfies

1. all tasks are Block MDPs;
2. all tasks satisfy [Assumption 3.3](#) and [Assumption 3.4](#);
3. the latent dynamics are exactly the same for all source and target tasks.

For any pre-training algorithm A , there exists $P_{\kappa}^? g_{k=1:K} \in 2^{M_{\kappa}}$, such that with probability at least $\frac{1}{2}$, A will output a feature \hat{f} , such that for any policy taking the functional form of

$$(s) = f \hat{f}^{\wedge}(s; a) g_{a2A}; f r(s; a) g_{a2A}, \text{ we have}$$

$$V^? - V \geq \frac{1}{2}.$$

Proof of [Theorem D.3](#) Consider a tabular MDP with $H = 2$, two latent states $z_1; z_2$ for $h = 1$ and two latent states $z_3; z_4$ for $h = 2$.

- For $h = 1$, let there be two actions $a_1; a_2$. Let the observation state space $\mathcal{S} = R_1 \cup R_2 \cup B_1 \cup B_2$, where in task 1 one can only observe R_2 and in task 2 one can only observe $B_1 \cup B_2$. Correspondingly, $o_1(s|z)$ is only supported on $R_1 \cup R_2$ (i.e., $o_1(R_i|z_i) = 1$) and similar for task 2. Let the latent state transition be such that $P(z_3|z_1; a_1) = P(z_3|z_2; a_2) = 1$, and $P(z_4|z_1; a_2) = P(z_4|z_2; a_1) = 1$. All rewards are 0 for $h = 1$.
- For $h = 2$, in state z_3 , all actions have reward 1, and in state z_4 all actions have reward 0.
- The initial state distribution is $d_0(z_1) = d_0(z_2) = \frac{1}{2}$.

Algorithm 3 Multi-task REPLEARN

- 1: Input: Datasets $\{D_k\}_{k=1:K}$, model class $\mathcal{G}_k, k = 1; \dots; K$.
- 2: Compute $\text{MLE}(\hat{\theta}_{1:k}) := \underset{\theta \in \mathcal{G}_k}{\text{argmax}} E_{D_k} \log(\theta(s; a))$.
- 3: Return $\hat{\theta}$.

Now, consider a 2-element feature class $\mathcal{F}_{1;2}$ for $h = 1$, such that

$$\begin{aligned} \mathcal{F}_1 &= \{R_1 = 1; R_2 = 2; B_1 = 1; B_2 = 2\} \\ \mathcal{F}_2 &= \{R_1 = 1; R_2 = 2; B_1 = 2; B_2 = 1\} \end{aligned}$$

Denote $\theta_i(s; a) = e_{\theta_i(s; a)}$ for $i \in [1; 2]$. In addition, define $\mathcal{F}_{3;4}$ where

$$\begin{aligned} \mathcal{F}_3 &= \{z_3 = (1; 0); z_4 = (0; 1)\} \\ \mathcal{F}_4 &= \{z_3 = (1; 0); z_4 = (0; 1)\} \end{aligned}$$

Notice that \mathcal{F}_1 and \mathcal{F}_2 are merely permutations of one another and so given any single task data, the two hypothesis will not be distinguishable by any means. Therefore, for any algorithm, there is at least probability $1/2$ that it will choose the wrong hypothesis. Suppose θ_1 is the correct hypothesis and θ_2 is the one that the algorithm picks (i.e., $\hat{\theta} = \theta_2$). Let task 3 be such that any state emits to R_2 and $B_1 = B_2$ each with probability $1/2$ (i.e., $\theta_3(R_i|z_i) = \theta_3(B_i|z_i) = 0.5$). This construction satisfies [Assumption 3.3](#) and [Assumption 3.4](#).

Then, for any policy that only make decision based on $\theta(s; a)$ and $\theta(s; a)$, would output the same action for observations in R_1 and B_2 , or for B_1 and R_2 . However, notice that the optimal policy, which would try to go to z_3 from either z_1 or z_2 , will pick a_1 at R_1 and B_1 while picking a_2 at R_2 and B_2 , which means that the optimal policy will not agree on R_1 and B_2 , and it also will not agree on R_2 and B_1 . Thus clearly, no such policy as defined above is capable of capturing the optimal policy. From the reward perspective, notice that $d(z_1) = d(z_2) = 1/2$ and $d(R_1) = d(R_2) = d(B_1) = d(B_2) = 1/4$. Since $(R_1) = (B_2)$, the agent will only be able to collect reward at one of the R_1 and B_2 (but not at both). Similarly, since $(R_2) = (B_1)$, the agent will only be able to collect reward at one of the R_2 and B_1 by reaching z_3 (but not at both). This means that will have average reward $1/4$. Since the optimal policy will be able to collect reward at all $R_2; B_1; B_2$, it will have average reward $1/2$. This concludes the proof. \square

[Theorem D.2](#) and [Theorem D.3](#) show that it's impossible to allow online learning in the source tasks without much stronger assumptions. In our paper, we show that our [Assumption 5.1](#), which ensures reachability in the raw states, is sufficient to establish an end-to-ending online transfer learning result. However, it is unclear if [Assumption 5.1](#) is necessary for online learning. We leave this as an important direction of future work.

E REWARD-FREE REP-UCB

In this section, we adapt the Rep-UCB algorithm ([Uehara et al., 2021](#)) for reward-free exploration in a single task. We drop all task subscripts as this section is for a single task only, i.e. think about the task as being each source task. The original Rep-UCB algorithm was for finite-horizon discounted MDPs, so we modify it to work for our undiscounted and finite-horizon setting. Our goal is to prove that Rep-UCB can learn a model that satisfies strong TV guarantees, i.e. [Theorem E.1](#) and [\(5\)](#). Note that FLAMBE ([Agarwal et al., 2020](#), Theorem 2) can be used for this directly, but at a worse (polynomial) sample complexity. Thus, we do a bit more work to derive a new model-learning algorithm for low-rank MDPs, based on Rep-UCB, that is more sample efficient in the source tasks.

A finite-horizon analysis of Rep-UCB was done in BRIEE ([Zhang et al., 2022](#)), so here we just need to replace BRIEE's RepLearn with that of the MLE, which is how we learn θ and b , as in Rep-UCB. Recall the notation of ([Zhang et al., 2022](#)),

$$\theta_{h;n}(s; a) = \frac{1}{n} \sum_{i=0}^{X-1} d_h^{b_i}(s) \text{Unif}(a)$$

Algorithm 4 REWARDFREE REP-UCB

- 1: Input: Regularizer η_n , bonus scaling β_n , model class $\mathcal{M} = \{P_h^? : h = 1, \dots, H\}$, number of episodes N .
- 2: Initialize b_0 as random and $D_{h;0}^0 = \emptyset$.
- 3: for episode $n = 1; 2; \dots; N$ do
- 4: Data collection from b_{n-1} : for $h = 1; 2; \dots; H-1$,

$$s \sim d_h^{b_{n-1}}; a \sim \text{Unif}(A); s^0 \sim P_h^?(s; a);$$

$$\mathbf{e} \sim d_h^{b_{n-1}}; \mathbf{a} \sim \text{Unif}(A); \mathbf{e}^0 \sim P_h^?(s; \mathbf{a}); \mathbf{a}^0 \sim \text{Unif}(A); \mathbf{e}^{00} \sim P_h^?(s^0; \mathbf{a}^0);$$

$$D_{h;n} = D_{h;n-1} \cup \{f(s; a; s^0)g; D_{h;n}^0 = D_{h;n-1}^0 \cup \{f(s^0; \mathbf{a}^0; \mathbf{e}^{00})g\};$$
 For $h = 0$, only collect $D_{0;n}$.
- 5: Learn model via MLE: for all $h = 0; 1; \dots; H-1$,

$$\hat{P}_{h;n} = (\hat{b}_{h;n}; \hat{b}_{h;n}) = \underset{h; \mathcal{M}_h}{\text{argmax}} E_{D_{h;n}} [\sum_{h; \mathcal{M}_h} \log P_h^?(s; a)^T \cdot \log P_h^?(s; a)]:$$
- 6: Update exploration bonus: for all $h = 0; 1; \dots; H-1$,

$$\hat{b}_{h;n}(s; a) = \frac{1}{n} \sum_{i=0}^{n-1} \hat{b}_{h;n}^{b_{i-1}}(s; a)$$

$$\hat{b}_{h;n} = \frac{1}{n} \sum_{(s;a) \in \mathcal{D}_{h;n}} \hat{b}_{h;n}(s; a) \hat{b}_{h;n}(s; a)^T + \eta_n I:$$
- 7: Learn policy $b_n = \underset{P_n, \hat{b}_n}{\text{argmax}} V_{P_n, \hat{b}_n}$ and let \hat{P}_n be its value.
- 8: Let $\hat{b} = \underset{n}{\text{argmin}} \sum_{n=2}^N \hat{P}_n$.
- 9: Output: $\hat{b}; \hat{P}_b$.

$$\hat{b}_{h;n}(s; a) = \frac{1}{n} \sum_{i=0}^{n-1} \hat{b}_{h;n}^{b_{i-1}}(s; a) = \frac{1}{n} \sum_{i=0}^{n-1} E_{\mathbf{e} \sim d_h^{b_{i-1}}; \mathbf{a} \sim \text{Unif}(A)} [P_h^?(s; \mathbf{e}; \mathbf{a}) \text{Unif}(\mathbf{a})]$$

$$\hat{b}_{h;n}(s; a) = \frac{1}{n} \sum_{i=0}^{n-1} d_h^{b_i}(s; a)$$

$$\hat{b}_{h;n} = n E_{(s;a)} (\hat{b}_{h;n}(s; a) \hat{b}_{h;n}(s; a)^T + \eta_n I):$$

By using MLE (Uehara et al., 2021, Lemma 18) to learn models, with probability at least $1 - \delta$ for any $n = 1; 2; \dots; N$ and $h = 0; 1; \dots; H-1$, we have

$$\max_{h;n} E_{D_{h;n}} \|\hat{P}_{h;n}(s; a) - P_h^?(s; a)\|_{TV}^2 \leq \frac{2}{n} E_{D_{h;n}} \|\hat{b}_{h;n}(s; a) - P_h^?(s; a)\|_{TV}^2 \leq \frac{2}{n}; \quad (3)$$

where

$$\frac{2}{n} = O\left(\frac{\log(jMj/nH)}{n}\right);$$

and $jMj = \max_{h \in [H]} \sum_{h'} |j_{h'}|$. We also adopt the same choice of η_n parameters as BRIEE, which we assume from now on.

$$\eta_n = \left(\frac{d \log(jMj/nH)}{n}\right)$$

$$\eta_n = \frac{p}{njAj^2 + nd}:$$

As in Rep-UCB, we posit standard assumptions about realizability and normalization, on the (source) task of interest.

Assumption E.1 For any $h = 0; 1; \dots; H-1$, we have $\|P_h^?\|_2 \leq \frac{1}{\beta_h}$ and $\|R_h^?\|_2 \leq \frac{1}{\beta_h}$. For any $\beta_h \geq \frac{1}{k_2}$, $k_2 \geq 1$. For all $\beta_h \geq \frac{1}{k_2}$ and any function $g: S \rightarrow \mathbb{R}$, we have $\|g\|_s \leq k_2 \|g\|_1 \leq k_1 \frac{1}{\beta_h} \|g\|_1$.

Lemma E.1 Let r be any reward function. Suppose we ran Algorithm 4 with line 7 having reward $r + \hat{b}_n$ instead of just \hat{b}_n . Then, for any $\delta \in (0, 1)$, w.p. at least $1 - \delta$, we have

$$\sum_{n=0}^{N-1} V_{\hat{P}_n, r + \hat{b}_n} - V_{P^?, r} \leq O\left(H^2 d^2 jAj^{1.5} \frac{p}{N \log(jMj/NH)}\right)$$

Proof. Start from the third equation of Zhang et al. (2022, Theorem A.4). Following their proof until the last page of their proof, we arrive at the following: for any $1 \leq n \leq N$,

$$\begin{aligned} & \mathbb{E}_{\mathbf{p}_n; \mathbf{r} + \mathbf{b}_n} V_{\mathbf{p}_n; \mathbf{r}}^{b_n} - \mathbb{E}_{\mathbf{p}_n; \mathbf{r}} V_{\mathbf{p}_n; \mathbf{r}}^{b_n} \\ & \leq \sum_{h=0}^{H-1} \mathbb{E}_{\mathbf{s}; \mathbf{a}} \mathbb{E}_{\mathbf{p}_n; \mathbf{r}} \left[\sum_{h_n=0}^{H-h-1} k_h(\mathbf{s}; \mathbf{a}) \left(\frac{1}{jA_j} \frac{1}{d^{2n}} + \frac{1}{jA_j} \frac{1}{d^{2n}} \right) \right] \\ & + (2H+1) \sum_{h=0}^{H-1} \mathbb{E}_{\mathbf{s}; \mathbf{a}} \mathbb{E}_{\mathbf{p}_n; \mathbf{r}} \left[\sum_{h_n=0}^{H-h-1} k_h(\mathbf{s}; \mathbf{a}) \left(\frac{1}{njA_j} \frac{1}{d^{n+1}} + (2H+1) \frac{1}{jA_j} \frac{1}{d^n} \right) \right] \end{aligned}$$

By elliptical potential arguments, we have

$$\sum_{n=0}^{H-1} \mathbb{E}_{\mathbf{s}; \mathbf{a}} \mathbb{E}_{\mathbf{p}_n; \mathbf{r}} \left[\sum_{h_n=0}^{H-h-1} k_h(\mathbf{s}; \mathbf{a}) \right] \leq \frac{S}{dN \log \left(1 + \frac{N}{d} \right)}$$

Thus, summing over n , noting that n_1, n_2, \dots, n_H are increasing in n , we can combine the above to get,

$$\begin{aligned} & \sum_{n=0}^{H-1} \mathbb{E}_{\mathbf{p}_n; \mathbf{r} + \mathbf{b}_n} V_{\mathbf{p}_n; \mathbf{r}}^{b_n} - \mathbb{E}_{\mathbf{p}_n; \mathbf{r}} V_{\mathbf{p}_n; \mathbf{r}}^{b_n} \\ & \leq \frac{S}{dN \log \left(1 + \frac{N}{d} \right)} \left(H \frac{1}{jA_j} \frac{1}{d^{2N}} + H^2 \frac{1}{NjA_j} \frac{1}{d^{N+1}} \right) \\ & \leq \frac{S}{dN \log \left(1 + \frac{N}{d} \right)} \left(H^p \frac{1}{NjA_j^3} \frac{1}{d^{N+1}} + H^2 \frac{1}{NjA_j} \frac{1}{d^{N+1}} \right) \\ & \leq \frac{S}{dN \log \left(1 + \frac{N}{d} \right)} \left(H^2 \frac{1}{djA_j^3} \frac{1}{\log(jMjNH)} + d^3 \frac{1}{\log(jMjNH)} \right) \\ & \leq O \left(H^2 d^2 jA_j^{1.5} \frac{1}{N \log(jMjNH)} \right) \end{aligned}$$

□

This gives the following useful corollary for reward free exploration.

Lemma E.2 For any $\epsilon \in (0, 1)$ w.p. at least $1 - \epsilon$ we have

$$\| \mathbf{v}_b \| \leq O \left(H^2 d^2 jA_j^{1.5} \frac{1}{N \log(jMjNH)} \right)$$

Proof. By definition of \mathbf{v}_b , we have

$$\frac{N}{2} \mathbf{v}_b = \sum_{n=N-2}^{H-1} V_{\mathbf{p}_n; \mathbf{b}_n}^{b_b} - \sum_{n=0}^{H-1} V_{\mathbf{p}_n; \mathbf{b}_n}^{b_b};$$

which is bounded by the previous lemma and the fact that $V_{\mathbf{p}_0; \mathbf{r}=0}^{b_b} = 0$, since in Algorithm 4, the reward function is zero. □

Conditioning on this, we now show that the environment \mathcal{E}_b has low TV error for any policy-induced distribution.

Theorem E.1 For any policy π , we have

$$\sum_{h=0}^{H-1} \mathbb{E}_{\mathbf{d}_p; \mathbf{r}} \left[\sum_{h_n=0}^{H-h-1} P_h(\mathbf{s}; \mathbf{a}) \left\| \mathbf{p}_{h; \mathbf{b}}(\mathbf{s}; \mathbf{a}) - \mathbf{p}_h(\mathbf{s}; \mathbf{a}) \right\|_{TV} \right] \leq O \left(H^3 d^2 jA_j^{1.5} \frac{1}{N \log(jMjNH)} \right) := \epsilon_{TV}$$

Proof. In this proof, let $\mathbb{P} = \mathbb{P}_b$, which is the returned environment from the algorithm. Let $r(s; a) = P_h^?(s; a) - \mathbb{P}_h(s; a)_{TV} \in [0; 2]$. Then,

$$\begin{aligned} & \sum_{h=0}^{H-1} E_{d_{P^?,h}} E_{d_{\mathbb{P},h}} [r(s; a)] \\ &= \sum_{h=0}^{H-1} V_{P^?,r} V_{\mathbb{P},r} \\ &= \sum_{h=0}^{H-1} E_{d_{\mathbb{P},h}} E_{P_h^?(s;a)} E_{\mathbb{P}_h(s;a)} V_{P^?,r,h+1}(s^0) \quad (\text{Simulation lemma}) \\ & \leq \sum_{h=0}^{H-1} E_{d_{\mathbb{P},h}} P_h^?(s; a) - \mathbb{P}_h(s; a)_{TV} : \end{aligned}$$

Thus,

$$\begin{aligned} & \sum_{h=0}^{H-1} E_{d_{P^?,h}} P_h^?(s; a) - \mathbb{P}_h(s; a)_{TV} \\ & \leq (2H + 1) \sum_{h=0}^{H-1} E_{d_{\mathbb{P},h}} P_h^?(s; a) - \mathbb{P}_h(s; a)_{TV} \quad (\text{by (Zhang et al., 2022, Lemma A.1)}) \\ & \leq H \sum_{h=0}^{H-1} E_{d_{\mathbb{P},h}} \left[\sum_{i=0}^q \frac{r}{jA_j} \frac{\log(jM_j NH)}{N} \right] \\ & \leq H \sum_{h=0}^{H-1} V_{\mathbb{P},b_b} + \sum_{h=0}^{H-1} \frac{r}{2jA_j} \frac{\log(jM_j NH)}{N} \\ & \leq H \sum_{h=0}^{H-1} H^2 d^2 jA_j^{1.5} \frac{\log(jM_j NH)}{N} + \sum_{h=0}^{H-1} \frac{r}{jA_j} \frac{\log(jM_j NH)}{N} \quad (\text{by Lemma E.2}) \\ & \leq 2O \sum_{h=0}^{H-1} H^3 d^2 jA_j^{1.5} \frac{\log(jM_j NH)}{N} : \end{aligned}$$

□

This also gives us a guarantee on the TV distance between the visitation distributions induced by $\mathbb{P}^?$ vs. by \mathbb{P} .

Lemma E.3 Suppose \mathbb{P} satisfies the following for all $h = 0; 1; \dots; H - 1$,

$$\sum_{h=0}^{H-1} E_{d_{P^?,h}} \mathbb{P}_h(s; a) - P_h^?(s; a)_{TV} \leq \epsilon_h \quad (4)$$

Then, for any $h = 0; 1; \dots; H - 1$, we have

$$\sum_{t=0}^{H-1} d_{\mathbb{P},h} - d_{P^?,h} \leq \sum_{t=0}^{H-1} \epsilon_t$$

Note, for $h = 0$, the sum is empty so the right hand side is 0.

Proof. We proceed by induction for $t = 0; 1; \dots; H - 1$. For the base case $t = 0$, no transition has been taken, so $d_{\mathbb{P},0} = d_{P^?,0}$. Now let $t = 0; 1; \dots; H - 2$ be arbitrary, and suppose that the claim is true for t (IH). We want to show the claim holds for $t + 1$. One key fact we'll use is that, for any measure, we have $k_{TV} = \sup_{k \in \mathcal{K}} |E[f]|$. Below we use the notation that $f(s; a) = E_{a(s)} f(s; a)$.

$$k_{\mathbb{P},h+1} - k_{h+1} \leq k_{TV}$$

$$\begin{aligned}
 &= \sup_{k \in \mathcal{K}_1} E_{d_{\mathcal{P},h+1}} [f(s; a)] - E_{d_{h+1}} [f(s; a)] \\
 &= \sup_{k \in \mathcal{K}_1} E_{(\mathbf{s}; \mathbf{a}) \sim d_{\mathcal{P},h}; (s; a) \sim \mathcal{P}_h(\mathbf{s}; \mathbf{a})} [f(s; h+1)] - E_{(\mathbf{s}; \mathbf{a}) \sim d_h; (s; a) \sim P_h^2(\mathbf{s}; \mathbf{a})} [f(s; h+1)] \\
 &\quad \sup_{k \in \mathcal{K}_1} E_{(\mathbf{s}; \mathbf{a}) \sim d_{\mathcal{P},h}} E_{(\mathbf{s}; \mathbf{a}) \sim d_h} E_{\mathcal{P}_h(\mathbf{s}; \mathbf{a})} f(s; h+1) \\
 &\quad + \sup_{k \in \mathcal{K}_1} E_{(\mathbf{s}; \mathbf{a}) \sim d_h} E_{\mathcal{P}_h(\mathbf{s}; \mathbf{a})} f(s; a) - E_{P_h^2(\mathbf{s}; \mathbf{a})} f(s; h+1) \\
 &\leq \sum_{t=0}^H \left(\sup_{k \in \mathcal{K}_1} E_{(\mathbf{s}; \mathbf{a}) \sim d_h} E_{\mathcal{P}_h(\mathbf{s}; \mathbf{a})} E_{P_h^2(\mathbf{s}; \mathbf{a})} f(s; h+1) \right) \quad (\text{by IH and Jensen}) \\
 &\leq \sum_{t=0}^H \left(\sup_{k \in \mathcal{K}_1} E_{(\mathbf{s}; \mathbf{a}) \sim d_h} E_{\mathcal{P}_h(\mathbf{s}; \mathbf{a})} E_{P_h^2(\mathbf{s}; \mathbf{a})} f(s; h+1) \right) \quad (\text{by (4) and } k \in \mathcal{K}_1)
 \end{aligned}$$

as desired. □

Thus, when combined with [Theorem E.1](#), we have for $r \in (0, 1]$ and any policy π ,

$$\sup_{k \in \mathcal{K}_{\mathcal{P},h}} |d_{\mathcal{P},h} - d_{P_h^2}|_{TV} = O\left(H^3 d^2 j A_j^{1.5} \frac{\log(jMj NH)}{N} \right) = \epsilon_{TV} \quad (5)$$

In other words, the sample complexity needed for a model-err of ϵ is

$$O\left(\frac{H^6 d^4 j A_j^3 \log(jMj NH)}{\epsilon_{TV}^2} \right)$$

Note this is much better than FLAMBE's guarantee ([Agarwal et al., 2020](#), Theorem 2) which requires,

$$O\left(\frac{H^{22} d^7 j A_j^9 \log(jMj NH)}{\epsilon_{TV}^{10}} \right)$$

F REWARD-FREE EXPLORATION

In this section, we show that the mixture policy returned by [Algorithm 2](#) has good coverage. Recall that [Algorithm 2](#) contains two main steps:

Step 1 Learn a model \mathcal{P} . This was the focus of the previous section, where our model-UCB method obtained a strong TV guarantee [\(5\)](#) by requiring number of episodes at most,

$$N_{\text{REWARDFREE}} = O\left(\frac{H^6 d^4 j A_j^3 \log(jMj NH)}{\epsilon_{TV}^2} \right)$$

Step 2 Run LSVI-UCB ([Algorithm 6](#)) in the learned model \mathcal{P} with reward at the t -th episode being $b_{h,t}$ and `UNIFORMACTIONS = TRUE`. The optimistic bonus pushes the algorithm to explore directions that are not well-covered yet by the mixture policy up to this point. With elliptical potential, we can establish that this process will terminate in polynomial number of steps.

We now focus on Step 2. Let π_h^{t+1} denote rolling-in for h steps and taking uniform actions on the $h+1$ step, thus inducing a distribution over $\mathcal{S}_{h+1}; \mathcal{A}_{h+1}$. We abuse the notation a little and use π_h^{t+1} for a policy that just takes one uniform action from the initial distribution.

Lemma F.1 Let $\epsilon \in (0, 1)$ and run REWARDFREE ([Algorithm 2](#)). Let $\Sigma_{h,N}$ be the empirical covariance at the N -th iteration of LSVI-UCB ([Algorithm 6](#)). Then, w.p. at least $1 - \epsilon$ we have,

$$\sup_{h=0}^H \sum_{t=1}^N E_{\pi_h^{t+1}; \mathcal{P}} b_h(s_h; a_h) - \sum_{h=1}^H \text{Ad}^{1.5} H^3 \frac{\log(dNH)}{\epsilon} = N$$

Proof. In this proof, we'll treat the empirical MDP \mathcal{M}_h^e , as that is the environment we're running in. Thus, we abuse notation and \mathcal{M}_h^e is the model-based perspective of the linear MDP, where $b_{h,e}$ is $\sum_{k=1}^e b_h(s_h^k; a_h^k)$ (s_{h+1}^k). Also, in Algorithm 2, we set reward to be zero, but for the purpose of this analysis, suppose the reward function is precisely the (unscaled) bonus in LSVI-UCB, i.e. $r_{h,e}(s_h; a_h) = b_{h,e}(s_h; a_h)$. This does not change the algorithm at all since the β -scaling of the bonus dominates this reward in the definition of \mathcal{M}_h^e , but thinking about the reward in this way will make our analysis simpler.

Recall the high-level proof structure of reward free guarantee of linear MDP (with known features) (Wang et al., 2020, Lemma 3.2).

Step 1 Show that $\|b_{h,e}\| \leq 2V_h$ and w.p.1, for all $h; e$,

$$\|b_{h,e}\| = \sup_{f \in \mathcal{F}_h} \sum_{k=1}^e b_h(s_h^k; a_h^k) - P_h^?(s_h; a_h) f(s_h; a_h)$$

This step only uses self-normalized martingale bounds. So, line 9 can use any martingale sequence of states and actions, and this claim still holds, with $b_{h,e}$ using the appropriate covariance under the data.

Step 2 Show optimism conditioned on Step 1. Specifically, for $e = 1; 2; \dots; N$, we have $E_{d_0} V_0^?(s_0; r_e) - V_{0,e}(s_0) \geq 0$. To show this, we need that $b_{h,e}(s_h) = b_{h,e}(s_h; a_h^e(s_h)) \geq b_{h,e}(s_h; a_h^?(s_h))$ (this is for the unclipped case of optimism), which we have satisfied in the algorithm, i.e. a_h^e is greedy w.r.t. \mathcal{M}_h^e .

Step 3 Bound the sum $\sum_{e=1}^P V_{h,e}$, where we decompose it as a sum of expected bonuses with the expectation is under \mathcal{P}^e .

Step 3 is the only place where we use the fact that s_h^k are data sampled from rolling out. For Step 1 and 2, please refer to existing proofs in (Agarwal et al., 2019; Jin et al., 2020b; Wang et al., 2020).

Now we show Step 3 for our modified algorithm with uniform actions. First, let us show a simulation lemma. For any episode $e = 1; 2; \dots; N$, for any s_0 , recalling definition of reward being $b_{h,e}$, we have

$$\begin{aligned} V_{0,e}(s_0) &\leq (1 + \beta) b_{0,e}(s_0; a_0^e(s_0)) + P_{0,e}^?(s_0; a_0^e(s_0)) V_{1,e} \\ &\leq (1 + 2\beta) b_{0,e}(s_0; a_0^e(s_0)) + P_{0,e}^?(s_0; a_0^e(s_0)) V_{1,e}; \end{aligned}$$

where the first inequality is due to the thresholding $b_{h,e}$'s and the second inequality is due to Step 1. Continuing in this fashion, we have

$$E_{d_0} V_{0,e}(s_0) \leq (1 + 2\beta)^H E_{\mathcal{P}^e} [b_{h,e}(s_h; a_h)]:$$

Summing over $e = 1; 2; \dots; N$, we have

$$\begin{aligned} E_{d_0} \sum_{e=1}^N V_{0,e}(s_0) &\leq \sum_{e=1}^N \sum_{h=0}^{H-1} E_{\mathcal{P}^e} [b_{h,e}(s_h; a_h)] \\ &\leq A \sum_{h=0}^{H-1} E_{\mathcal{P}^e} [b_{h,e}(s_h; a_h)] \end{aligned}$$

For each $h = 0; 1; \dots; H-1$, apply Azuma's inequality to the martingale difference sequence $e = E_{\mathcal{P}^e} [b_{h,e}(s_h; a_h)] - b_{h,e}(s_h^e; a_h^e)$. The envelope is at most $2V_h$ So, w.p.1,

$$A \sum_{h=0}^{H-1} b_{h,e}(s_h^e; a_h^e) + A \sqrt{N \log(H/\delta)}$$

Now apply a self-normalized elliptical potential bound to the first term, giving that

$$\sum_{h=0}^H \sum_{e=1}^N \mathbb{E} \left[\sum_{i=1}^d \frac{1}{N} \sum_{e=1}^N \frac{b_{h,e}(s_h^e; a_h^e)^2}{\sum_{i=1}^d \sum_{e=1}^N b_{h,e}(s_h^e; a_h^e)^2} \right] \leq H \frac{d}{N} \log(N).$$

Thus, we finally have

$$\sum_{e=1}^N \mathbb{E}_{d_0} \left[\sum_{i=1}^d \frac{1}{N} \sum_{e=1}^N \frac{b_{0,e}(s_0)^2}{\sum_{i=1}^d \sum_{e=1}^N b_{0,e}(s_0)^2} \right] \leq H \frac{d}{N} \log(NH).$$

Consider any episode $e = 1; 2; \dots; N$. By definition, $b_{h,N}(s; a) = b_{h,e}(s; a)$, so for all s, a we have pointwise that $b_{h,N}(s; a) \leq b_{h,e}(s; a)$. Hence, for all s , we have $V_0^?(s; r^N) \leq V_0^?(s; r^e)$, and further using optimism, we have

$$N \mathbb{E}_{d_0} [V_0^?(s_0; r^N)] \leq \sum_{e=1}^N \mathbb{E}_{d_0} [V_0^?(s_0; r^e)] \leq \sum_{e=1}^N \mathbb{E}_{d_0} \left[\sum_{i=1}^d \frac{1}{N} \sum_{e=1}^N \frac{b_{0,e}(s_0)^2}{\sum_{i=1}^d \sum_{e=1}^N b_{0,e}(s_0)^2} \right] \leq H \frac{d}{N} \log(NH).$$

Now consider any h and policy π_h , and consider rolling it out for $h+1$ steps and taking a random action. Then we have

$$\mathbb{E}_{\pi_h} \left[\sum_{i=1}^d \frac{1}{h+1} \sum_{h=1}^N \frac{b_h(s_h; a_h)^2}{\sum_{i=1}^d \sum_{h=1}^N b_h(s_h; a_h)^2} \right] \leq \mathbb{E}_{d_0} \left[\sum_{i=1}^d \frac{1}{h+1} \sum_{h=1}^N \frac{b_{0,h}(s_0)^2}{\sum_{i=1}^d \sum_{h=1}^N b_{0,h}(s_0)^2} \right] \leq H \frac{d}{h+1} \log(NH).$$

Summing over h incurs an extra H factor on the right. This concludes the proof. \square

Lemma F.2 (One-step back for Linear MDP) Suppose $\mathcal{P}_h = (P_h; R_h)$ is a linear MDP. Suppose π is any mixture of policies, and let $\Sigma_h := n \mathbb{E}_{\pi} \left[\sum_{i=1}^d \frac{1}{n} \sum_{h=1}^n \frac{b_h(s_h; a_h)^2}{\sum_{i=1}^d \sum_{h=1}^n b_h(s_h; a_h)^2} \right] + I$ denote the unnormalized covariance. For any $g: S \rightarrow \mathbb{R}$, policy π , and $h = 0; 1; \dots; H-2$, we have

$$\mathbb{E} [g(s_{h+1}; a_{h+1})] \leq \mathbb{E} \left[\frac{1}{n} \sum_{h=1}^n \frac{b_h(s_h; a_h)^2}{\sum_{i=1}^d \sum_{h=1}^n b_h(s_h; a_h)^2} \right] \frac{n \mathbb{E}_{\pi} [g(s_{h+1}; a_{h+1})^2] + d \text{kgk}_1^2}{\lambda_{\min}(\Sigma_h)}.$$

Proof.

$$\begin{aligned} \mathbb{E} [g(s_{h+1}; a_{h+1})] &= \mathbb{E} \left[\sum_{i=1}^d \frac{1}{n} \sum_{h=1}^n \frac{b_h(s_h; a_h)^2}{\sum_{i=1}^d \sum_{h=1}^n b_h(s_h; a_h)^2} g(s_{h+1}; a_{h+1}) \right] \\ &\leq \mathbb{E} \left[\sum_{i=1}^d \frac{1}{n} \sum_{h=1}^n \frac{b_h(s_h; a_h)^2}{\sum_{i=1}^d \sum_{h=1}^n b_h(s_h; a_h)^2} \right] \frac{n \mathbb{E}_{\pi} [g(s_{h+1}; a_{h+1})^2] + d \text{kgk}_1^2}{\lambda_{\min}(\Sigma_h)} \end{aligned}$$

where

$$\begin{aligned} &\sum_{i=1}^d \frac{1}{n} \sum_{h=1}^n \frac{b_h(s_h; a_h)^2}{\sum_{i=1}^d \sum_{h=1}^n b_h(s_h; a_h)^2} \\ &= n \mathbb{E}_{\pi} \left[\sum_{i=1}^d \frac{1}{n} \sum_{h=1}^n \frac{b_h(s_h; a_h)^2}{\sum_{i=1}^d \sum_{h=1}^n b_h(s_h; a_h)^2} \right] \\ &\leq n \mathbb{E}_{\pi} \left[\sum_{i=1}^d \frac{1}{n} \sum_{h=1}^n \frac{b_h(s_h; a_h)^2}{\sum_{i=1}^d \sum_{h=1}^n b_h(s_h; a_h)^2} \right] + d \text{kgk}_1^2 \end{aligned}$$

\square

Under reachability, we can show that small (squared) bonuses and spectral coverage, in the sense of having lower bounded eigenvalues, are somewhat equivalent.

Lemma F.3 Let Σ_h be a symmetric positive definite matrix and define the bonus $b_h(s; a) = \frac{1}{\lambda_{\min}(\Sigma_h)} \sum_{i=1}^d \frac{1}{n} \sum_{h=1}^n \frac{b_h(s_h; a_h)^2}{\sum_{i=1}^d \sum_{h=1}^n b_h(s_h; a_h)^2}$. Then we have

1. For any policy π , $\mathbb{E}_{d_h} [b_h^2(s; a)] \geq \frac{1}{\min(\lambda_{\min}(\Sigma_h))}$. That is, coverage implies small squared bonus.
2. Suppose reachability under π (Assumption 3.3), then we have the converse: there exists $\epsilon > 0$ such that for any policy π , $\mathbb{E}_{d_h} [b_h^2(s; a)] \geq \frac{\epsilon}{\min(\lambda_{\min}(\Sigma_h))}$. That is, small squared bonus implies coverage.

Proof. The first claim follows directly from Cauchy-Schwartz. Indeed, for any policy we have

$$\mathbb{E}_{d_h} b_h^2(s; a) \leq \mathbb{E}_{d_h} \|k_h^?(s; a)\|_2^2 \leq \|k_h^?\|_2^2 = \frac{1}{\min(\cdot)}.$$

For the second claim, [Assumption 3.3](#) implies that there exist a policy such that for all vectors $v \in \mathbb{R}^d$ with $\|v\|_2 = 1$, we have $\mathbb{E}_{d_h} (k_h^?(s; a)^\top v)^2 \geq \frac{1}{d} \sum_{i=1}^d v_i^2$. Now decompose $v = \sum_{i=1}^d v_i v_i^\top$, where $v_i; v_i$ are eigenvalue/vector pairs with $\|v_i\|_2 = 1$ and $v_1^\top v_2^\top \dots v_d^\top = d$. Then substituting this into the definition of the bonus, we have

$$\begin{aligned} \mathbb{E}_{d_h} b_h^2(s; a) &= \sum_{i=1}^d \frac{1}{d} \mathbb{E}_{d_h} (k_h^?(s; a)^\top v_i)^2 \\ &= \frac{1}{d} \mathbb{E}_{d_h} (k_h^?(s; a)^\top v_d)^2 \\ &\geq \frac{1}{\min(\cdot)}. \end{aligned}$$

□

We now prove our main lemma for reward-free exploration, [Lemma 4.2](#).

Lemma 4.2 (Reward-free Exploration) Fix any source task $k \in [K]$. Suppose [Assumptions 3.2](#) and [3.3](#). Then, for any $\epsilon \in (0, 1)$, w.p. $1 - \epsilon$, REWARDFREE ([Algorithm 2](#)) with $N_{\text{LSVI-UCB}} = \lceil e A^3 d^6 H^8 \epsilon^{-2} \rceil$ and $N_{\text{REWARDFREE}} = \lceil e A^3 d^4 H^6 \log \frac{1}{\epsilon} \rceil N_{\text{LSVI-UCB}}^2$ returns a ϵ -exploratory policy k where $\epsilon = e^{-1} A^{-3} d^{-5} H^{-7} \epsilon^{-2}$. The sample complexity here is $N_{\text{REWARDFREE}}$ episodes in the source task.

Proof of Lemma 4.2 In this proof, let

$$\begin{aligned} \tilde{b}_h &= N_{\text{LSVI-UCB}} \mathbb{E}_{h-1}^{+1} \left(b_h^?(s_h; a_h) - b_h^?(s_h; a_h)^\top \right)^2 + I; \\ b_h &= N_{\text{LSVI-UCB}} \mathbb{E}_{h-1}^{+1} \left(b_h(s_h; a_h) - b_h(s_h; a_h)^\top \right)^2 + I; \end{aligned}$$

where $\tilde{b}_h = dH \log(N_{\text{LSVI-UCB}}) \geq 1$. This setting of \tilde{b}_h satisfies the precondition for the Concentration of Inverse Covariances [Zanette et al. \(2021, Lemma 39\)](#), which implies w.p. at least that

$$b_h^{-1} \leq 2 \sum_{e=1}^{N_{\text{LSVI-UCB}}} \left(b_h(s_h^e; a_h^e) - b_h(s_h^e; a_h^e)^\top \right)^2 + I \leq 2 \frac{1}{N_{\text{LSVI-UCB}}};$$

where we've also used the fact that $\tilde{b}_h \geq 1$, so $(A + I)^{-1} \leq (A + \tilde{b}_h I)^{-1}$.

Under this event, for any, we have,

$$\sum_{h=1}^H \mathbb{E}_{\tilde{p}} \left(b_h(s_h; a_h) - b_h(s_h; a_h)^\top \right)^2 \leq \sum_{h=1}^H \mathbb{E}_{\tilde{p}} \left(b_h(s_h; a_h) - b_h(s_h; a_h)^\top \right)^2 \leq \frac{1}{N_{\text{LSVI-UCB}}}. \quad (6)$$

Now let $h = 0; 1; \dots; H - 2$ be arbitrary. By [Assumption 3.3](#) (there exists some policy with coverage) such that,

$$\begin{aligned} \frac{1}{\min(\cdot)} &\leq \mathbb{E}_e \left(b_{h+1}^?(s_{h+1}; a_{h+1}) - b_{h+1}^?(s_{h+1}; a_{h+1})^\top \right)^2 && \text{(by Lemma F.3)} \\ \mathbb{E}_e \left(b_{h+1}^?(s_{h+1}; a_{h+1}) - b_{h+1}^?(s_{h+1}; a_{h+1})^\top \right)^2 &\leq \frac{1}{\min(\cdot)} && \text{(by (1))} \\ \mathbb{E}_{e, \tilde{p}} \left(b_h(s_h; a_h) - b_h(s_h; a_h)^\top \right)^2 &\leq \frac{1}{A(2d + \epsilon_{\text{TV}} N_{\text{LSVI-UCB}}) + \epsilon_{\text{TV}}} && \text{(by Corollary F.1)} \end{aligned}$$

$$\begin{aligned}
 & \mathbb{E}_{\pi} \left[\sum_{h=1}^H \sum_{s \in \mathcal{S}} b_h(s; a) \right] \leq \frac{C}{A^{1.5} d^2 H^3} \frac{1}{\log(dHN_{\text{LSVI-UCB}})} + 1 = N_{\text{LSVI-UCB}} \quad (\text{by } \|\cdot\|_{\text{TV}} = 1 = N_{\text{LSVI-UCB}}) \\
 & \cdot A^{1.5} d^2 H^3 \frac{1}{\log(dHN_{\text{LSVI-UCB}})} = N_{\text{LSVI-UCB}} + 1 = N_{\text{LSVI-UCB}} \quad (\text{by (6) and Lemma F.1}) \\
 & \cdot A^{1.5} d^2 H^3 \frac{1}{\log(dHN_{\text{LSVI-UCB}})} = N_{\text{LSVI-UCB}}
 \end{aligned}$$

Recall that $\sum_{h=1}^H \sum_{s \in \mathcal{S}} b_h(s; a) = dH \log(N_{\text{LSVI-UCB}})$, we have,

$$\begin{aligned}
 & \min_{\pi} \mathbb{E}_{\pi} \left[\sum_{h=1}^H \sum_{s \in \mathcal{S}} b_{h+1}(s; a) \right] \leq \frac{C}{A^{1.5} d^2 H^3} \frac{1}{\log(dHN_{\text{LSVI-UCB}})} + 1 \\
 & = \frac{1}{N_{\text{LSVI-UCB}}} \frac{C}{A^{1.5} d^2 H^3} \frac{1}{\log(dHN_{\text{LSVI-UCB}})} + 1 \\
 & \quad \& \frac{C}{A^{1.5} d^2 H^3} \frac{1}{N_{\text{LSVI-UCB}}} \frac{dH}{N_{\text{LSVI-UCB}}};
 \end{aligned}$$

where we've omitted the log terms for simplicity in the &. Now we optimize $N_{\text{LSVI-UCB}}$ to maximize this bound. For $a, b > 0$, to maximize a function of the form $f(x) = \frac{a}{x} + \frac{b}{x}$, it's best to set x such that $\frac{a}{x^2} = \frac{2b}{x}$, resulting in value $f(x^*) = \frac{a^2}{4b}$. Setting,

$$\begin{aligned}
 x &= N_{\text{LSVI-UCB}}; \\
 a &= \frac{C}{A^{1.5} d^2 H^3}; \\
 b &= dH;
 \end{aligned}$$

Hence, we need to set

$$N_{\text{LSVI-UCB}} = e \frac{A^3 d^6 H^8}{2};$$

which results in a \min lower bound of

$$\min_{\pi} \mathbb{E}_{\pi} \left[\sum_{h=1}^H \sum_{s \in \mathcal{S}} b_{h+1}(s; a) \right] \geq e \frac{a^2}{b} = e \frac{C^2}{A^3 d^5 H^7};$$

Finally, we used the fact that $\|\cdot\|_{\text{TV}} = 1 = N_{\text{LSVI-UCB}}$, which is set by the choice of $N_{\text{REWARDFREE}}$ in the lemma statement to satisfy (5).

The above proves coverage of π for $h = 0; 1; \dots; H-2$. Finally to argue for $h=H-1$, which is simply taking a random action at time t , we can simply invoke Assumption 3.3 for $t=0$ to get a policy that

$$\mathbb{E}_{\pi} \left[\sum_{s \in \mathcal{S}} b_0(s; a) \right] \geq \frac{1}{A} \mathbb{E}_{\pi} \left[\sum_{s \in \mathcal{S}} b_0(s; a) \right] \geq \frac{1}{A};$$

□

Corollary F.1 Let $\{b_h\}_{h=1}^H$ be defined as in the proof of Lemma 4.2. For any $h = 0; 1; \dots; H-2$ and any policy π , we have

$$\mathbb{E}_{\pi} \left[\sum_{h=1}^H \sum_{s \in \mathcal{S}} b_h(s; a) \right] \geq \frac{1}{A} \mathbb{E}_{\pi} \left[\sum_{h=1}^H \sum_{s \in \mathcal{S}} b_h(s; a) \right] \geq \frac{1}{A} \frac{C}{A^{1.5} d^2 H^3} \frac{1}{\log(dHN_{\text{LSVI-UCB}})} + 1;$$

Intuitively, this means that coverage in the learned features implies coverage in the true features.

Proof. For shorthand, let $N = N_{\text{LSVI-UCB}}$. Apply Lemma F.2 (one-step back) to the learned model and the function $f(s; a) = \frac{1}{k_{h+1}} \sum_{s \in \mathcal{S}} b_{h+1}(s; a)$, which is bounded by $\frac{1}{k_{h+1}} \leq 1$. We have,

$$\begin{aligned}
 & \mathbb{E}_{\pi} \left[\sum_{h=1}^H \sum_{s \in \mathcal{S}} b_{h+1}(s; a) \right] \geq \frac{1}{A} \mathbb{E}_{\pi} \left[\sum_{h=1}^H \sum_{s \in \mathcal{S}} b_{h+1}(s; a) \right] \geq \frac{1}{A} \frac{C}{A^{1.5} d^2 H^3} \frac{1}{\log(dHN_{\text{LSVI-UCB}})} + 1 \\
 & \mathbb{E}_{\pi} \left[\sum_{h=1}^H \sum_{s \in \mathcal{S}} b_h(s; a) \right] \geq \frac{1}{A} \mathbb{E}_{\pi} \left[\sum_{h=1}^H \sum_{s \in \mathcal{S}} b_h(s; a) \right] \geq \frac{1}{A} \frac{C}{A^{1.5} d^2 H^3} \frac{1}{\log(dHN_{\text{LSVI-UCB}})} + 1 + d
 \end{aligned}$$

$$\begin{aligned} & \mathbb{E}_{\mathbf{p}; \mathbf{k}} \mathbb{E}_{\mathbf{b}_h(\mathbf{s}_h; \mathbf{a}_h)} \mathbb{E}_{\mathbf{b}_{h+1}} \frac{r}{N \sum_j A_j \mathbb{E}_{\mathbf{k}_{h+1}} \mathbb{E}_{\mathbf{s}_{h+1}; \mathbf{a}_{h+1}} \mathbb{E}_{\mathbf{k}_{h+1}}^2 + N \sum_j A_j \|\mathbf{T}_V\| + d} \\ & \mathbb{E}_{\mathbf{p}; \mathbf{k}} \mathbb{E}_{\mathbf{b}_h(\mathbf{s}_h; \mathbf{a}_h)} \mathbb{E}_{\mathbf{b}_{h+1}} \frac{p}{d \sum_j A_j + N \sum_j A_j \|\mathbf{T}_V\| + d}; \end{aligned}$$

where we used the fact that

$$\begin{aligned} & \mathbb{E}_{\mathbf{k}_{h+1}} \mathbb{E}_{\mathbf{s}_{h+1}; \mathbf{a}_{h+1}} \mathbb{E}_{\mathbf{k}_{h+1}}^2 \mathbb{E}_{\mathbf{s}_{h+1}; \mathbf{a}_{h+1}} \\ &= \text{Tr} \left(\mathbb{E}_{\mathbf{k}_{h+1}} \mathbb{E}_{\mathbf{s}_{h+1}; \mathbf{a}_{h+1}} \mathbb{E}_{\mathbf{k}_{h+1}} \mathbb{E}_{\mathbf{s}_{h+1}; \mathbf{a}_{h+1}}^\top \right) = \text{Tr} \left(\mathbb{E}_{\mathbf{k}_{h+1}} \mathbb{E}_{\mathbf{s}_{h+1}; \mathbf{a}_{h+1}} \mathbb{E}_{\mathbf{k}_{h+1}}^\top \mathbb{E}_{\mathbf{s}_{h+1}; \mathbf{a}_{h+1}} \right) + \mathbf{I} \\ &= \frac{1}{N} \text{Tr}(\mathbf{I} + \mathbf{M}) \frac{d}{N}; \end{aligned}$$

where \mathbf{M} is a positive definite matrix. Thus, doing an initial change from \mathbf{p}_{h+1} to \mathbf{p}_{h+1} concludes the proof. \square

G REPRESENTATION TRANSFER

First, we prove [Lemma 4.1](#), restated below.

Lemma 4.1 Suppose [Assumption 3.4](#) and that for all source tasks $\mathcal{K} \subseteq [K]$ we have a \min -exploratory policy $\pi_{\mathcal{K}}$. Then, for any $\epsilon \in (0, 1)$, learning features \mathbf{b} using the cross-sampling procedure of [Algorithm 1](#) satisfies [\(2\)](#) w.p. $1 - \epsilon$. Furthermore, for any $h = 0, 1, \dots, H - 1$, there exist $\mathbf{e}_h : \mathcal{S} \rightarrow \mathbb{R}^d$ such that for any function $g : \mathcal{S} \rightarrow [0, 1]$, $\| \mathbb{E}_{\mathbf{p}; \mathbf{k}} \mathbb{E}_{\mathbf{b}_h(\mathbf{s}_h; \mathbf{a}_h)} \mathbb{E}_{\mathbf{e}_h} \mathbb{E}_{\mathbf{k}; h} \mathbb{E}_{\mathbf{s}_h; \mathbf{a}_h} g(\mathbf{s}_h) d(\mathbf{s}_h) - \mathbb{E}_{\mathbf{p}; \mathbf{k}} \mathbb{E}_{\mathbf{b}_{k;h}} \mathbb{E}_{\mathbf{s}_h; \mathbf{a}_h} g(\mathbf{s}_h) d(\mathbf{s}_h) \|_2 \leq \frac{\epsilon}{\sqrt{N}}$ and

$$\sup_{\mathbf{p}; \mathbf{k}} \mathbb{E}_{\mathbf{b}_h(\mathbf{s}_h; \mathbf{a}_h)} \mathbb{E}_{\mathbf{e}_h} \mathbb{E}_{\mathbf{k}; h} \mathbb{E}_{\mathbf{s}_h; \mathbf{a}_h} g(\mathbf{s}_h) d(\mathbf{s}_h) - \mathbb{E}_{\mathbf{p}; \mathbf{k}} \mathbb{E}_{\mathbf{b}_{k;h}} \mathbb{E}_{\mathbf{s}_h; \mathbf{a}_h} g(\mathbf{s}_h) d(\mathbf{s}_h) \leq \frac{\epsilon}{\sqrt{N}}$$

Proof of Lemma 4.1 Fix an arbitrary ϵ . Denote $\mathbf{e}_h(\mathbf{s}^0) = \frac{1}{K} \sum_{k=0}^{K-1} \mathbf{b}_{k;h}(\mathbf{s}^0)$. First, note that

$$\begin{aligned} & \max_{g: \mathcal{S} \rightarrow [0, 1]} \mathbb{E}_{\mathbf{p}; \mathbf{k}} \mathbb{E}_{\mathbf{b}_h(\mathbf{s}_h; \mathbf{a}_h)} \mathbb{E}_{\mathbf{e}_h} \mathbb{E}_{\mathbf{k}; h} \mathbb{E}_{\mathbf{s}_h; \mathbf{a}_h} g(\mathbf{s}_h) d(\mathbf{s}_h) \\ &= \max_{g: \mathcal{S} \rightarrow [0, 1]} \mathbb{E}_{\mathbf{p}; \mathbf{k}} \mathbb{E}_{\mathbf{b}_{k;h}} \mathbb{E}_{\mathbf{s}_h; \mathbf{a}_h} g(\mathbf{s}_h) d(\mathbf{s}_h) \\ &= \max_{\mathbf{s}} \mathbb{E}_{\mathbf{p}; \mathbf{k}} \mathbb{E}_{\mathbf{b}_{k;h}} \mathbb{E}_{\mathbf{s}_h; \mathbf{a}_h} g(\mathbf{s}_h) d(\mathbf{s}_h) \quad (\text{Since } \mathbb{E}_{\mathbf{p}; \mathbf{k}} \mathbb{E}_{\mathbf{b}_{k;h}} \mathbb{E}_{\mathbf{s}_h; \mathbf{a}_h} g(\mathbf{s}_h) d(\mathbf{s}_h) \text{ is } \mathbf{p}\text{-} \bar{d} \text{ by 3.2}) \\ &= \frac{\epsilon}{\sqrt{N}} \end{aligned}$$

For any $h = 0, 1, \dots, H - 1$, we have

$$\begin{aligned} & \mathbb{E}_{\mathbf{p}; \mathbf{k}} \mathbb{E}_{\mathbf{b}_h(\mathbf{s}_h; \mathbf{a}_h)} \mathbb{E}_{\mathbf{e}_h} \mathbb{E}_{\mathbf{k}; h} \mathbb{E}_{\mathbf{s}_h; \mathbf{a}_h} g(\mathbf{s}_h) d(\mathbf{s}_h) \\ &= \mathbb{E}_{\mathbf{p}; \mathbf{k}} \mathbb{E}_{\mathbf{b}_h(\mathbf{s}_h; \mathbf{a}_h)} \mathbb{E}_{\mathbf{e}_h} \mathbb{E}_{\mathbf{k}; h} \mathbb{E}_{\mathbf{s}_{h+1}; \mathbf{a}_{h+1}} \mathbb{E}_{\mathbf{b}_{k;h}(\mathbf{s}_{h+1})} \mathbb{E}_{\mathbf{s}_h; \mathbf{a}_h} g(\mathbf{s}_h) d(\mathbf{s}_h) \\ &= \mathbb{E}_{\mathbf{p}; \mathbf{k}} \mathbb{E}_{\mathbf{b}_h(\mathbf{s}_h; \mathbf{a}_h)} \mathbb{E}_{\mathbf{e}_h} \mathbb{E}_{\mathbf{k}; h} \mathbb{E}_{\mathbf{s}_{h+1}; \mathbf{a}_{h+1}} \mathbb{E}_{\mathbf{b}_{k;h}(\mathbf{s}_{h+1})} \mathbb{E}_{\mathbf{s}_h; \mathbf{a}_h} g(\mathbf{s}_h) d(\mathbf{s}_h) \\ &= \max_{\mathbf{k}=1} \mathbb{E}_{\mathbf{p}; \mathbf{k}} \mathbb{E}_{\mathbf{b}_h(\mathbf{s}_h; \mathbf{a}_h)} \mathbb{E}_{\mathbf{e}_h} \mathbb{E}_{\mathbf{k}; h} \mathbb{E}_{\mathbf{s}_h; \mathbf{a}_h} g(\mathbf{s}_h) d(\mathbf{s}_h) \end{aligned}$$

First consider the case when $h = 0$. At $h = 0$, the distribution under $\mathbf{p}_{\mathbf{k}}$ is the same as $\mathbf{p}_{k;h}$, and so, we directly get that the above quantity is at most $\frac{\epsilon}{\sqrt{N}}$, which proves the $h = 0$ case.

Now consider any $h = 1; 2; \dots; H - 1$. To simplify notation, let us denote

$$\begin{aligned} \text{err}_{k;h}(s_h; a_h) &= b_h(s_h; a_h) - b_{k;h}(s_h; a_h) - \int_{\mathcal{Z}} b_h(s_h; a_h) - b_{k;h}(s_h; a_h) \text{TV}; \\ w_{k;h} &= \int_{\mathcal{S}_h} d_{K;h-1}(s_h) E_{a_h \sim p_h(s_h)} \text{err}_{k;h}(s_h; a_h); \\ \mu_{k;h} &= E_{k;P_k} \int_{\mathcal{S}_h} d_{K;h-1}(s_h) \int_{\mathcal{A}_h} b_h(s_h; a_h) \text{TV}; \end{aligned}$$

Note that $\min_k (\mu_{k;h})$ by assumption. Now continuing from where we left off, we take a one-step back as follows,

$$\begin{aligned} & \max_{k=1}^{K-1} E_{P_K} \text{err}_{k;h}(s_h; a_h) \\ = & \max_{k=1}^{K-1} E_{P_K} \int_{\mathcal{S}_h} d_{K;h-1}(s_h; a_h); w_{k;h} \\ & \max_{k=1}^{K-1} E_{P_K} \int_{\mathcal{S}_h} d_{K;h-1}(s_h; a_h) \mu_{k;h-1} kw_{k;h} \mu_{k;h-1} \end{aligned}$$

By $\mu_{k;h}$ guarantee of $\mu_{k;h}$, and Jensen's inequality to push the square inside,

$$\begin{aligned} & \frac{p}{\min_{k=1}^{K-1}} \max_{k=1}^{K-1} q \frac{E_{\mathcal{S}_h; a_h} \int_{\mathcal{S}_h} d_{K;h-1}(s_h; a_h); a_h \sim p_h(s_h) \text{err}_{k;h}(s_h; a_h)^2}{E_{\mathcal{S}_h; a_h} \int_{\mathcal{S}_h} d_{K;h-1}(s_h; a_h); a_h \sim p_h(s_h) \text{err}_{k;h}(s_h; a_h)^2} \\ & \frac{A^{1=2}}{p} \frac{\max_{k=1}^{K-1} q}{\min_{k=1}^{K-1}} \frac{E_{\mathcal{S}_h; a_h} \int_{\mathcal{S}_h} d_{K;h-1}(s_h; a_h); a_h \sim p_h(s_h) \text{err}_{k;h}(s_h; a_h)^2}{E_{\mathcal{S}_h; a_h} \int_{\mathcal{S}_h} d_{K;h-1}(s_h; a_h); a_h \sim p_h(s_h) \text{err}_{k;h}(s_h; a_h)^2} \end{aligned}$$

By Assumption 3.4, the expectation over $\mathcal{P}_{k;h-1}$ is a linear combination of expectations over $\mathcal{P}_{j;h-1}$,

$$\begin{aligned} & \frac{A^{1=2}}{p} \frac{\max_{k=1}^{K-1} q}{\min_{k=1}^{K-1}} \frac{\sum_{j=1}^U \int_{\mathcal{S}_h} d_{K;h-1}(s_h; a_h); a_h \sim p_h(s_h) \text{err}_{k;h}(s_h; a_h)^2}{\sum_{j=1}^U \int_{\mathcal{S}_h} d_{K;h-1}(s_h; a_h); a_h \sim p_h(s_h) \text{err}_{k;h}(s_h; a_h)^2} \\ & \frac{A^{1=2}}{p} \frac{\max_{k=1}^{K-1} q}{\min_{k=1}^{K-1}} \frac{\sum_{j=1}^U \int_{\mathcal{S}_h} d_{K;h-1}(s_h; a_h); a_h \sim p_h(s_h) \text{err}_{k;h}(s_h; a_h)^2}{\sum_{j=1}^U \int_{\mathcal{S}_h} d_{K;h-1}(s_h; a_h); a_h \sim p_h(s_h) \text{err}_{k;h}(s_h; a_h)^2} \\ & \frac{A^{1=2}}{p} \frac{\max_{k=1}^{K-1} q}{\min_{k=1}^{K-1}} \frac{\sum_{j=1}^U \int_{\mathcal{S}_h} d_{K;h-1}(s_h; a_h); a_h \sim p_h(s_h) \text{err}_{k;h}(s_h; a_h)^2}{\sum_{j=1}^U \int_{\mathcal{S}_h} d_{K;h-1}(s_h; a_h); a_h \sim p_h(s_h) \text{err}_{k;h}(s_h; a_h)^2}; \end{aligned}$$

where we used the MLE guarantee (2) in the last step. □

Algorithm 5 Transfer learning with online access

PRE-TRAINING PHASE

Input: num. LSVI-UCB episodes $N_{\text{LSVI-UCB}}$, num. model-learning episodes $N_{\text{REWARDFREE}}$, size of cross-sampled datasets, failure probability δ .

- 1: for source task $k = 1; \dots; K - 1$ do
- 2: Find policy cover $\pi_k = \text{REWARDFREE}(P_k^?; N_{\text{LSVI-UCB}}; N_{\text{REWARDFREE}}; \delta)$. (Algorithm 2)
- 3: for source task $k = 1; \dots; K - 1$ do
- 4: For each $h = 0; 1; \dots; H - 1$, sample \mathcal{D}_k as n i.i.d. $(s_h; a_h; s_{h+1})$ tuples from π_k .
- 5: For each $h = 0; 1; \dots; H - 1$, learn $b_h = \text{Multi-task REPLEARN}(f_{k;h}, g_{k2[K-1]})$. (Algorithm 3)

DEPLOYMENT PHASE

Additional Input: number of deployment episodes \bar{s}

- 1: Set $\bar{d} = H \bar{s} + dH \frac{p}{\log(dHT)}$.
- 2: Run LSVI-UCB $f_{h=0}^{H-1}; r = r_K; T; \bar{d}$ in the target task $\mathcal{P}_K^?$ (Algorithm 6).

Next we state an analogous lemma for when we don't need generative access to the source task, but instead assume Assumption 5.1, and Assumption 5.2.

Lemma G.1 Suppose [Assumption 5.1](#), and [Assumption 5.2](#). Now take the setup of [Lemma 4.1](#) with the only difference being that it is learned as in [Algorithm 5](#). Then, the same guarantee of [Lemma 4.1](#) holds with a slightly different right hand side for the bound on the TV-error,

$$\sup_{\mathbb{P}_K} \mathbb{E} [b_h(s_h; a_h) - b_h(\cdot)] - \mathbb{E} [b_h(s_h; a_h) - b_{k;h}(\cdot)]_{TV} \leq \frac{\max_{k=1,2} K \frac{1-2}{n}}{(\text{raw min})^{1-2}};$$

Proof of [Lemma G.1](#) Fix an arbitrary \cdot . Denote $b_h(s^0) = \prod_{k=0}^{K-1} b_{k;h}(s^0) \wedge_{k;h}(s^0)$. Then, some algebra with importance sampling gives us the bound,

$$\mathbb{E}_{\mathbb{P}_K} [b_h(s_h; a_h) - b_h(\cdot)] - \mathbb{E} [b_h(s_h; a_h) - b_{k;h}(\cdot)]_{TV} \leq 3$$

$$\mathbb{E}_{\mathbb{P}_K} [b_h(s_h; a_h) - b_{k;h}(s_{h+1})] - \mathbb{E} [b_h(s_h; a_h) - b_{k;h}(s_{h+1})]_{TV} \leq 5$$

$$\max_{k=1}^{K-1} \mathbb{E}_{\mathbb{P}_K} [b_h(s_h; a_h) - b_{k;h}(\cdot)] - \mathbb{E} [b_h(s_h; a_h) - b_{k;h}(\cdot)]_{TV}$$

$$\max_{k=1}^{K-1} \frac{\mathbb{E}_{\mathbb{P}_K} [b_h(s_h; a_h) - b_{k;h}(\cdot)] - \mathbb{E} [b_h(s_h; a_h) - b_{k;h}(\cdot)]_{TV}^2}{\mathbb{E}_{\mathbb{P}_K} [b_h(s_h; a_h) - b_{k;h}(\cdot)] - \mathbb{E} [b_h(s_h; a_h) - b_{k;h}(\cdot)]_{TV}}$$

By [Assumption 5.1](#), [Assumption 5.2](#), for any $s; a$, we have $\frac{d_{K;h}(s;a)}{d_{k;h}(s;a)}$ $\frac{1}{\text{raw min}}$, where we used the coverage-under- ϵ -assumption in the last inequality. In other words, for each $k=1,2,\dots,K-1$, we have

$$\frac{d_{K;h}}{d_{k;h}} \leq \frac{1}{\text{raw min}}, \text{ hence we can importance sample,}$$

$$\frac{\max_{k=1}^{K-1} \mathbb{E}_{\mathbb{P}_K} [b_h(s_h; a_h) - b_{k;h}(\cdot)] - \mathbb{E} [b_h(s_h; a_h) - b_{k;h}(\cdot)]_{TV}^2}{(\text{raw min})^{1-2}}$$

$$\frac{\max_{k=1}^{K-1} K \frac{1-2}{n}}{(\text{raw min})^{1-2}};$$

□

H PROOFS FOR DEPLOYMENT PHASE

H.1 ONLINE RL LEMMAS

Lemma H.1 (Self-normalized Martingale) Consider iterations $F_i, g_{i=1,2,\dots}$, so that $\mathbb{E}[g_i | F_{i-1}] = 0$ and $\{g_i | F_{i-1}, g_{i=1,2,\dots}\}$ are sub-Gaussian with parameter σ . Let $\{X_i | g_{i=1,2,\dots}\}$ be random variables in a Hilbert space \mathcal{H} . Suppose a linear operator $\Sigma_0 : \mathcal{H} \rightarrow \mathcal{H}$ is positive definite. For any t , denote $\Sigma_t = \Sigma_0 + \sum_{i=1}^t X_i X_i^T$. Then w.p. at least $1 - \delta$, we have,

$$\| \Sigma_t^{-1} \sum_{i=1}^t X_i \|^2 \leq 2 \log \frac{\det(\Sigma_t) \det(\Sigma_0)^{-1}}{\delta};$$

Proof. Lemma A.8 of ([Agarwal et al., 2019](#)). □

Lemma H.2 Let $\Sigma_t = I + \sum_{i=1}^t x_i x_i^T$ for $x_i \in \mathbb{R}^d$ and $\|x_i\| \leq 1$. Then $\Sigma_t^{-1} \sum_{i=1}^t x_i \leq d$.

Proof. Lemma D.1 of ([Jin et al., 2020b](#)). □

H.2 PROOF OF MAIN RESULTS

Let $(x)_y$ refer to the clamping operator, i.e. $(x)_y = \min\{x, y\}$. Let M_V be the maximum possible value in the MDP with the given reward function.

Algorithm 6 LSVI-UCB

- 1: Input: Features $\mathbf{b}_h, \mathbf{g}_{h=0;1;\dots;H-1}$, reward $r_h, \mathbf{g}_{h=0;1;\dots;H-1}$, number of episodes N , bonus scaling parameter, UNIFORMACTIONS = FALSE.
 - 2: for episode $e = 1; 2; \dots; N$ do
 - 3: Initialize $\mathbf{v}_{h,e}(s) = 0; \forall s$
 - 4: for step $h = H-1; H-2; \dots; 0$ do
 - 5: Learn best predictor for \mathbf{b}_{h+1}^e ,

$$\mathbf{x}_{h,e}^1 = \arg \min_{\mathbf{x}} \sum_{k=1}^K \mathbf{b}_h(s_h^k; a_h^k) \mathbf{b}_h(s_h^k; a_h^k)^\top \mathbf{x} + I;$$

$$\mathbf{w}_{h,e}^1 = \arg \min_{\mathbf{w}} \sum_{k=1}^K \mathbf{b}_h(s_h^k; a_h^k) \mathbf{v}_{h+1,e}(s_{h+1}^k);$$
 - 6: Set bonus and value functions,

$$\mathbf{b}_{h,e}(s; a) = \mathbf{b}_h(s; a)_{h,e}^1;$$

$$\mathbf{Q}_{h,e}(s; a) = \mathbf{w}_{h,e}^1 \sum_{h=0}^H \mathbf{b}_h(s; a) + r_h(s; a) + \mathbf{v}_{h,e}(s; a);$$

$$\mathbf{v}_{h,e}(s) = \max_a \mathbf{Q}_{h,e}(s; a)_{M_V};$$
 - 7: Set $\pi_h^e(s) = \arg \max_a \mathbf{Q}_{h,e}(s; a)$.
 - 8: Execute π^e to collect a trajectory $(s_h^e; a_h^e)_{h=0}^{H-1}$.
 - 9: If UNIFORMACTIONS = TRUE, discard a_h^e and draw freshly sampled uniform actions independently for all h , i.e. $a_h^e \sim \text{Unif}(A)$.
 - 10: Return: uniform mixture $\pi = \text{Uniform}(\pi_{e=1}^N)$.
-

Previously, [Jin et al. \(2020b\)](#) analyzed LSVI-UCB under point-wise model-misspecification. Here, we show that similar guarantees hold under a more general ρ -distribution model-misspecification "ms", captured by [Assumption H.1](#).

Assumption H.1 Suppose for every $h = 0; 1; \dots; H-1$, there exists \mathbf{I}_h such that for any policy π ,

$$\mathbb{E} \left[\mathbf{e}_h(\pi)^\top \mathbf{b}_h(s_h; a_h) \right] \leq \mathbf{I}_h \left(\sum_{j=1}^K \mathbb{P}_h^{\pi}(j | s_h; a_h) \right)_{TV} \text{ "ms"}$$

We further assume that $\sup_{s,a,h} \mathbf{e}_h(\pi)^\top \mathbf{b}_h(s; a) \leq M$ and $\|\mathbf{e}_h\|_2 \leq M$ and $\|\mathbf{I}_h\|_2 \leq M$ for $h = 0; 1; \dots; H-1$. For some positive constant M .

In other words, we only need the model to be accurate on average under the occupancy distributions realizable by policies. We also make a slight generalization on the regularization constant λ which is set to 1 in the original linear MDP definition ([Jin et al., 2020b](#)). Later, we will later instantiate the above assumption with our transferred model $\mathbf{b}_{k,h}(s^0) = \sum_{k=1}^K \mathbb{P}_{k,h}(s^0) \mathbf{b}_{k,h}(s^0)$, then for any $s; a$, we have

$$\begin{aligned} \|\mathbf{e}_h(\pi)^\top \mathbf{b}_h(s; a)\|_{TV} &= \sum_{s^0} \sum_{k=1}^K \mathbb{P}_{k,h}(s^0) \|\mathbf{b}_{k,h}(s^0)^\top \mathbf{b}_h(s; a)\|_{TV} \\ &= \sum_{s^0} \sum_{k=1}^K \mathbb{P}_{k,h}(s^0) \|\mathbf{b}_{k,h}(s^0)^\top \mathbf{b}_h(s; a)\|_{TV} \\ &= \sum_{k=1}^K \max_{s^0} \mathbb{P}_{k,h}(s^0) \|\mathbf{b}_{k,h}(s^0)^\top \mathbf{b}_h(s; a)\|_{TV} \quad (\text{by } \|\mathbf{b}_{k,h}(s^0)^\top \mathbf{b}_h(s; a)\|_{TV} \leq 1) \end{aligned}$$

Also,

$$\begin{aligned} \mathbf{k}^T \mathbf{e}_h \mathbf{k}_2 &= \sum_{s^0} \sum_{k=1}^{K-1} \mathbf{b}_{k;h}(s^0) \mathbf{b}_{k;h}(s^0)^T \mathbf{f}(s^0) \\ &= \sum_{k=1}^{K-1} \max_{s^0} \mathbf{b}_{k;h}(s^0) \mathbf{b}_{k;h}(s^0)^T \mathbf{f}(s^0) \\ &= \sum_{k=1}^{K-1} \mathbf{p}_{k;h}^{-1} \mathbf{k}_1 : \quad (\text{by } \mathbf{k}^T \mathbf{b}_{k;h} \mathbf{k}_2 \leq \mathbf{p}_{k;h}^{-1} \mathbf{k}_1) \end{aligned}$$

So we will set $M = \dots$

Note that we only need the existence of $\mathbf{b}_h(\cdot)^T \mathbf{b}_h(s; a)$ need not be a valid probability kernel. In fact, it may even be negative valued.

In this section, we make a model-based analysis of LSVI. Similar approaches have been used in prior works, e.g. [Lykouris et al. \(2021\)](#); [Agarwal et al. \(2019\)](#); [Zhang et al. \(2022\)](#). For simplicity, we suppose that S is finite, but may be exponentially large, as we suffer no dependence on $|S|$. The proof can be easily extended to infinite state spaces by replacing inner products with integrals.

Consider the following quantity,

$$\mathbf{b}_{h;e} = \sum_{k=1}^{K-1} (\mathbf{s}_{h+1}^k) \mathbf{b}_h(\mathbf{s}_h^k; \mathbf{a}_h^k)^T \left(\arg \min_{\mathbf{s} \in \mathcal{R}^S} \sum_{k=1}^{K-1} \mathbf{b}_h(\mathbf{s}_h^i; \mathbf{a}_h^i) (\mathbf{s}_{h+1}^i) \mathbf{k}_2 + \mathbf{k}_F^2 \right);$$

where (s) is a one-hot encoding of the state s . In words, this is the best choice for linearly (in $\mathbf{b}_h(s; a)$) predicting $E_{s^0 \sim P_h^*(s; a)}[\mathbf{f}(s^0)] = P_h^*(s^0 | s; a)$. We highlight that this is just a quantity for analysis and not computed in the algorithm. Finally, denote

$$\begin{aligned} \mathbf{b}_{h;e} &= \mathbf{b}_{h;e} \mathbf{b}_h; \\ \mathbf{p}_h &= \mathbf{e}_h \mathbf{b}_h; \end{aligned}$$

We will also sometimes use the shorthand $\mathbf{b}_h(s; a)$ for $E_{s^0 \sim P_h^*(s; a)}[\mathbf{f}(s^0)]$.

For each $h = 0; 1; \dots; H-1$, let V_h denote the class of functions

$$\mathcal{V}_h = \left\{ \max_a \mathbf{w}^T \mathbf{b}_h(s; a) + r_h(s; a) + \epsilon \|\mathbf{b}_h(s; a)\|_{M_V} \mid \mathbf{w} \in \mathbb{R}^O, \|\mathbf{w}\|_2 \leq N M_V; \epsilon \in [0; B]; \mathbf{I} \text{ symmetric} \right\}$$

The motivation behind this construction is that it satisfies the key property that all of the learned value functions $\mathbf{b}_{h;e}$ during [Algorithm 6](#) are captured in this class.

Lemma H.3 For any $h = 0; 1; \dots; H-1$,

1. $\sup_s \|\mathbf{b}_{h;e}(s)\|_{M_V} \leq M_V$.
2. For any $\epsilon = 1; 2; \dots; N$, we have $\mathbf{b}_{h;e} \in \mathcal{V}_h$.
3. If $\mathbf{v} \in \mathcal{V}_h$, we have $\sup_s \|\mathbf{v}(s)\|_{M_V} \leq M_V$.

Proof. Recall that

$$\begin{aligned} \mathbf{b}_{h;e}(s) &= \max_a \sum_{k=1}^{K-1} \mathbf{w}_{h;e}^T \mathbf{b}_h(s; a) + r_h(s; a) + \mathbf{b}_{h;e}(s; a) \\ \text{where } \mathbf{w}_{h;e} &= \sum_{k=1}^{K-1} \mathbf{b}_h(\mathbf{s}_h^k; \mathbf{a}_h^k) \mathbf{v}_{h+1;e}(\mathbf{s}_{h+1}^k); \end{aligned}$$

From the thresholding, we have

$$\|\mathbf{b}_{h;e}(s)\|_{M_V} \leq M_V$$

We can bound the norm $\|w_{h,e}\|$ as follows,

$$\|w_{h,e}\| \leq \sum_{k=1}^h \sup_s |b_{h+1,e}(s_{h+1}^k)| \leq N \sup_s |b_{h+1,e}(s)| \leq NM_V$$

We also required B , and we regularized the covariance with λ_{\min} is at least 1. Hence $\mathcal{E}_{h,e}$ satisfies all the conditions to be \mathcal{M}_h . \square

Now we control the metric entropy of \mathcal{M}_h in ℓ_1 , i.e. $d(f_1; f_2) = \sup_s |f_1(s) - f_2(s)|$ for $f_i \in \mathcal{M}_h$. Lemma H.4 Let $\epsilon > 0$ be arbitrary and let \mathcal{N}_ϵ be the smallest-net with ϵ of \mathcal{M}_h . Then,

$$\log |\mathcal{N}_\epsilon| \leq d \log(1 + 6L/\epsilon) + \log(1 + 6B/\epsilon) + d^2 \log(1 + 18B^2/\epsilon^2)$$

Proof. Let $f_1, f_2 \in \mathcal{M}_h$. Then,

$$\begin{aligned} & |f_1(s) - f_2(s)| \\ & \leq \max_a (w_1 - w_2)^T b_h(s; a) + \sum_{i=1}^h |b_h(s; a_i) - b_h(s; a_{i+1})| \\ & \leq \|w_1 - w_2\|_2 + \max_a \sum_{i=1}^h |b_h(s; a_i) - b_h(s; a_{i+1})| \\ & \leq \|w_1 - w_2\|_2 + j \sum_{i=1}^h \max_a |b_h(s; a_i) - b_h(s; a_{i+1})| \leq (\min_i \lambda_i)^{-1} \\ & \leq \|w_1 - w_2\|_2 + j \sum_{i=1}^h \frac{1}{\lambda_i}; \end{aligned}$$

where we used for any $a, b \geq 0$, we have $\frac{a-b}{a+b} \leq \frac{a}{a+b} \leq \frac{a}{b}$. Now proceeding like the Lemma 8.6 in the RL Theory Monograph (Agarwal et al., 2019), we have the result. \square

In this section, we'll use the following bonus scaling parameter,

$$:= O \left(\frac{1}{Nd^m} M_V + M_V M \frac{1}{d \log(dNM_V)} \right) \quad (7)$$

The following high probability event $\mathcal{E}_{\text{model}}$ is a key step in our proof. Essentially, Theorem H.1 guarantees that, for all functions f_h , the model we learn is an accurate predictor of the expectation, up to a bonus and some vanishing terms.

For all the following lemmas and theorems, suppose Assumption H.1 and the bonus scaling as in (7). Throughout the section, $b_h(\cdot)$ refers to indicator functions of the trajectory, where $s_h = (s_0; s_1; \dots; s_h)$. As before, the expectations $\mathbb{E}[g(s_h)]$ are with respect to the distribution of trajectories when π is executed in the environment \mathcal{M} .

Theorem H.1 Let $\lambda \in (0, 1)$. Then, w.p. $1 - \lambda$, for any time h , episode e , indicator functions $b_1; \dots; b_h$, and policy π , we have

$$\sup_{f \in \mathcal{M}_h} \mathbb{E} \left[\sum_{i=1}^h |b_{h,e}(s_h; a_h) - P_h^?(s_h; a_h) f_h(s_h)| \right] \leq \mathbb{E} \left[\sum_{i=1}^h |b_{h,e}^e(s_h; a_h) - P_h^?(s_h; a_h) f_h(s_h)| \right] + \sum_{i=1}^h \lambda_i \quad (\mathcal{E}_{\text{model}})$$

Proof. Condition on the outcome of Lemma H.5, which implies that w.p. $1 - \lambda$, for any $h; e; \pi; s_h$, we have

$$\sup_{f \in \mathcal{M}_h} \mathbb{E} \left[\sum_{i=1}^h |b_{h,e}(s_h; a_h) - P_h^?(s_h; a_h) f_h(s_h)| \right] \leq \mathbb{E} \left[\sum_{i=1}^h |b_{h,e}^e(s_h; a_h) - P_h^?(s_h; a_h) f_h(s_h)| \right]$$

Also, for any $h; e; \pi; s_h$, by Assumption H.1, we have (w.p. 1) that

$$\sup_{f \in \mathcal{M}_h} \mathbb{E} \left[\sum_{i=1}^h |P_h^?(s_h; a_h) - P_h^?(s_h; a_h) f_h(s_h)| \right] \leq \sum_{i=1}^h \lambda_i \quad \#$$

$$\mathbb{E} \sup_{f \in \mathcal{F}_h} \mathbb{P}_h(s_h; a_h) P_h^?(s_h; a_h) f$$

Combining these two yields the result, as

$$\sup_{f \in \mathcal{F}_h} \mathbb{E} \mathbb{P}_{h;e}(s_h; a_h) P_h^?(s_h; a_h) f - \sup_{f \in \mathcal{F}_h} \mathbb{E} \mathbb{P}_h(s_h; a_h) P_h^?(s_h; a_h) f + \sup_{f \in \mathcal{F}_h} \mathbb{E} \mathbb{P}_{h;e}(s_h; a_h) \mathbb{P}_h(s_h; a_h) f - \sup_{f \in \mathcal{F}_h} \mathbb{E} \mathbb{P}_h(s_h; a_h) \mathbb{P}_h(s_h; a_h) f$$

□

Lemma H.5 Suppose Assumption H.1 and the bonus scaling is set as in (7). For any $\epsilon \in (0, 1)$, w.p. at least $1 - \epsilon$, we have for any time t , episode e , and policy π ,

$$|b_{h;e}(s_h; a_h) - \mathbb{E} \mathbb{P}_{h;e}(s_h; a_h) P_h^?(s_h; a_h) f| \leq \epsilon$$

Proof. Consider any $h; e$. Define $\mathbf{H}_h^k := (s_{h+1}^k) + P_h^?(s_{h+1}^k; j_h^k; a_h^k)$, so that $\mathbb{E}[\mathbf{H}_h^k | \mathcal{H}_{k-1}] = 0$, where \mathcal{H}_{k-1} contains the states and actions before episode k . As what follows, we slightly abuse notation, as $\mathbf{P}(s; a) \mathbf{b}^T(s; a)$ will denote the outer product, and hence a $d \times d$ quantity.

$$\begin{aligned} b_{h;e} &= \sum_{k=1}^K (s_{h+1}^k) b_h(s_h^k; a_h^k)^T \\ &= \sum_{k=1}^K P_h^?(s_h^k; a_h^k) \mathbb{P}_h(s_h^k; a_h^k) b_h(s_h^k; a_h^k)^T + \sum_{k=0}^K \mathbb{P}_h(s_h^k; a_h^k) \mathbf{H}_h^k b_h(s_h^k; a_h^k)^T \\ &= \sum_{k=1}^K P_h^?(s_h^k; a_h^k) \mathbb{P}_h(s_h^k; a_h^k) b_h(s_h^k; a_h^k)^T + e_{h;e} \sum_{k=0}^K \mathbf{H}_h^k b_h(s_h^k; a_h^k)^T \end{aligned}$$

Rearranging, we have

$$b_{h;e} - e_{h;e} = \sum_{k=0}^K P_h^?(s_h^k; a_h^k) \mathbb{P}_h(s_h^k; a_h^k) b_h(s_h^k; a_h^k)^T - \sum_{k=0}^K \mathbf{H}_h^k b_h(s_h^k; a_h^k)^T$$

Now let $f \in \mathcal{F}_h$ be arbitrary. For any $s_h; a_h$, multiply the above with $b_h(s_h; a_h)$ and multiply with f , we have

$$\begin{aligned} & \mathbb{P}_{h;e}(s_h; a_h) P_h^?(s_h; a_h) f \\ &= f^T (b_{h;e} - e_{h;e}) b_h(s_h; a_h) \\ &= f^T \left\{ \sum_{k=1}^K P_h^?(s_h^k; a_h^k) \mathbb{P}_h(s_h^k; a_h^k) b_h(s_h^k; a_h^k)^T - \sum_{k=0}^K \mathbf{H}_h^k b_h(s_h^k; a_h^k)^T \right\} b_h(s_h; a_h) \\ &+ f^T e_{h;e} b_h(s_h; a_h) \\ &+ f^T \sum_{k=1}^K \mathbf{H}_h^k b_h(s_h^k; a_h^k)^T b_h(s_h; a_h) \end{aligned}$$

We can deterministically bound Term (b) as follows,

$$\begin{aligned} & \sup_{f \in \mathcal{V}_h} f^T e_{h,e} b_{h,e}(s_h; a_h) \\ &= \sup_{f \in \mathcal{V}_h} (f^T e_{h,e})^T b_{h,e}(s_h; a_h) \\ & \sup_{f \in \mathcal{V}_h} \sum_{k=1}^K f^T e_{h,e} b_{h,e}(s_h; a_h) \\ & \leq \sum_{k=1}^K \sup_{f \in \mathcal{V}_h} f^T e_{h,e} b_{h,e}(s_h; a_h) \end{aligned} \quad (\text{by Assumption H.1})$$

This term will be lower order compared to the other two.

We now derive the bound for Term (c) for any fixed $f \in \mathcal{V}_h$. Observe that

$$\begin{aligned} f^T \sum_{k=1}^K b_h(s_h^k; a_h^k) b_{h,e}(s_h; a_h) &= \sum_{k=1}^K f^T b_h(s_h^k; a_h^k) b_{h,e}(s_h; a_h) \\ &= \sum_{k=1}^K f^T b_h(s_h^k; a_h^k) (f^T b_h(s_h^k; a_h^k))^{-1} b_{h,e}(s_h; a_h) \end{aligned}$$

Now we argue w.p.1, for any $e; h$ we have

$$\sum_{k=1}^K b_h(s_h^k; a_h^k) (f^T b_h(s_h^k; a_h^k))^{-1} \leq 2kV_h k_1 \sqrt{\frac{1}{2 \log(1+\epsilon) + d \log(N+1)}};$$

which implies the claim about a_h . Indeed, we can apply Lemma H.1. Checking the preconditions, $E_{P_h^?}(s_h; a_h) f^T b_h(s_h^k; a_h^k) = 0$, $\sum_{k=1}^K f^T b_h(s_h^k; a_h^k) f^T b_h(s_h^k; a_h^k) = 2kV_h k_1$, $\det(I) = 1$, and $\det(I) = \det(I) (e+1)^d$ since the largest eigenvalue is 1. So, w.p. at least $1-\epsilon$, for all e , we have the above inequality.

Thus, for any fixed $f \in \mathcal{V}_h$, w.p. 1, for all $e; h$ we have,

$$\begin{aligned} & \sum_{h,e} b_{h,e}(s_h; a_h) f^T b_h(s_h; a_h) \\ & \leq \text{Term(a)} + \text{Term(b)} + \text{Term(c)} \\ & \leq 4kV_h k_1 \sqrt{\frac{1}{2 \log(1+\epsilon) + d \log(N)}} + \sum_{h,e} b_{h,e}(s_h; a_h) \\ & + kV_h k_1 M \sqrt{\frac{1}{d}} b_{h,e}(s_h; a_h) \\ & + 4kV_h k_1 \sqrt{\frac{1}{2 \log(1+\epsilon) + d \log(N)}} b_{h,e}(s_h; a_h) \\ & \leq \sum_{h,e} b_{h,e}(s_h; a_h) \left(\frac{1}{dN} + kV_h k_1 M \sqrt{\frac{1}{2 \log(1+\epsilon) + d \log(N)}} \right) \end{aligned}$$

Now we apply a covering argument. Namely, union bound the above argument to every element in an ϵ -net of \mathcal{V}_h . For any $f \in \mathcal{V}_h$, let f^ϵ be its neighbor in the net such that $\|f - f^\epsilon\| \leq \epsilon$, so we have

$$\sum_{h,e} b_{h,e}(s_h; a_h) f^T b_h(s_h; a_h) \leq \sum_{h,e} b_{h,e}(s_h; a_h) f^{\epsilon T} b_h(s_h; a_h) + \sum_{h,e} b_{h,e}(s_h; a_h) (f^\epsilon - f)^T b_h(s_h; a_h)$$

and

$$\sum_{h,e} b_{h,e}(s_h; a_h) f^{\epsilon T} b_h(s_h; a_h) \leq \sum_{h,e} b_{h,e}(s_h; a_h) (f^\epsilon - f)^T b_h(s_h; a_h) + \sum_{h,e} b_{h,e}(s_h; a_h) f^{\epsilon T} b_h(s_h; a_h)$$

Setting $\epsilon = \frac{1}{N}$, the metric entropy is of the order $d \log(N(M_V + B)) + \log(BN) + d^2 \log(BdN)$. The error incurred with this epsilon net is a constant, which is lower order.

Thus, we have

$$\delta_{s_h; a_h} : \sup_{f \in \mathcal{V}_h} \sum_{h,e} b_{h,e}(s_h; a_h) f^T b_h(s_h; a_h)$$

$$\begin{aligned} & \cdot \frac{P}{dNkV_hk_1} \text{ms} + kV_hk_1 M \frac{P}{\log(1=) + d\log(M_V) + d^2\log(BdN)} b_{h,e}(s_h; a_h) \\ & \cdot \frac{P}{dNM_V} \text{ms} + M_V M \frac{P}{\log(1=) + d\log(M_V) + d^2\log(BdN)} b_{h,e}(s_h; a_h) \end{aligned}$$

Note that $\frac{P}{dNkV_hk_1}$ scales as $\frac{P}{\log B}$, so one can find a valid δ by solving δ for B .

□

Lemma H.6 Let $f \in \mathcal{V}_h$. For any $\delta \in (0, 1)$, w.p. at least $1 - \delta$, for any time h , episode e , we have

$$\begin{aligned} \delta s_h; a_h : f^\top \sum_{k=1}^{\infty} P_h^?(s_h^k; a_h^k) \mathbb{E}_h(b_h(s_h^k; a_h^k)) b_h(s_h^k; a_h^k)^\top \frac{1}{h,e} b_h(s_h; a_h) \\ \leq 4kV_hk_1 (1 + M) \frac{P}{\log(1=) + d\log(N) + d^2\log(BdN)} + \frac{P}{dNkV_hk_1} \text{ms} b_{h,e}(s_h; a_h); \end{aligned}$$

Proof. First observe that

$$\begin{aligned} & f^\top \sum_{k=1}^{\infty} (P_h^?(s_h^k; a_h^k) \mathbb{E}_h(b_h(s_h^k; a_h^k)) b_h(s_h^k; a_h^k)^\top \frac{1}{h,e} b_h(s_h; a_h) \\ & = \sum_{h,e}^{1=2} \sum_{k=1}^{\infty} b_h(s_h^k; a_h^k) f^\top (P_h^?(s_h^k; a_h^k) \mathbb{E}_h(b_h(s_h^k; a_h^k)) \frac{1}{h,e} b_h(s_h; a_h) \\ & \leq \sum_{k=1}^{\infty} b_h(s_h^k; a_h^k) \mathbb{E}_k b_{h,e}(s_h; a_h); \end{aligned}$$

where $\mathbb{E}_k = P_h^?(s_h^k; a_h^k) \mathbb{E}_h(b_h(s_h^k; a_h^k)) f$.

Now we will argue that w.p.1, for all $e; h$,

$$\sum_{k=1}^{\infty} b_h(s_h^k; a_h^k) \mathbb{E}_k \leq 4kV_hk_1 (1 + M) \frac{P}{\log(1=) + d\log(N) + d^2\log(BdN)} + \frac{P}{dNkV_hk_1} \text{ms};$$

which will imply the claim for all $s_h; a_h$.

Apply self-normalized martingale concentration (Lemma H.1) to $\sum_{k=1}^{\infty} b_h(s_h^k; a_h^k) \mathbb{E}_k$ and $\mathbb{E}_k = \mathbb{E}[\mathbb{E}_k | H_{k-1}]$, where the expectation is over $(s_h^k; a_h^k)$ in the definition of \mathbb{E}_k . To see sub-Gaussianity, bound the envelope $\mathbb{E}_k \leq k f k_1 k P_h^?(s_h^k; a_h^k) \mathbb{E}_h(b_h(s_h^k; a_h^k)) k_{TV} k V_h k_1 (1 + M)$, and thus $\mathbb{E}_k \leq 2kV_hk_1 (1 + M)$. Now compute the determinant $\det(\mathbb{E}_k) = 1$ and since $\max_{h,e}(\mathbb{E}_k) \leq e + 1$, we have that $\log \det(\mathbb{E}_k) \leq d \log(e + 1)$. Hence, w.p. at least $1 - \delta$, we have

$$\delta e : \sum_{k=1}^{\infty} b_h(s_h^k; a_h^k) (\mathbb{E}_k - \mathbb{E}[\mathbb{E}_k | H_{k-1}]) \leq 2kV_hk_1 (1 + M) \frac{P}{2\log(1=) + d\log(N + 1)};$$

By Assumption H.1 applied to \mathbb{E}_k (the data-generating policy for episode e), we have $\mathbb{E}[\mathbb{E}_k | H_{k-1}] \leq \frac{kV_hk_1}{i} \text{ms}$. Recall for any scalars α_i and vectors x_i , we have $\sum_{i=1}^k \alpha_i x_i k^2 \leq \sum_{i=1}^k \alpha_i^2 k x_i k^2$. Thus,

$$\begin{aligned} & \sum_{k=1}^{\infty} b_h(s_h^k; a_h^k) \mathbb{E}[\mathbb{E}_k | H_{k-1}] \\ & \leq \frac{\sum_{k=1}^{\infty} \sum_{h,e}^{1=2} \sum_{k=1}^{\infty} b_h(s_h^k; a_h^k) k^2}{\sum_{h,e}^{1=2} \sum_{k=1}^{\infty} \mathbb{E}[\mathbb{E}_k | H_{k-1}]^2} \end{aligned}$$

$$P \frac{d}{d} \overline{(e-1)kV_h k_1} \text{ "ms : } \quad (\text{by Lemma H.2})$$

Combining these two bounds concludes the proof. □

Lemma H.7 (Optimism) Suppose (E_{model}) holds. Let $\gamma = kV_h k_1 \text{ "ms}$. Then, for any episode $e = 1; 2; \dots; N$, we have

$$\delta_h = 0; 1; \dots; H-1 : E \left[Q_h^?(s_h; a_h) - Q_{h,e}(s_h; a_h) - \gamma \sum_{i=h}^H \delta_i \right]$$

and

$$\delta_h = 0; 1; \dots; H-1 : E \left[V_h^?(s_h) - V_{h,e}(s_h) - \gamma \sum_{i=h}^H \delta_i \right]$$

where

$$Q_h(s_h) := \mathbb{E} \left[\sum_{i=h}^H r_i(s_i; a_i) \mid s_h \right]$$

$$V_h(s_h) = \sum_{h^0=0}^{\gamma^h} V_h^0(s_{h^0})$$

Abusing notation, $\gamma_i(\cdot)$ is the constant function 1.

In particular, we have that

$$E_{d_0} \left[V_0^?(s_0) - V_{0,e}(s_0) - \gamma \sum_{i=0}^H \delta_i \right]$$

Proof. Fix any episode e . We prove both claims via induction on $h = H; H-1; \dots; 1; 0$. The base case holds trivially since $Q_{H,e}$ and $V_H^?$ are zero at every state by definition. Indeed, we have that for any s , including s_H , that

$$E \left[P_{H-1}^?(s_{H-1}; a_{H-1}) (V_H^?(s_H) - V_{H,e}(s_H) - \gamma \sum_{i=H}^H \delta_i) \right]$$

$$= E \left[(0 - 0) \gamma \sum_{i=H}^H \delta_i \right] = 0$$

Now let's show the inductive step. Let $h = H-1; H-2; \dots; 1; 0$ be arbitrary and suppose the inductive hypothesis. So suppose (E_{optimism}) holds at $h+1$ (we don't even need Q -optimism in the future), i.e.

$$E \left[P_h^?(s_h; a_h) (V_{h+1}^?(s_{h+1}) - V_{h+1,e}(s_{h+1}) - \gamma \sum_{i=h+1}^H \delta_i) \right] = E \left[P_h^?(s_h; a_h) (V_{h+1}^?(s_{h+1}) - V_{h+1,e}(s_{h+1}) - \gamma \sum_{i=h+1}^H \delta_i) \right] \quad (\text{IH})$$

Recalling that $Q_{h,e}(s_h; a_h) = r_h(s_h; a_h) + \mathbb{E} \left[P_{h+1}^?(s_{h+1}; a_{h+1}) V_{h+1,e}(s_{h+1}) + b_{h,e}(s_h; a_h) \right]$, we have

$$E \left[P_h^?(s_h; a_h) (Q_{h,e}(s_h; a_h) - Q_h(s_h)) - \gamma \sum_{i=h}^H \delta_i \right]$$

$$= E \left[P_h^?(s_h; a_h) (V_{h+1}^?(s_{h+1}) - V_{h+1,e}(s_{h+1}) - b_{h,e}(s_h; a_h)) - \gamma \sum_{i=h}^H \delta_i \right]$$

$$= E \left[P_h^?(s_h; a_h) (V_{h+1}^?(s_{h+1}) - V_{h+1,e}(s_{h+1}) - b_{h,e}(s_h; a_h)) - \gamma \sum_{i=h}^H \delta_i \right] + (H-h-1) \quad (\text{by (IH)})$$

$$= E \left[P_h^?(s_h; a_h) (V_{h+1}^?(s_{h+1}) - V_{h+1,e}(s_{h+1}) - b_{h,e}(s_h; a_h)) - \gamma \sum_{i=h}^H \delta_i \right] + (H-h-1)$$

$$+ (H-h-1) = (H-h); \quad (\text{by } (E_{\text{model}}) \text{ and } V_{h+1,e} \leq 2V_h \text{ (Lemma H.3)})$$

which proves the Q -optimism claim.

Now let's prove V -optimism.

$$E \left[V_h^?(s_h) - V_{h,e}(s_h) - \gamma \sum_{i=h}^H \delta_i \right]$$

$$\begin{aligned}
 &= \mathbb{E} \sum_{h=0}^{H-1} \left[Q_h^2(s_h; a_h) - Q_{h,e}(s_h; b_h^e(s_h)) \right] \\
 &= \mathbb{E} \sum_{h=0}^{H-1} \left[(Q_h^2(s_h; a_h) - M_{V_{h-1}}(s_h)) + (M_{V_{h-1}}(s_h) - Q_{h,e}(s_h; b_h^e(s_h))) \right] \\
 &+ \mathbb{E} \sum_{h=0}^{H-1} \left[Q_h^2(s_h; a_h) - Q_{h,e}(s_h; b_h^e(s_h)) \right] \\
 &= \mathbb{E} \sum_{h=0}^{H-1} \left[Q_h^2(s_h; a_h) - Q_{h,e}(s_h; b_h^e(s_h)) \right] \\
 &= \mathbb{E} \sum_{h=0}^{H-1} \left[Q_h^2(s_h; a_h) - Q_{h,e}(s_h; b_h^e(s_h)) \right] \\
 &= \mathbb{E} \sum_{h=0}^{H-1} \left[Q_h^2(s_h; a_h) - Q_{h,e}(s_h; b_h^e(s_h)) \right]
 \end{aligned}$$

by Q-optimism. □

Remark H.1 We did not require $b_{h,e}$ to be a valid transition! It is in general unbounded and can even have negative entries!

Lemma H.8 (Simulation) For any episode $e = 1; 2; \dots; N$, we have

$$\mathbb{E}_{d_0} \left[V_{0,e}^i(s_0) - V_0^e(s_0) \right] \leq \sum_{h=0}^{H-1} \mathbb{E}_{b_{h,e}} \left[b_{h,e}(s_h; a_h) + (P_h^2(s_h; a_h) - P_h^e(s_h; a_h)) V_{h+1,e}^i \right]$$

Proof. We progressively unravel the left hand side. For any

$$\begin{aligned}
 &V_{0,e}^i(s_0) - V_0^e(s_0) \\
 &= Q_{0,e}(s_0; \theta_0^e(s_0)) - Q_0^e(s_0; \theta_0^e(s_0)) \\
 &= b_{0,e}(s_0; \theta_0^e(s_0)) + (P_{0,e}^2(s_0; \theta_0^e(s_0)) - P_0^2(s_0; \theta_0^e(s_0))) V_{1,e}^i + P_0^2(s_0; \theta_0^e(s_0)) V_{1,e}^i - V_1^e;
 \end{aligned}$$

where the inequality is due to the thresholding on the value function. Now, perform this recursively on the $P_0^2(s_0; \theta_0^e(s_0)) V_{1,e}^i - V_1^e$ term. Doing this unravelling H times gives the result. □

Theorem H.2 Suppose [Assumption H.1](#) and let β and β_{bonus} be defined as [\(7\)](#). Let $\gamma \in (0, 1)$. Then w.p. at least $1 - \delta$, we have that the regret of LSVI is sublinear,

$$\sum_{e=0}^{N-1} V_{0,e}^i \leq \mathcal{O} \left(d H N M_V^p \frac{1}{\log(HN)} + d^{1.5} H^p \overline{N M_V} M \log(d H N) \right)$$

where \mathcal{O} hides \log dependence.

Proof. We first condition on the high-probability event $\mathcal{E}_{\text{model}}$, which occurs w.p. at least $1 - \delta$. Fix any arbitrary episode e . By optimism [Lemma H.7](#) and the simulation lemma [Lemma H.8](#),

$$\mathbb{E}_{d_0} \left[V_0^2(s_0) - V_0^{b^e}(s_0) \right] \leq \mathbb{E}_{d_0} \left[V_0^e(s_0) - V_0^{b^e}(s_0) \right] + H \sum_{h=0}^{H-1} \mathbb{E}_{b_{h,e}} \left[b_{h,e}(s_h; a_h) + (P_h^2(s_h; a_h) - P_h^e(s_h; a_h)) V_{h+1,e}^i \right] + H$$

Applying [\(E_{model}\)](#) with no indicators, i.e. $\mathbb{1}_{h=h} = 1$ always, gives,

$$\mathbb{E}_{b^e} \sum_{h=0}^{H-1} [2 b_{h,e}(s_h; a_h)] + 2 H :$$

Now, summing over $e = 1; 2; \dots; N$, we have

$$\begin{aligned}
 &\sum_{e=1}^N \mathbb{E}_{d_0} \left[V_0^2(s_0) - V_0^{b^e}(s_0) \right] \\
 &\leq 2 H N + 2 \sum_{h=0}^{H-1} \sum_{e=1}^N \mathbb{E}_{b_{h,e}} [b_{h,e}(s_h; a_h)]
 \end{aligned}$$

By Azuma's inequality applied to the martingale difference $b_{h,e} = E_{b^e} [b_{h,e}(s_h; a_h)] - b_{h,e}(s_h^e; a_h^e)$, which has envelope bounded by H , it implies w.p.1

$$2HN + 2 \sum_{h=0}^{H-1} \sum_{e=1}^N b_{h,e}(s_h^e; a_h^e) + 4 \sqrt{p \over N \log(HN)}:$$

It remains to bound the sum of expected bonuses. By Lemma H.2, we know that almost surely,

$$\sum_{h=0}^{H-1} \sum_{e=1}^N b_{h,e}(s_h^e; a_h^e) \leq H \sqrt{p \over dN \log(N)}:$$

So, putting everything together,

$$\begin{aligned} & \sum_{e=1}^N E_{d_0} V_0^e(s_0) - V_0^{b^e}(s_0) \\ & \leq HN + H \sqrt{p \over dN \log(N)} + \sqrt{p \over N \log(HN)} \\ & \leq HN + \sqrt{p \over dNM_V} \kappa_1 + M_V M \sqrt{p \over \log(dHN)} \leq H \sqrt{p \over dN \log(HN)} \\ & = HN + dHN M_V \sqrt{p \over \log(HN)} \kappa_1 + d^{1.5} H \sqrt{p \over NM_V} M \sqrt{\log(dHN)} \end{aligned}$$

Note that $H \sqrt{p \over dN \log(HN)} = H M_V \sqrt{p \over M \log(HN)}$ is of lower order (with respect to N), we can simply drop it. This concludes the proof. \square

Corollary H.1 By setting $\epsilon = 1/N$, we have that expected regret also has the same rate as above.

Proof. The expected regret by law of total probability, since regret is at most

$$E[\text{Reg}_N] = E[\text{Reg}_N | \mathcal{E}_{\text{model}}] + NH(1 - P(\mathcal{E}_{\text{model}})) \\ = E[\text{Reg}_N | \mathcal{E}_{\text{model}}] + H:$$

Since H is lower-order, we have the same rate. \square

I PROOF OF MAIN THEOREMS

First we prove Theorem 4.1 and Theorem 4.2.

Theorem 4.1 (Regret under generative source access) Suppose Assumptions 3.1-3.4 and 4.1, and suppose the input policies π_k are ϵ_{\min} -exploratory. Then, for any $\epsilon \in (0, 1)$, w.p.1 $1 - \epsilon$, REPTRANSFER when deployed in the target task has regret at most

$$dH^2 \sqrt{p \over T \log(1/\epsilon)}, \text{ with at most } \frac{1}{\epsilon_{\min}} A \frac{3}{\epsilon_{\max}} KT \log \frac{j}{j-1} + K \log j \text{ generative accesses per source task, with}$$

Theorem 4.2 (Regret for REPTRANSFER with REWARDFREE subroutine) Suppose Assumptions 3.1-3.4 and 4.1 hold, and let $\epsilon \in (0, 1)$. Let $f_{k, g_{k=1}^K}$ be exploratory policies learned from running REWARDFREE on each source task with $N_{\text{LSVI-UCB}}$ and $N_{\text{REWARDFREE}}$ set as in Lemma 4.2. Then, w.p.1 $1 - \epsilon$, running REPTRANSFER with $f_{k, g_{k=1}^K}$ has regret in the target task of $\mathcal{O} \left(dH^2 \sqrt{p \over T \log(1/\epsilon)} \right)$, with at most $\mathcal{O} \left(A^4 \frac{3}{\epsilon_{\max}} d^5 H^7 K^2 T^2 (\log(j/j-1) + K \log j) \right)$ generative accesses per source task.

Proof of Theorem 4.1 and Theorem 4.2 For the regret bound, set $M_V = H$ and $M = \frac{1}{\epsilon_{\min}}$ and apply Theorem H.2. This choice of M is valid by the argument following Assumption H.1. This gives us a regret bound of

$$\mathcal{O} \left(dH^2 T \kappa_1 + d^{1.5} H^2 \sqrt{p \over T \log(1/\epsilon)} \right);$$

where κ_1 can be made smaller than $\sqrt{p \over T}$, in which the second term dominates.

Now, we calculate the pre-training phase sample complexity in a source task.

First, let's calculate the reward-free model learning sample complexity, i.e. this is the number of samples required for learning θ_k . Recall that we need this to be sufficiently large such that $\|TV = 1 - N_{LSVI-UCB}$. As required by Lemma 4.2, we need,

$$\begin{aligned} N_{\text{REWARDFREE}} &= \Theta \left(A^3 d^4 H^6 \log(j \text{ } j \text{ } j =) N_{\text{LSVI-UCB}}^2 \right) \\ &= \Theta \left(A^3 d^4 H^6 \log(j \text{ } j \text{ } j =) A^3 d^6 H^8 \right) \\ &= \Theta \left(A^9 d^{16} H^{22} \log(j \text{ } j \text{ } j =) \right) : \end{aligned}$$

Second, we calculate the cross-sampling sample complexity. Recall that the number of samples in each pairwise dataset. In order to reduce to $1 = \bar{T}$, by Lemma 4.1, we need

$$\begin{aligned} 1 = \bar{T} \quad n_{\text{ms}} &= A^3_{\text{max}} K = \min_{1=2} \frac{p}{n} \\ &= A^3_{\text{max}} K = \min_{1=2} \frac{1}{n} \log \frac{j \text{ } j}{j} + K \log j \text{ } j \quad (\text{by (2)}) \end{aligned}$$

which implies that we need

$$n \geq \frac{1}{\min} A^3_{\text{max}} K T \log \frac{j \text{ } j}{j} + K \log j \text{ } j :$$

Incorporating the coverage result from Lemma 4.2 gives,

$$n \geq A^4_{\text{max}} d^5 H^7 K T^2 \log \frac{j \text{ } j}{j} + K \log j \text{ } j :$$

Since each task is in at most 2 pairwise datasets, each of size n , the total pre-training sample complexity per task is at most,

$$\begin{aligned} &N_{\text{REWARDFREE}} + (K - 2) n \\ &= \Theta \left(A^9 d^{16} H^{22} \log(j \text{ } j \text{ } j =) + A^4_{\text{max}} d^5 H^7 K^2 T^2 \log \frac{j \text{ } j}{j} + K \log j \text{ } j \right) : \end{aligned}$$

□

Now we prove Theorem 5.2, restated below.

Theorem 5.2 (Regret with online access). Suppose Assumptions 3.1-3.4, 5.1, 5.2 hold. W.p., Algorithm 5 with appropriate parameters achieves a regret in the target $d^{1:5} H^2 \bar{T} \log(1 =)$, with at most $\text{poly}(A; \text{max}; d; H; K; T; \text{raw}; \text{min}; \log(j \text{ } j \text{ } j =))$ online queries in the source tasks.

Proof of Theorem 5.2 We follows the same format as the proof of Theorem 4.1. The regret bound is identical.

Now let's compute the pre-training sample complexity. The regret bound requires us to set $1 = \bar{T}$. Here, our n_{ms} comes from Lemma G.1, so

$$1 = \bar{T} \quad n_{\text{ms}} = \frac{\text{max} K^{1=2}}{(\text{raw} \text{ } \text{min})^{1=2}} \frac{1}{n} \log \frac{j \text{ } j}{j} + K \log j \text{ } j ;$$

which implies we need

$$n \geq \frac{2_{\text{max}} K}{\text{raw} \text{ } \text{min}} \log \frac{j \text{ } j}{j} + K \log j \text{ } j :$$

Plugging in the coverage of Lemma 4.2,

$$\begin{aligned} n &\geq \frac{2_{\text{max}} K T}{\text{raw}} \log \frac{j \text{ } j}{j} + K \log j \text{ } j \quad A^3 d^5 H^7 2^1 \\ &= \frac{A^3 2_{\text{max}} d^5 H^7 K T}{\text{raw}^2} \log \frac{j \text{ } j}{j} + K \log j \text{ } j : \end{aligned}$$

Here, we only collect one dataset, so the total pre-training sample complexity is

$$N_{\text{REWARDFREE}} + n = \Theta \left(A^9 d^{16} H^{22} \log(j) \log(j) + A^3 \frac{2}{\max} d^5 H^7 K T_{\text{raw}}^{-1} \log \frac{j}{j} + K \log j \right) :$$

□

J EXPERIMENT DETAILS

J.1 CONSTRUCTION OF COMBLOCK

In this section we first introduce the vanilla Combination lock (Comblock) environment that is widely used as the benchmark for algorithms for Block MDPs. We provide a visualization of the comblock environment in Fig. 1(a). Concretely, the environment has a horizon of 3 latent states $z_{i,h}; i \in \{0, 1, 2\}$ for each timestep h and 10 actions. Among the three latent states, we denote $z_{0,h}$ as the good states which leads to the final reward and the bad states. At the beginning of the task, the environment will uniformly and independent sample 1 out of the 10 actions for each good state $z_{0,h}$ and $z_{1,h}$ for each timestep h , and we denote these actions $a_{i,h}$ as the optimal actions (corresponding to each latent state). These optimal actions, along with the task itself, determines the dynamics of the environment. At each good latent state $z_{i,h}$, if the agent takes the correct action, the environment transits to the either good state at the next timestep (i.e., $z_{i,h+1}$) with equal probability. Otherwise, if the agent takes any 9 of the bad actions, the environment will transition to the bad state $z_{2,h+1}$ deterministically, and the bad states transit to only bad states at the next timestep deterministically. There are two situations where the agent receives a reward: one is upon arriving the good states at the last timestep, the agent receives a reward of 1. The other is upon the first ever transition into the bad state, the agent receives an “anti-shaped” reward of 0.1 with probability 0.5. Such design makes greedy algorithms without strategic exploration such as policy optimization methods easily fail. For the initial state distribution, the environment starts for $z_{i,0}$ with equal probability. The dimension of the observation is $d_{\text{obs}} = (H + j) \times (S_j + 1) \times e$. For the emission distribution, given a latent state $z_{i,h}$, the observation is generated by first concatenate the one-hot vectors of the state and the horizon, adding i.i.d. $(0,0.1)$ noise at each entry, appending 0 at the end if necessary. Then finally we apply a linear transformation on the observation with a Hadamard matrix. Note that without a good feature or strategic exploration, it takes actions to reach the final goal with random actions.

J.2 CONSTRUCTION OF TRANSFER SETUP IN SECTION. 6

In this section we introduce the detailed construction of the experiment in Section. 6. For the source environment, we simply generate 5 random vanilla comblock environment described in Section. J.1. Note that in this way we ensure that the emission distribution shares across the sources, but the latent dynamics are different because the optimal actions are independently randomly selected. For the target environment, for each timestep h , we randomly acquire the optimal actions from one of the sources and set it to be the optimal action of the target environment at time h if the selected optimal actions are different for the two good states. Otherwise we keep sampling until they are different. Note that under such construction, since we fix the emission distribution, Assumption 3.4 is satisfied if we set $\beta = 1$ for the source environment where we select the optimal action and 0 for the other sources, at each timestep. To see how Assumption 5.1 is satisfied, recall that comblock environment naturally satisfies Assumption 3.3, and identical emission implies that the conditional ratio of all observations between source and target is 1.

J.3 CONSTRUCTION OF TRANSFER SETUP IN SECTION. 6

Now we introduce the construction of the Comblock with Partitioned Observation (Comblock-PO) environment. Comparing with the vanilla comblock environment, the major difference is in

the observation space. In this setting, the size of the observation depends on the number of source environments K . Let the size of the original observation space be $|O|$, the size of the observation for comblock-PO is $K|O|$. For the k -th source environment, where $k \in [K]$, the environment first generates the $|O|$ -dimensional observation vector as in the original comblock, and then embed it to the $(k-1)|O|$ -th to $k|O|$ -th entries of the $K|O|$ -dimensional observation vector, where it is 0 everywhere else. Thus we can see that the observation space for each source environment is disjoint (and thus the name partitioned observations). For the target environment, since the latent dynamics are the same, we only need to design the emission distribution: for each latent state, we assign the emission distribution uniformly at random from one of the sources.

J.4 IMPLEMENTATION DETAILS

Our implementation builds on BRIE (Zhang et al., 2022). In the Multi-task REPLEARN stage, we require our learned feature to predict the Bellman backup of all the sources simultaneously. Therefore, in each iteration we have k discriminators and k sets of linear weights (instead of 1 in BRIE), where k is the number of source environments. For the deployment stage we implement LSVI following Algorithm 6.

To create the training dataset for Multi-task REPLEARN, for each (i, j) environment pairs where $i \in [K]$, $j \in [K]$, we collect 500 samples for each timestep. For each (i, i) environment pairs, we collect $500 \cdot (k-1) \cdot k$ samples for each timestep, where k denotes the number of sources. Thus we ensure that the number of samples from cross transition of different environments is the same as the number of samples from cross transition of the same environment. For the online setting, we simply sample $1000 \cdot (k-1) \cdot k$ samples for each (i, i) cross transition to ensure that the total number of samples is the same for G-TRANSFER and O-REPTTRANSFER.

To sample the initial state action pair $(i, s; a)$ pair as in (1), for 90% of the samples, we follow the policy from each source environment trained using BRIE. For the remaining 10%, we follow the same policy to state s , and then take a uniform random policy to get a . With this sampling scheme we ensure that Assumption 3.3 is satisfied. In the setting of Section 6, we follow a more simple procedure to ensure that the samples are more balanced among the three states: we skip the first sampling step from environment (i.e., samples given $(s; a)$), and simply reset environment to s , where s is one of the three states with equal probability, and generate the observation accordingly. Note that such visitation distribution is also possible in the online setting with a more nuanced sampling procedure, and in the experiment we use the same sampling procedure for G-TRANSFER and O-REPTTRANSFER for a fair comparison.

J.5 HYPERPARAMETERS

In this section, we record the hyperparameters we try and the final hyperparameter we use for each baseline. The hyperparameters for REPTTRANSFER in Section 6 is in Table 3. The hyperparameters for REPTTRANSFER in Section 6 is in Table 4. The hyperparameters for BRIE is in Table 5. We use the same set of hyperparameters for G-TRANSFER and O-REPTTRANSFER.

Table 3: Hyperparameters for **ERT**TRANSFER in Comblock.

	Value Considered	Final Value
Decoder learning rate	{1e-2}	1e-2
Discriminatorf learning rate	{1e-2}	1e-2
Discriminatorf hidden layer size	{256}	256
RepLearn Iteratiof	{30}	30
Decoder number of gradient steps	{64}	64
Discriminatorf number of gradient steps	{64}	64
Decoder batch size	{256}	256
Discriminatorf batch size	{512}	512
RepLearn regularization coef cient	{0.01}	0.01
Decoder softmax temperature	{1}	1
Decoder ₀ softmax temperature	{0.1}	0.1
LSVI bonus coef cient	{1, $\frac{H}{5}$ }	1
LSVI regularization coef cient	{1}	1
Buffer size	{1e5}	1e5
Update frequency	{50}	50
Optimizer	{SGD}	SGD

Table 4: Hyperparameters for **ERT**TRANSFER in Comblock-PO.

	Value Considered	Final Value
Decoder learning rate	{1e-2}	1e-2
Discriminatorf learning rate	{1e-2}	1e-2
Discriminatorf hidden layer size	{256,512}	256
Discriminatorf hidden layer number	{2,3}	3
RepLearn Iteratiof	{30,40,50,100,150}	50
Decoder number of gradient steps	{64,80,128,256}	64
Discriminatorf number of gradient steps	{64,80,128,256}	64
Decoder batch size	{256,512}	512
Discriminatorf batch size	{256,512}	512
RepLearn regularization coef cient	{0.01}	0.01
Decoder softmax temperature	{1}	1
Decoder ₀ softmax temperature	{0.1,1}	1
LSVI bonus coef cient	{1, $\frac{H}{5}$ }	1
LSVI regularization coef cient	{1}	1
Buffer size	{1e5}	1e5
Update frequency	{50}	50
Optimizer	{SGD, Adam}	Adam

J.6 VISUALIZATIONS

In this section we provide a comprehensive visualization of the decoders for all baselines in the target environment. We observe that the behaviors of all baselines are similar across the 5 random seeds. Thus to avoid redundancy, we only show the visualization from 1 random seed. We provide an example in Fig. 3 on how to interpret the visualization: let the emission function of the target environment be α , and let the decoder that we are evaluating be β and to generate the blue block in Fig. 3, we sample 30 observations $s_n, g_{n=1}^{30}$ from the target environment at $t_{1:13}$, the latent state 1 (the title of the subplot) from timestep t_3 (the x-axis). Concretely, $s_n, g_{n=1}^{30} \sim \alpha(\cdot | z_{1:13})$. The blue block denotes the three-dimensional decoded latent states from these 30 observations: $\hat{z} = \frac{1}{30} \sum_{n=1}^{30} (s_n)$.

Table 5: Hyperparameters for REE in Comblock and Comblock-PO.

	Value Considered	Final Value
Decoder learning rate	{1e-2}	1e-2
Discriminatorf learning rate	{1e-2}	1e-2
Discriminatorf hidden layer size	{256}	256
RepLearn IteratioT	{30}	30
Decoder number of gradient steps	{64}	64
Discriminatorf number of gradient steps	{64}	64
Decoder batch size	{512}	512
Discriminatorf batch size	{512}	512
RepLearn regularization coef cient	{0.01}	0.01
Decoder softmax temperature	{1}	1
Decoder ₀ softmax temperature	{0.1}	0.1
LSVI bonus coef cient	{ $\frac{H}{5}$ }	$\frac{H}{5}$
LSVI regularization coef cient	{1}	1
Buffer size	{1e5}	1e5
Update frequency	{50}	50

J.6.1 VISUALIZATIONS FROM SECTION. 6

We record the visualization of the 5 sources from Fig. 3 to Fig. 7, REPTTRANSFER in Fig.8; G-REPTTRANSFER in Fig. 9; running REE on target in Fig. 10.

Figure 3: Visualization of decoders from source 1. Note the collapse happens at timestep 5, 9 and 17.

Figure 4: Visualization of decoders from source 2. Note the collapse happens at timestep 1 and 10.

J.6.2 VISUALIZATIONS FROM SECTION. 6

We record the visualization of the 2 sources from Fig. 11 and Fig. 12. REPTTRANSFER in Fig. 13; G-REPTTRANSFER in Fig. 14; running BREE on target in Fig. 15. Note that the features collapse at some timesteps in Fig. 14 and Fig. 15, but this is acceptable because the optimal actions at those timesteps are the same for the collapsed states.

Figure 5: Visualization of decoders from source 3. Note the collapse happens at timestep 14 and 15.

Figure 6: Visualization of decoders from source 4. Note the collapse happens at timestep 7, 16, 24.

Figure 7: Visualization of decoders from source 3. Note the collapse happens at timestep 13 and 16.

Figure 8: Visualization of decoders from ~~CERT~~TRANSFER

Figure 9: Visualization of decoders from GERT RANSFER

Figure 10: Visualization of decoders from running BE on target.

Figure 11: Visualization of decoders from source environment 1.

Figure 12: Visualization of decoders from source environment 2.

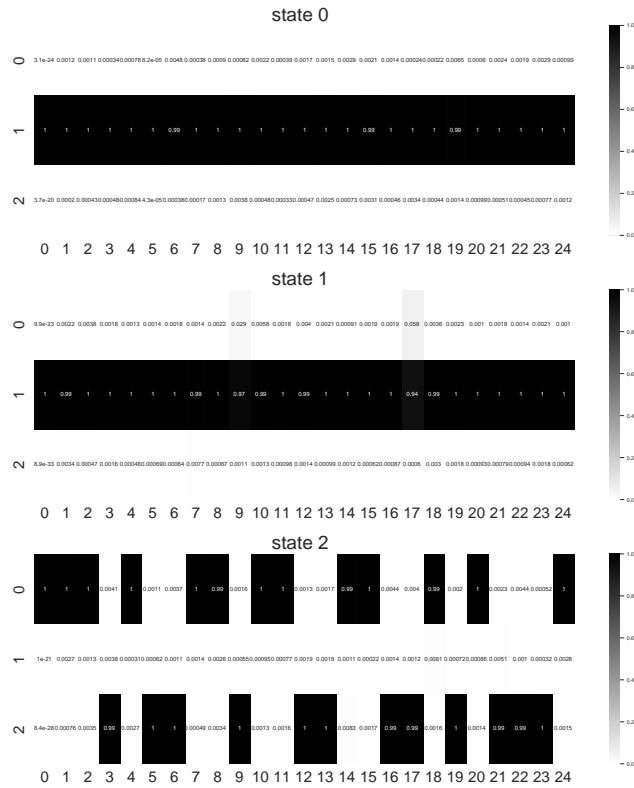


Figure 13: Visualization of decoders from O-REPTRANSFER.

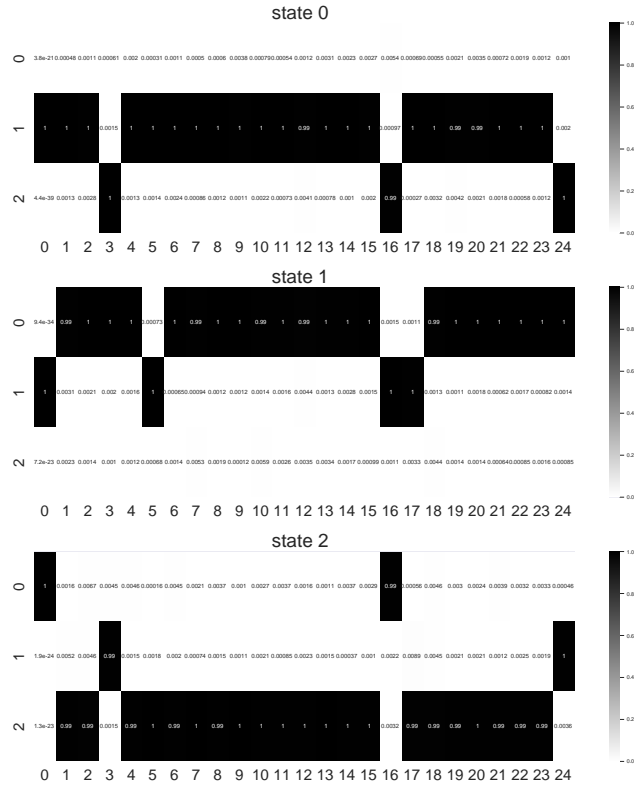


Figure 14: Visualization of decoders from G-REPTRANSFER.

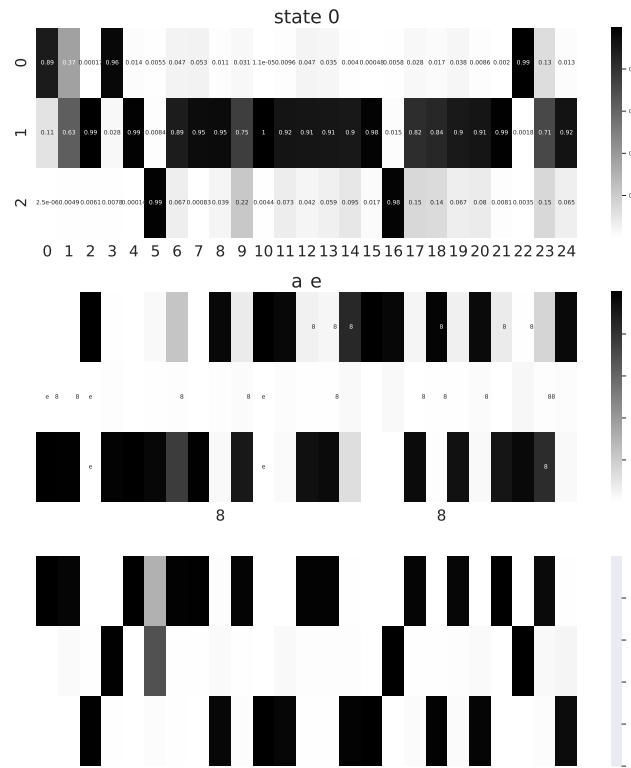


Figure 15: Visualization of decoders running BRIEE in the target.