
Making Linear MDPs Practical via Contrastive Representation Learning

Tianjun Zhang^{*1} Tongzheng Ren^{*23} Mengjiao Yang¹³ Joseph E. Gonzalez¹ Dale Schuurmans³⁴ Bo Dai³

Abstract

It is common to address the curse of dimensionality in Markov decision processes (MDPs) by exploiting low-rank representations. This motivates much of the recent theoretical study on linear MDPs. However, most approaches require a given representation under unrealistic assumptions about the normalization of the decomposition or introduce unresolved computational challenges in practice. Instead, we consider an alternative definition of linear MDPs that automatically ensures normalization while allowing efficient representation learning via contrastive estimation. The framework also admits confidence-adjusted index algorithms, enabling an efficient and principled approach to incorporating optimism or pessimism in the face of uncertainty. To the best of our knowledge, this provides the first practical representation learning method for linear MDPs that achieves both strong theoretical guarantees and empirical performance. *Theoretically*, we prove that the proposed algorithm is sample efficient in both the online and offline settings. *Empirically*, we demonstrate superior performance over existing state-of-the-art model-based and model-free algorithms on several benchmarks.

1. Introduction

In reinforcement learning (RL), the goal of an agent is to learn a policy that maximizes the expected reward accumulated during sequential interaction with an unknown environment. Most RL algorithms are developed under the framework of Markov decision processes (MDPs) (Puterman, 2014), where performance can be evaluated by “regret”, *i.e.*, the difference in accumulated reward between the best policy and the policies produced by the algorithm.

^{*}Equal contribution ¹UC Berkeley ²UT Austin ³Google Brain ⁴University of Alberta. Correspondence to: Tianjun Zhang <tianjunz@berkeley.edu>, Tongzheng Ren <tongzheng@utexas.edu>, Bo Dai <bodai@google.com>.

It is known that in the tabular case where the number of states $|\mathcal{S}|$ and actions $|\mathcal{A}|$ are finite, there exist algorithms that can achieve the asymptotic regret lower bound $\Omega\left(\sqrt{|\mathcal{S}||\mathcal{A}|} \cdot \sqrt{T}\right)$ (Jaksch et al., 2010; Dann & Brunskill, 2015; Osband & Van Roy, 2016; Azar et al., 2017; Jin et al., 2018) after T iterations, which is unimprovable in the worst case. Unfortunately, the polynomial dependence on $|\mathcal{S}|$ and $|\mathcal{A}|$ impedes the application of the regret bounded RL algorithms in practice, where the number of states and actions can be very large or even infinite. This difficulty is known as the “curse of dimensionality” in control and RL (Bellman, 1966).

To tackle the curse of dimensionality, low-rank representations via function approximation have been used to exploit additional structures in an MDP. One prominent direction is the linear MDP framework (Yang & Wang, 2020; Jin et al., 2020), which leverages additional assumptions to reduce the algorithmic and theoretical challenge of coping with the dimensionality issue in RL. In particular, assuming a *known* low-rank feature that can represent the transition operator and reward function in a linear MDP, then simple optimistic algorithms can achieve regret that is *independent* of the cardinality of the state and action sets. Therefore, the linear MDP framework partially breaks the curse of dimensionality, provides a deeper understanding of the empirical success of deep RL (Mnih et al., 2015; Levine et al., 2016), and can deliver performance improvements. However, the requirement that the low-rank features be known beforehand is difficult to achieve in practice, which has limited the practical impact of linear MDPs in real applications.

Recently, there have been important new progress in making linear MDPs more practical by introducing representation learning techniques that satisfy provable guarantees. Agarwal et al. (2020) proposed the first theoretically-grounded representation learning algorithm for linear MDPs, which achieves polynomial sample complexity given a maximum likelihood estimation (MLE) oracle for stochastic transition dynamics. Later, Uehara et al. (2021) improved their algorithm by exploiting upper confidence bounds (UCB) to reduce sample complexity under the same MLE oracle. These efforts have revealed the possibility of achieving statistically efficient representation learning in linear MDPs. However, these algorithms have not addressed the underlying

ing computational issues, leaving a gap between the theory and practice of linear MDPs. Specifically, the transition dynamics in the definition of a linear MDP do not necessarily form a density model, meaning that there might not be a computationally efficient way to implement an MLE oracle. Moreover, the classical formulation of a linear MDP requires an extra normalization condition on the representation, which is also computationally challenging.

Our contributions. In this paper, we address the lingering computational efficiency issues in representation learning for linear MDPs. We provide solutions to the aforementioned difficulties, bridge the gap between theory and practice, and eventually establish a sound yet practical foundation for learning in linear MDPs. More specifically:

- We clarify the importance of the normalization condition and propose a variant of linear MDPs wherein it is easy to enforce normalization (Section 4.1);
- We develop the *ConTrastive Representation Learning algorithm*, CTRL that can implicitly satisfy the density requirement (Section 4.2);
- We incorporate the representation learning techniques in optimistic and pessimistic confidence-adjusted RL algorithms, CTRL-UCB and CTRL-LCB, for both online and offline RL problems (Section 4.3), and establish their sample complexities respectively (Section 5);
- We conduct a comprehensive comparison to existing state-of-the-art model-based and model-free RL algorithms on several benchmarks for the online and offline settings, demonstrating superior empirical performance of the proposed CTRL (Section 6).

1.1. Related Work

There have been several attempts to achieve provably sound and effective representation learning within more restricted models that make extra assumptions. For example, Du et al. (2019); Misra et al. (2020) considered representation learning in *block MDPs*; Modi et al. (2021) proposed a model-free representation learning algorithm for *low-nonnegative rank MDPs*. These are special case of linear MDPs and are quite limited in terms of expressiveness, as shown in (Agarwal et al., 2020). We emphasize that these algorithms also primarily focus on sample complexity, and do not necessarily address all the computational complexity challenges.

One exception is Ren et al. (2021), which investigates an alternative proposal for exploiting linear MDP structure to achieve provable algorithms in practical problems based on a different solution to the computational intractability issue. Their proposal is to maintain a *posterior* of the control model under specific *state-independent noise*, e.g., Gaussian noise, exploiting the explicit Fourier transform of the Gaussian distribution to recover an equivalent *infinite-*

dimensional linear MDP posterior. By contrast, the method proposed here only requires a *point estimate* of the features. Moreover, we do not restrict the dynamics and noise model, attaining greater *generality*, while learning a *low-rank* feature representation that is more compact and effective than a random features approximation.

Beyond theoretically-justified representation learning techniques for RL, many representation learning methods have been proposed for different purposes, such as: reconstruction (Watter et al., 2015; Li et al., 2019; Hafner et al., 2019), bi-simulation (Ferns et al., 2004; Gelada et al., 2019; Castro, 2020; Zhang et al., 2020b), invariance preservation (Zhang et al., 2020a), successor predictive occurrence backup (Dayan, 1993; Barreto et al., 2016; Kulkarni et al., 2016), and spectral decomposition of various state-state transition operators (Mahadevan & Maggioni, 2007; Wu et al., 2018; Duan et al., 2018). Contrastive losses have also been introduced for representation learning in RL. For instance, (Yang et al., 2020) extends Contrastive Predictive Coding (Oord et al., 2018) for planning with specially designed positive pairs for *state* representation learning; (Nachum & Yang, 2021; Yang et al., 2021) considered a contrastive loss for training an energy-based model with offline data to produce an *infinite-dimensional* representation for imitation learning in *finite-action MDPs*. Note that, beyond the various differences between these approaches to representation learning (state vs. state-action modeling, imitation learning vs. policy optimization, finite vs. continuous actions), a major differentiator of this paper is that CTRL is naturally developed with UCB or LCB to support exploration or conservative exploitation respectively, while previous work considered passive data collection that was *independent* of such optimistic or pessimistic considerations. This property of our proposed method allows provable sample complexity bounds for both online and offline settings.

2. Preliminaries

We first provide a brief introduction to linear MDPs and noise contrastive estimation, which will play an important role in the algorithm design.

2.1. Markov Decision Processes (MDPs)

Markov Decision Processes (MDPs) are a standard model of sequential decision making considered in reinforcement learning. An MDP is specified by a tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, r, P, \rho, \gamma)$, where \mathcal{S} is the state space, \mathcal{A} is the action space, $r : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ is the reward function, $P : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ is the transition operator, ρ is an initial state distribution and $\gamma \in (0, 1]$ is the discount factor. A policy π can be defined as $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$ where $\Delta(\mathcal{A})$ denotes the distributions over domain \mathcal{A} . Following the standard notation, we define the value function $V_{P,r}^{\pi}(s) :=$

$\mathbb{E}_{P,\pi} [\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) | s_0 = s]$ and the action-value function $Q_{P,r}^{\pi}(s, a) = \mathbb{E}_{P,\pi} [\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) | s_0 = s, a_0 = a]$, which are the expected discounted cumulative rewards under transition P and reward r when executing policy π . From these two definitions, it is straightforward to establish the following recursive relationship:

$$Q_{P,r}^{\pi}(s, a) = r(s, a) + \gamma \mathbb{E}_{s' \sim P(\cdot | s, a)} [V_{P,r}^{\pi}(s')], \quad (1)$$

$$V_{P,r}^{\pi}(s) = \mathbb{E}_{a \sim \pi(s)} [Q_{P,r}^{\pi}(s, a)]. \quad (2)$$

For simplicity, with slightly abuse of notation, we define $V_{P,r}^{\pi} = \mathbb{E}_{s \sim \rho} [V_{P,r}^{\pi}(s)]$. Furthermore, we additionally define the occupancy measure:

$$d_P^{\pi}(s, a) = (1 - \gamma)^{-1} \mathbb{E}_{s_0 \sim \rho, a_0 \sim \pi(s_0)} [\sum_{t=0}^{\infty} \gamma^t \mathbf{1}(s, a)], \quad (3)$$

where $\mathbf{1}(\cdot)$ is the indicator function. Clearly, d_P^{π} is a probability measure on (s, a) that satisfies

$$\mathbb{E}_{(s,a) \sim \rho \times \pi} [Q_{P,r}^{\pi}(s, a)] = \mathbb{E}_{(s,a) \sim d_P^{\pi}(s,a)} [r(s, a)], \quad (4)$$

as shown in Nachum & Dai (2020).

When $|\mathcal{S}|$ and $|\mathcal{A}|$ are large or even infinite, it is necessary to incorporate function approximation to exploit MDP structure, to address the curse of dimensionality. One of the most prominent forms of structured MDP that allows for simple function approximation is linear MDPs, which add the following structural assumption:

Assumption 2.1 (Linear MDP (Jin et al., 2020)). An MDP $\mathcal{M} = (\mathcal{S}, \mathcal{A}, H, P, r)$ is called a *linear MDP* if there exists two feature maps $\phi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^d$, $\mu : \mathcal{S} \rightarrow \mathbb{R}^d$ and a vector $\theta_r \in \mathbb{R}^d$, such that

$$P(s' | s, a) = \langle \phi(s, a), \mu(s') \rangle, \quad r(s, a) = \langle \phi(s, a), \theta_r \rangle \quad (5)$$

and ϕ and μ also satisfy the normalization conditions:

$$\forall (s, a), \|\phi(s, a)\|_2 \leq 1, \|\theta_r\|_2 \leq \sqrt{d} \quad (6)$$

$$\forall g : \mathcal{S} \rightarrow \mathbb{R}, \|g\|_{L_{\infty}} \leq 1 \Rightarrow \left\| \int_{\mathcal{S}} \mu(s') g(s') ds' \right\|_2 \leq \sqrt{d} \quad (7)$$

An essential property of a linear MDP is that $Q_{P,r}^{\pi}(s, a)$ for any policy π is linear w.r.t. $\phi(s, a)$ (Jin et al., 2020); that is, when (5) is applied to the right hand side of (1) we obtain:

$$\begin{aligned} Q_{P,r}^{\pi}(s, a) &= r(s, a) + \gamma \int V_{P,r}^{\pi}(s') P(s' | s, a) ds' \\ &= \left\langle \phi(s, a), \theta_r + \gamma \int V_{P,r}^{\pi}(s') \mu(s') ds' \right\rangle. \end{aligned}$$

Consequently, one can naturally incorporate linear function approximation when learning a Q function in a linear MDP, simplifying practical application.

We emphasize that, although the MDP is referred to as ‘‘linear’’ here, the model components $\phi(s, a)$ and $\mu(s')$ can be flexibly parametrized by deep models, and thus can be viewed as energy-based models (LeCun et al., 2006) with special potential functions, *i.e.*, $\log(\langle \phi(s, a), \mu(s') \rangle)$.

2.2. Noise Contrastive Estimation

Noise contrastive estimation (NCE) (Gutmann & Hyvarinen, 2010; 2012) was first proposed to estimate unnormalized statistical models without requiring explicit calculation of the partition function. For an (unnormalized) statistical model $p_f(x) = \frac{f(x)}{Z_f}$ (where $f(x) > 0$ for all x in the support and $Z_f = \int f(x) dx$), maximum likelihood estimation (MLE) is generally intractable due to the presence of Z_f . NCE was developed to estimate f from data $\{(x_i)\}_{i \in [n]}$ by the following M-estimator:

$$\begin{aligned} (\hat{f}_n^{\text{NCE}}, \hat{\gamma}_n^{\text{NCE}}) &= \arg \max_{f \in \mathcal{F}, \gamma \in \mathbb{R}} \frac{1}{n} \sum_{i=1}^n \left[\log(h(x_i; f, \gamma)) \right. \\ &\quad \left. + \sum_{j=1}^K \log(1 - h(y_{i,j}; f, \gamma)) \right], \quad (8) \end{aligned}$$

where $\{y_{i,j}\}_{i \in [n], j \in [K]}$ are artificially generated data sampled from a predefined noise distribution $q(y)$ that has full support on the domain of interest, and $h(x; f, \gamma)$ is defined by

$$h(x; f, \gamma) = \frac{f(x) \exp(-\gamma)/q(x)}{f(x) \exp(-\gamma)/q(x) + K}. \quad (9)$$

Here $\hat{\gamma}_n^{\text{NCE}}$ is an estimator of the normalizing constant Z_f . The objective function (8) is also known as the binary-based objective. It can be implemented via a logistic regression that distinguishes real data from generated noise.

The binary objective works well for an unconditional model. However, as Ma & Collins (2018) point out, when (8) is used to estimate a *conditional* model $p(x|u) = \frac{f(x,u)}{Z_f(u)}$ (where $Z_f(u) = \int f(x, u) dx$), NCE cannot provide a consistent estimator, due to the different values of $Z_f(u)$ for different choices of the conditioning variable value u . To resolve this issue, a ranking-based objective has been proposed to provide consistent estimation under conditioning:

$$\hat{f}_n^{\text{NCE}} = \arg \max_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n \log \frac{f(x_i, u_i)/q(x_i)}{\frac{f(x_i, u_i)}{q(x_i)} + \sum_{j=1}^K \frac{f(y_{i,j}, u_i)}{q(y_{i,j})}}, \quad (10)$$

where $\{(x_i, u_i)\}_{i=1}^n$ are the given data. This estimator (10) can be viewed as a multi-class extension to distinguishing the real data from generated noise.

Using standard asymptotic M-estimator theory, Gutmann & Hyvarinen (2010; 2012); Riou-Durand & Chopin (2018); Ma & Collins (2018) have provided consistency and asymptotic normality results for the parametric model under the binary-based objective, while Ma & Collins (2018) provided similar results for the ranking-based objective.

3. The Difficulties of Practical Linear MDP

Here we explicitly identify the computational difficulties in applying linear MDP, paving the way to making it practical.

3.1. Normalization Conditions

The normalization condition in Assumption 2.1 might seem strange at first glance. One could ask why there is a constraint $\|\int_{\mathcal{S}} \mu(s')g(s')ds'\|_2 \leq \sqrt{d}$ (for all $\|g\|_{L_\infty} \leq 1$) instead of any other function of d ? To motivate this choice, consider the tabular MDP, where $|\mathcal{S}| < \infty$, $|\mathcal{A}| < \infty$, which can be viewed as a special linear MDP with $d = |\mathcal{S}||\mathcal{A}|$, $\phi(s, a) = \mathbf{e}_{(s,a)}$ for \mathbf{e} the canonical basis of $\mathbb{R}^{|\mathcal{S}||\mathcal{A}|}$, and $\mu(s')_{(s,a)} = P(s'|s, a)$. It is immediate in this case that $\|\phi(s, a)\|_2 = 1$ and $\int_{\mathcal{S}} \mu(s')ds' = \mathbf{1} \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|}$ (where $\mathbf{1}$ denotes the vector with all entries equal to 1), which automatically satisfies the normalization condition. This constraint plays a critical role in obtaining a polynomial dependence on d in a worst case regret bound in linear MDPs, as it guarantees that (Jin et al., 2020, Lemma B.1)

$$\left\| \int_{\mathcal{S}} V_{P,r}^\pi(s')\mu(s')ds' \right\|_2 \leq (1 - \gamma)^{-1}\sqrt{d}$$

noting that $\|V_{P,r}^\pi\|_{L_\infty} \leq (1 - \gamma)^{-1}$. Such a guarantee is crucial both for the theoretical analysis and practical implementation, as it ensures the linear coefficients needed to express the Q function are not too large, which guarantees low statistical complexity in estimation.

However, when we want to learn ϕ and μ using a function approximator like a neural network, there is no practical mechanism for guaranteeing or even checking the normalization condition. Note that condition (7) requires exact knowledge of the Lebesgue integral $\int_{\mathcal{S}} \mu(s')g(s')ds'$ for an *indefinite* $g(s')$. Beyond a tabular MDP, even if g were known ahead of time, checking the constraint would be difficult: in a discrete space \mathcal{S} , the Lebesgue integral involves a summation over the entire set, whereas numerical methods would be needed in a continuous \mathcal{S} , which can be computationally inefficient or even fail in high dimensional spaces. More significantly, guaranteeing (7) requires the solution to an optimization on the functional g . Given that evaluating the Lebesgue integral for a given g would already be challenging, there is little hope for solving the optimization of g with a computationally efficient method.

Therefore, this normalization condition can be quite restrictive and impedes practical application of linear MDPs. Even in cases where the linear factorization condition (5) can be satisfied, ensuring the normalization condition (7) remains difficult. For example, Ren et al. (2021) show that for nonlinear control problems with Gaussian perturbation, special properties of the Gaussian kernel can be leveraged to obtain an infinite dimensional linear representation that satisfies (5); however, without further assumptions this does not provide adequate bounds for the term in (7). Such an observation is disappointing given that (5) provides a good criterion for learning the representation.

3.2. MLE Computation Oracle

In practice, one only has access to raw observations of state and action, which by themselves do not provide an adequate representation for a linear MDP. To resolve this issue, the seminal work (Agarwal et al., 2020) proposed to learn the representations ϕ and μ using MLE, but at a cost of a computationally inefficient policy-cover method to ensure sufficient exploration. Later, Uehara et al. (2021) replaced policy-cover exploration with confidence-adjusted index form exploration, which is far more computationally efficient. However, this does not completely solve the problem, since MLE for conditional model estimation itself can be computationally intractable, due to the constraints on the marginal regularity, *i.e.*,

$$\left\langle \phi(s, a), \int_{\mathcal{S}} \mu(s')ds' \right\rangle = 1, \quad \forall (s, a) \in \mathcal{S} \times \mathcal{A}.$$

It is not clear how to solve the optimization problem under such a hard constraint. One idea might be to normalize ϕ via $\frac{\phi(s,a)}{\int_{\mathcal{S}} \phi(s,a)\mu(s')ds'}$, which would naturally satisfy the marginal regularity condition and allow MLE without additional constraints. However, this creates new problems at decision making time. In particular, given a fixed (s, a) , one needs to compute the denominator $\int_{\mathcal{S}} \phi(s, a)\mu(s')ds'$ to extract the corresponding representation, which can be intractable when $|\mathcal{S}|$ is large or infinite. Moreover, since $\int_{\mathcal{S}} \phi(s, a)\mu(s')ds'$ depends on (s, a) , the integral would need to be estimated for several $(s, a) \in \mathcal{S} \times \mathcal{A}$ to choose an action, making the computational issue even more severe. Such an approach also causes trouble in the representation learning stage when gradient based optimization is used to solve the MLE, since the gradient of the integral would need to be estimated for each gradient step. Hence, the MLE computation oracle assumed in (Agarwal et al., 2020; Uehara et al., 2021) is not generally practical from a computational perspective.

4. The Proposed Solutions

In this section, we propose a new approach that addresses the two difficulties revealed for learning linear MDPs.

4.1. Modified Linear MDP

To ensure the normalization condition in Assumption 2.1, we introduce an auxiliary probability measure $p(s')$ in conjunction with $\mu(s')$ to obtain a modified linear MDP:

$$P(s'|s, a) = \langle \phi(s, a), p(s')\mu(s') \rangle. \quad (11)$$

Here $p(s')$ can be any probability measure with full support \mathcal{S} , which plays a similar role to the base measure in exponential families and generalized energy based models (Arbel et al., 2020). That is, $p(s')$ will remain fixed but $\mu(s')$ will be learned from data. Note that the introduction of $p(s')$ does not compromise expressiveness in classically consid-

Algorithm 1 CTRL-UCB: Online Exploration with Representation Learning

- 1: **Input:** Regularizer λ_n , parameter α_n , Model class $\mathcal{M} = \{(\mu, \phi) : \mu \in \Psi, \phi \in \Phi\}$, Iteration N , Number of Negative Samples K .
- 2: Initialize $\pi_0(\cdot | s)$ to be uniform; set $\mathcal{D}_0 = \emptyset$
- 3: **for** episode $n = 1, \dots, N$ **do**
- 4: Collect a transition (s, a, s') where $s \sim d_{P^*}^{\pi_{n-1}}$, $a \sim U(\mathcal{A})$, $s' \sim P^*(\cdot | s, a)$.
- 5: $\mathcal{D}_n = \mathcal{D}_{n-1} \cup \{s, a, s'\}$
- 6: Learn representation $\hat{\phi}_n(s, a)$ via (10).
- 7: Update the empirical covariance matrix

$$\hat{\Sigma}_n = \sum_{s, a \in \mathcal{D}_n} \hat{\phi}_n(s, a) \hat{\phi}_n(s, a)^\top + \lambda_n I$$

- 8: Set the exploration bonus:

$$\hat{b}_n(s, a) = \min \left(\alpha_n \sqrt{\hat{\phi}_n(s, a)^\top \hat{\Sigma}_n^{-1} \hat{\phi}_n(s, a)}, 2 \right)$$

- 9: Update policy $\pi_n = \operatorname{argmax}_\pi V_{\hat{P}_n, r + \hat{b}_n}^\pi$
 - 10: **end for**
 - 11: **Return** π_1, \dots, π_N
-

ered settings. Take, for example, a tabular MDP as in Section 3.1: here we can still use $\phi(s, a) = \mathbf{e}_{s,a}$ and let $p(s')$ be a uniform distribution on \mathcal{S} and $\mu(s')_{(s,a)} = |\mathcal{S}|P(s' | s, a)$, which automatically ensures (11) for all (s, a, s') .

The key benefit of modifying the linear MDP in this manner is that we can now easily control the normalization condition. Specifically, for $\|g\|_{L_\infty} \leq 1$, by Jensen’s inequality, we have

$$\begin{aligned} \left\| \int_{\mathcal{S}} p(s') \mu(s') g(s') ds' \right\|_2^2 &\leq \int_{\mathcal{S}} p(s') \|\mu(s') g(s')\|_2^2 ds' \\ &\leq \int_{\mathcal{S}} p(s') \|\mu(s')\|_2^2 ds', \end{aligned}$$

which can now be easily approximated via Monte Carlo estimation instead of numerical integration. This upper bound provides a straightforward way to control normalization by introducing a regularizer $\mathbb{E}_{p(s')} [\|\mu(s')\|_2^2]$ into the NCE losses (8) and (10). Alternatively, one could also ensure boundedness simply by introducing a bounded activation, e.g., $\frac{1}{1+\exp(\cdot)}$, $\exp(-(\cdot)^2)$, or $\tanh(\cdot)$, at the final layer in a neural network parametrization of $\mu(\cdot)$.

4.2. NCE instead of MLE

Due to the computational intractability of MLE discussed in Section 3.2, we exploit noise-contrastive estimation (NCE) to learn the representations ϕ and μ in (11), which bypasses the need to compute the inte-

gral $\int_{\mathcal{S}} \phi(s, a)^\top \mu(s') p(s') ds'$. Specifically, given transition data $\{(s_i, a_i, s'_i)\}_{i=1}^n$, we solve a version of (10) where $x_i = s'_i$, $u_i = (s_i, a_i)$, $f(s', (s, a)) \propto \phi(s, a)^\top p(s') \mu(s')$. As introduced in Section 2, NCE can maximize an objective like (8) or (10) without requiring computation of the integral $\int_{\mathcal{S}} \phi(s, a) \mu(s') p(s') ds'$ during the training stage. We only need a predefined noise distribution $q(s')$ that has full support on the domain of interest to generate the negative samples. One can also use $q(s') = p(s')$ for simplicity.

A desirable property of NCE pointed out by (Ma & Collins, 2018) is that under the realizability assumption with a valid ground truth probability model, the optimal solution will automatically satisfy

$$\int_{\mathcal{S}} \langle \phi^*(s, a), p(s') \mu^*(s') \rangle = 1.$$

For numerical purposes, we can also enforce this marginalization constraint during the entire training path and not just the endpoint by adding the additional regularizer to the loss

$$\begin{aligned} &\mathbb{E}_{(s,a)} \left[\left(\log \int_{\mathcal{S}} \phi(s, a)^\top \mu(s') p(s') ds' \right)^2 \right] \\ &\approx \frac{1}{n} \sum_{i=1}^n \left(\log \left(\frac{1}{K} \sum_{j=1}^K \phi(s_i, a_i)^\top \mu(s'_{i,j}) \right) \right)^2, \end{aligned}$$

where $\{s'_{i,j}\}_{j=1}^K \sim p(s')$.

Another important property of NCE is that (8) and (10) have a close relationship with MLE, as we will show in Section 5. This will allow us to show that the NCE solution preserves statistical properties of MLE that are key to ensuring the sample complexity of policy optimization.

Finally, we note that there are other estimation methods for unnormalized models, such as score matching (Hyvärinen & Dayan, 2005), contrastive divergence (Hinton, 2002) and adversarial training (Dai et al., 2019), that can be used for representation learning in a linear MDP. However, for most of these estimators, it is not clear how to establish the estimation error bound in terms of the total variation or Hellinger distances, which breaks a key step in establishing the needed statistical property for policy optimization.

4.3. CTRL Algorithm Framework

With the newly proposed modifications, we now describe the algorithmic framework for CTRL. Central to CTRL is the learning of the representations ϕ and μ via NCE (Line 6 in Algorithm 1 and Line 2 in Algorithm 2). In the offline setting, we already have a dataset \mathcal{D} , hence we can directly perform NCE on \mathcal{D} to obtain the (normalized) representation $\hat{\phi}$. However, in the online setting, we need to collect a dataset with sufficient quality (Line 4). Following Uehara et al. (2021), we first draw a sample from the stationary distribution of the policy proposed in the previous iteration

Algorithm 2 CTRL-LCB: Offline Representation Learning

- 1: **Input:** Regularizer λ , Parameter α , Model class \mathcal{M} , Dataset \mathcal{D} , Number of Negative Samples K .
- 2: Learn representation $\hat{\phi}(s, a)$ via (10).
- 3: Set the empirical covariance matrix

$$\hat{\Sigma} = \sum_{(s,a) \in \mathcal{D}} \hat{\phi}(s, a) \hat{\phi}(s, a)^\top + \lambda I.$$

- 4: Set the reward penalty:

$$\hat{b}(s, a) = \min \left(\alpha \sqrt{\hat{\phi}(s, a)^\top \hat{\Sigma}^{-1} \hat{\phi}(s, a)}, 2 \right).$$

- 5: Solve $\hat{\pi} = \operatorname{argmax}_{\pi} V_{\hat{P}, r - \hat{b}}^{\pi}$.
- 6: **Return** $\hat{\pi}$

(i.e., π_{n-1} , which can be done by stopping with probability $1 - \gamma$ at each step when interacting with the environment), then uniformly sample an action to guarantee sufficient exploration. For the optimism and pessimism needed in the online and offline scenarios, it is sufficient to add the standard deviation as a reward bonus (Line 8 in Algorithm 1) or subtract it as a reward penalty (Line 4 in Algorithm 2). The detailed algorithms are given in Algorithm 1 and 2.

Planning Module: We need a planning module for the policy optimization components in Line 9 of Algorithm 1 and Line 5 of Algorithm 2, respectively. Here we provide a practical planning algorithm that exploits the linear structure of the learned representation $\phi(s, a)$, rather than adopt a planning module from existing model-based RL approaches (Chua et al., 2018; Kurutach et al., 2018; Hafner et al., 2019), which only use samples from the learned model and thus are too inefficient. We illustrate optimistic planning in CTRL-UCB as a concrete example; a pessimistic planner for CTRL-LCB is given in Appendix B.

Given the representation $\phi(s, a)$, the Bellman equation is

$$Q^{\pi}(s, a) = r(s, a) + \hat{b}_n(s, a) + \gamma \mathbb{E}_P[V^{\pi}(s')], \quad (12)$$

which implies that Q^{π} no longer lies in the linear space of $\phi(s, a)$ due to the bonus. Fortunately, if we augment the space with $\psi(s, a) = [\phi(s, a), \hat{b}_n(s, a)]$, we can still represent the Q -function linearly with $\psi(s, a)$. However, directly introducing $\hat{b}_n(s, a)$ leads to an extra $\mathcal{O}(d^2)$ memory cost. For practical considerations, we exploit the fact that $\hat{b}_n(s, a)$ is a nonlinear function of $\phi(s, a)$ and parametrize $Q(s, a) = w_1^\top \phi(s, a) + w_2^\top \sigma(w_3^\top \phi(s, a))$, where $\sigma(\cdot)$ is a nonlinear activation function used to approximate the $\hat{b}_n(s, a)$ feature. With this parametrization of the Q family, we conduct fitted Q -evaluation (Munos & Szepesvári, 2008; Antos et al., 2007) for policy evaluation, and use entropy-regularized policy gradient (Nachum et al., 2017; Haarnoja

et al., 2018) to improve the policy. See Appendix B for more details about the practical planning module.

4.4. Practical Implementation

We build our practical online algorithm on top of the actor-critic algorithm with entropy regularizer (Haarnoja et al., 2018). The only change we make is stopping the gradient back-propagation from the temporal-difference learning objective to the backbone of the critic function, i.e., freezing ϕ . We update the backbone network by entirely (10). For the planning module, since some of our tasks are continuous, we train a policy network with policy gradient steps. We list two important design choices here:

- We choose the negative distribution q in (10) to be a mixture of 50% samples from the replay buffer and 50% samples from random trajectories.
- For computation convenience, instead of random sample actions (Line 4 in Algorithm 1), we choose a mixture of random actions and on-policy actions, i.e., $\epsilon \pi_0(a|s) + (1 - \epsilon) \pi(a|s)$, where $\pi_0(a|s)$ is a uniform policy. Detailed ablations are in Appendix F.2.

5. Theoretical Justification

We follow the outline of (Uehara et al., 2021) to provide a theoretical justification for CTRL in terms of its sample complexity in both the online and offline settings.

Theorem 5.1 (Sample Complexity of Online CTRL-UCB). *Assume $P \in \mathcal{M}$ where \mathcal{M} is a finite model class, $|\mathcal{A}| < \infty$. Let $\alpha_n = O(\sqrt{(|\mathcal{A}| + d^2)\gamma \log(|\mathcal{M}|n/\delta)})$, $\lambda_n = O(d \log(|\mathcal{M}|n/\delta))$. As $K \rightarrow \infty$, CTRL-UCB needs at most $\tilde{O}\left(\frac{d^4 |\mathcal{A}|^2}{(1-\gamma)^5 \epsilon^2}\right)$ samples to identify a policy $\hat{\pi}$ such that*

$$V_{P^*, r}^{\pi} - V_{P^*, r}^{\hat{\pi}} \leq \epsilon.$$

Theorem 5.2 (Error Bound of Offline CTRL-LCB). *Assume $P \in \mathcal{M}$ where \mathcal{M} is a finite model class, $|\mathcal{A}| < \infty$. Define*

$$C_{\pi}^* = \operatorname{tr} \left(\mathbb{E}_{d_P} [\phi^*(s, a) \phi^*(s, a)^\top] \right. \\ \left. (\mathbb{E}_{\mathcal{D}} [\phi^*(s, a) \phi^*(s, a)^\top])^{-1} \right),$$

where \mathcal{D} is the empirical measure of the offline data $\{(s_i, a_i, r_i, s'_i)\}_{i=1}^n$. Let $\omega = \max_{a, s} \pi_b(a|s)$ where π_b is the behavior policy. Let $\alpha = \sqrt{(\omega + d^2)\gamma \log |\mathcal{M}|/\delta}$, $\lambda = O(d \log |\mathcal{M}|/\delta)$. As $K \rightarrow \infty$, for any policy π , we have that

$$V_{P^*, r}^{\pi} - V_{P^*, r}^{\hat{\pi}} = \tilde{O} \left(\frac{d^2 \omega}{(1-\gamma)^2} \sqrt{\frac{C_{\pi}^*}{n}} \right).$$

We use $\tilde{O}(\cdot)$ to denote order up to logarithmic factors. We provide a proof sketch, deferring details to Appendix A.

Proof Sketch. The main difference between our proposed method and (Uehara et al., 2021) is that we use NCE for the

representation learning, while Uehara et al. (2021) use an MLE oracle, which is computationally intractable. We first establish the consistency of NCE w.r.t. MLE:

Lemma 5.3. Consider the conditional model $p(x|u)$ defined as $\frac{f(x,u)}{Z_f(u)}$ where $Z_f(u) = \int f(x,u)dx$. Define

$$\hat{p}_n^{\text{MLE}} = \arg \max_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n \log f(x_i, u_i) / Z_f(u_i), \quad (13)$$

and $\hat{p}_n^{\text{NCE}}(x|u) = \frac{\hat{f}_n^{\text{NCE}}(x,u)}{\hat{Z}_n^{\text{NCE}}(u)}$. The estimator obtained with the ranking objective (10) satisfies the following property almost surely:

$$\lim_{K \rightarrow \infty} \hat{p}_n^{\text{NCE,ranking}} = \hat{p}_n^{\text{MLE}} \quad (14)$$

Moreover, if $Z_f(u)$ is a constant function that does not depend on u for all $f \in \mathcal{F}$, then the estimator obtained by (8) satisfies the following property almost surely:

$$\lim_{K \rightarrow \infty} \hat{p}_n^{\text{NCE,binary}} = \hat{p}_n^{\text{MLE}} \quad (15)$$

With Lemma 5.3, we know when $K \rightarrow \infty$, NCE and MLE will return the same estimator. Then following the proof of (Uehara et al., 2021), we can obtain the stated results. \square

We considered the sample complexity with a number of negative samples $K \rightarrow \infty$ in NCE. The non-asymptotic error of NCE in terms of the total variation distance or the Hellinger distance is still an open problem in density estimation. Standard empirical process theory (e.g. (Sugiyama et al., 2012, Chapter 14)) can only show the convergence rate in terms of the objective, but it can be non-trivial to bound the measure of interest with only convergence in the objective. As our purpose is making linear MDPs practical, we leave such theoretical nuances of the original NCE as future work. An alternative direction is assuming that the data have some additional structures like the multi-view structure (Sridharan & Kakade, 2008; Foster et al., 2008; Tosh et al., 2021), which can obtain a finite sample TV error bound of $\mathcal{O}(1/\sqrt{n})$ as shown in the concurrent work (Qiu et al., 2022). However, the multi-view data assumption may constrain the potential applicability of the proposed method. Theorem 5.1 and Theorem 5.2 rely on some conditions about the model class size, action space, and coverage assumption, which we discuss in the following.

Remark (Infinite model class): Agarwal et al. (2020); Uehara et al. (2021) focused on the case when the model class \mathcal{M} is finite and $P \in \mathcal{M}$, which simplifies the theoretical analysis. However, in practice, since we use a function approximator with sufficient expressive power like a neural network, it is likely that $P \in \mathcal{M}$ more generally. When satisfied, the theoretical results will still hold for infinite \mathcal{M} with an additional logarithmic complexity term, such as effective dimension (Zhang, 2002) or Rademacher complexity (Bartlett & Mendelson, 2002).

Remark (Infinite action space): The assumption $|\mathcal{A}| < \infty$ is restrictive. However, there are no known results that remove this dependence for representation learning due to the need for exploration. In principle, we can replace $|\mathcal{A}|$ with the volume of the action set, which would be statistically and computationally manageable provided the dimensionality of the action space was small.

Remark (Coverage assumption): C_π^* in the offline RL scenarios measures the quality of the offline data under the optimal representation ϕ^* . If the offline data span the optimal representation space well, then the spectrum of $(\mathbb{E}_{\mathcal{D}}[\phi^*(s,a)\phi^*(s,a)^\top])^{-1}$ is small, implying that C_π^* is small and we obtain a near optimal policy.

6. Experiments

In this section, we first demonstrate that NCE is capable of capturing environment transition dynamics in a handcrafted environment, and then show that CTRL-UCB achieves state-of-the-art performance in MuJoCo (Brockman et al., 2016) and DeepMind Control Suite (Tassa et al., 2018) benchmarks, compared to both model-based (e.g., PETS (Chua et al., 2018), ME-TRPO (Kurutach et al., 2018)) and model-free RL algorithms (e.g., SAC (Haarnoja et al., 2018)). We also instantiate an offline version of our algorithm CTRL-LCB with just one line of modification from CTRL-UCB and show that it achieves comparable performance with some SoTA algorithms (e.g., CQL (Kumar et al., 2020)). To the best of our knowledge, this is the first practical RL algorithm that leverages the linear MDP assumption to surpass current SoTA model-free and model-based RL algorithms, taking a significant step toward making linear MDP algorithms practical and bridging the gap between theoretical and empirical RL. Note that as discussed before, due to the fact that bonus and entropy term does not lie in the linear space of $\phi(s,a)$, we instantiate two version of our algorithm: CTRL-UCBL uses a linear layer for the Q -network while CTRL-UCBM adopts a 2-layer MLP. We adopt 2 or 3 layers of MLP as Q -network for baselines (e.g., SAC) as a fair comparison. Please refer to Appendix C for the details. We also provide all the learning curves in Appendix F.

6.1. Learning Maze Dynamics with NCE

We first show that learning with the NCE objective is capable of capturing the transition of the environment. The environment we use here is a four-room maze environment, where the state is the position of the agent and the action is the velocity. The transition can be expressed as $s' = s + at + \epsilon$, where t is a fixed time interval and $\epsilon \sim \mathcal{N}(0, I)$. We run CTRL-UCBL for $100K$ steps and the learned transition heatmap is visualized in Figure 1. We see that the high density region is centered around the red square (target position) and cover the green dots (samples drawn from the envi-

Table 1. Performance on various MuJoCo control tasks. All the results are averaged across 4 random seeds and a window size of 10K. Results marked with * is adopted from MBBL (Wang et al., 2019). CTRL achieves strong performance compared with baselines.

		HalfCheetah	Reacher	Humanoid-ET	Pendulum	I-Pendulum
Model-Based RL	ME-TRPO*	2283.7±900.4	-13.4±5.2	72.9±8.9	177.3±1.9	-126.2±86.6
	PETS-RS*	966.9±471.6	-40.1±6.9	109.6±102.6	167.9±35.8	-12.1±25.1
	PETS-CEM*	2795.3±879.9	-12.3±5.2	110.8±90.1	167.4±53.0	-20.5±28.9
	Best MBBL	3639.0±1135.8	-4.1±0.1	1377.0±150.4	177.3±1.9	0.0±0.0
Model-Free RL	PPO*	17.2±84.4	-17.2±0.9	451.4±39.1	163.4±8.0	-40.8±21.0
	TRPO*	-12.0±85.5	-10.1±0.6	289.8±5.2	166.7±7.3	-27.6±15.8
	SAC* (3-layer)	4000.7±202.1	-6.4±0.5	1794.4±458.3	168.2±9.5	-0.2±0.1
Representation RL	DeepSF	4180.4±113.8	-16.8±3.6	168.6±5.1	168.6±5.1	-0.2±0.3
	SPEDE	4210.3±92.6	-7.2±1.1	886.9±95.2	169.5±0.6	0.0±0.0
	CTRL-UCBL	4667.4±341.5	-7.3±0.7	1165.6±144.7	167.8±1.3	0.0±0.0
	CTRL-UCBM (2-layer)	4888.5±205.8	-6.0±0.2	1605.9±225.1	168.1±3.1	0.0±0.0
		Ant-ET	Hopper-ET	S-Humanoid-ET	CartPole	Walker-ET
Model-Based RL	ME-TRPO*	42.6±21.1	1272.5±500.9	-154.9±534.3	160.1±69.1	-1609.3±657.5
	PETS-RS*	130.0±148.1	205.8±36.5	320.7±182.2	195.0±28.0	312.5±493.4
	PETS-CEM*	81.6±145.8	129.3±36.0	355.1±157.1	195.5±3.0	260.2±536.9
	Best MBBL	275.4±309.1	1272.5±500.9	1084.3±77.0	200.0±0.0	312.5±493.4
Model-Free RL	PPO*	80.1±17.3	758.0±62.0	454.3±36.7	86.5±7.8	306.1±17.2
	TRPO*	116.8±47.3	237.4±33.5	281.3±10.9	47.3±15.7	229.5±27.1
	SAC* (3-layer)	2012.7±571.3	1815.5±655.1	834.6±313.1	199.4±0.4	2216.4±678.7
Representation RL	DeepSF	768.1±44.1	548.9±253.3	533.8±154.9	194.5±5.8	165.6±127.9
	SPEDE	806.2±60.2	732.2±263.9	986.4±154.7	138.2±39.5	501.6±204.0
	CTRL-UCBL	3032.6±495.3	775.2±342.7	1178.8±65.4	180.4±26.8	1670.5±824.2
	CTRL-UCBM (2-layer)	1806.8±1488.0	2910.6±323.1	1257.9±262.8	199.4±1.3	3241.2±761.4

Table 2. Performance of on various Deepmind Suite Control tasks. All the results are averaged across four random seeds and a window size of 10K. Comparing with SAC, our method achieves even better performance on sparse-reward tasks.

		cheetah_run	cheetah_run_sparse	walker_run	walker_run_sparse	humanoid_run	hopper_hop
Model-Free RL	PPO	227.7±57.9	5.4±10.8	51.6±1.5	0.0±0.0	1.1±0.0	0.7±0.8
	SAC (2-layer)	222.2±41.0	32.4±27.8	183.0±23.4	53.5±69.3	1.3±0.1	0.4±0.5
	SAC (3-layer)	595.2±96.0	419.5±73.3	700.9±36.6	311.5±361.4	1.2±0.1	28.6±19.5
Representation RL	DeepSF	295.3±43.5	0.0±0.0	27.9±2.2	0.1±0.1	0.9±0.1	0.3±0.1
	CTRL-UCBM (2-layer)	679.0±40.8	446.2±146.3	743.1±53.0	697.0±44.8	11.5±5.4	161.9±76.1




Figure 1. Heatmap of the Learned Transition: The agent transits from orange star to red square and the green dots are the samples drawn from environment simulation.

ronment), which means the representation learned by NCE captures the environment transition and this demonstrates the effectiveness of the representation learning with NCE.

6.2. Dense-Reward MuJoCo Benchmarks

We test our algorithm extensively on the dense-reward MuJoCo benchmark from MBBL. We compare CTRL-UCBL and CTRL-UCBM with model-based RL baselines like PETS, ME-TRPO and model-free RL baselines like SAC and PPO. We also list the best model-based RL results (ex-

cept for iLQR) in MBBL for comparison. All the algorithms run for 200K environment steps. The results are averaged across four random seeds and a window size of 10K.

In Table 1, we can see that CTRL-UCBL achieves the SoTA results among all model-based RL algorithms and significantly improves the prior baselines. Despite the improvement in terms of final return, prior model-based RL methods can hardly learn the behavior such as walking and moving the ant in these two tasks. By contrast, CTRL-UCBL learns the two behaviors successfully. We also compare our algorithm with the SoTA model-free RL method SAC. Our method achieves comparable sometimes better performance in most of the tasks. Lastly, comparing with two representation learning baselines: DeepSF (Barreto et al., 2016) and SPEDE (Ren et al., 2021), CTRL-UCBL also shows dominate performance. However, the performance of CTRL-UCBL is not satisfactory in Humanoid and Hopper environment. It is potentially due to the their hard-exploration nature and the linear parametrized Q failed to capture the bonus for exploration.

In order to evaluate whether the learned representation is

Table 3. Average normalized scores on D4RL locomotion tasks. CTRL-LCB show comparable or better performance especially in expert and medium settings.

Task Name	SAC	BCQ	BEAR	CQL	BRAC	CPC	CTRL-LCB
halfcheetah-medium-v0	-4.3	35.4	38.6	42.8	44.4	44.1	46.3
hopper-medium-v0	0.8	37.1	47.6	58.0	37.2	70.8	49.9
walker2d-medium-v0	0.9	19.2	33.2	79.2	69.7	68.3	75.3
halfcheetah-medium-expert-v0	1.8	59.2	51.7	62.4	24.0	66.6	99.3
hopper-medium-expert-v0	1.6	66.6	4.0	111.0	57.6	112.4	111.3
walker2d-medium-expert-v0	1.9	19.1	10.8	98.7	0.8	93.9	104.2
halfcheetah-medium-replay-v0	-2.4	29.5	36.2	46.2	44.3	45.0	47.5
hopper-medium-replay-v0	3.5	15.7	25.3	48.6	28.6	28.4	31.5
walker2d-medium-replay-v0	1.9	8.3	10.8	26.7	10.1	1.0	22.4

effective for learning the Q -function in practice, we also compare CTRL-UCBM with a 2-layer MLP with the SAC algorithm using a 3-layer MLP for Q -network. Our algorithm still achieves impressive performance compared to SAC. This suggests that learning the Q -function on top of the learned representation is much more effective.

6.3. Sparse-Reward DeepMind Control Suite

To better understand how CTRL perform in a sparse-reward setting, we conduct experiments on the DeepMind Control Suite. We compare CTRL-UCBM with model-free RL methods like SAC (including 2-layer MLP and 3-layer MLP as critic network) and PPO (3 layer of MLP as value network). In addition, we also compare with the representation learning method DeepSF. We run all the algorithms for $200K$ environment steps across four random seeds with a window size of $10K$. CTRL-UCBM outperforms SAC (2-layer) by a large margin even in relatively dense-reward settings like `cheetah-run` and `walker-run`. This margin becomes even larger in sparse-reward setting like `walker-run-sparse` and `hopper-hop`. Even comparing with SAC using 3-layer MLP, CTRL-UCBM also achieves better performance. This again shows the effectiveness of our method.

Table 4. Ablation study on bonus coefficient. Bonus boosts the performance especially in sparse-reward settings.

	<code>cheetah_run</code>	<code>walker_run</code>	<code>walker_run_sparse</code>
CTRL-UCBM ($\alpha = 0$)	656.9 \pm 53.5	703.7 \pm 109.0	118.4 \pm 236.5
CTRL-UCBM ($\alpha = 1$)	615.5 \pm 28.7	700.6 \pm 26.6	390.9 \pm 229.7
CTRL-UCBM ($\alpha = 5$)	679.0\pm40.8	743.1\pm53.0	697.0\pm44.8

6.4. Ablation Study

Our ablation study mainly cares about how does the bonus help in online exploration scenarios. To study this, we ablate on the scaling coefficient α used for balancing bonus and actual reward. The results in Table 6.3 show that in both dense-reward settings like `cheetah-run` and sparse-reward settings like `walker-run-sparse`, adding bonus will consistently improve the performance. Especially in `walker-run-sparse`, the performance boosts from no bonus to $\alpha = 5$ boosts by 583.4. Thus,

bonus term is necessary in sparse-reward settings and such optimism help the online learning process. We also show additional ablations on the dimension of feature and ranking vs. classification objective for NCE in Appendix F.

6.5. Offline Reinforcement Learning

Lastly, we instantiate our CTRL-LCB algorithm in the offline setting on the D4RL benchmark (Fu et al., 2020). Unlike online RL which involves data collection, we split the offline CTRL-LCB into two phases: representation learning phase using NCE and policy update phase. During representation learning phase, we pretrain $\phi(s, a)$ according to Algorithm 2 with for $200K$ steps. After pretraining, we fix the learned representations and use SAC with moderate policy regularization as in Wu et al. (2019) to learn the downstream tasks for for $1M$ steps across 4 random seeds.

We compare CTRL-LCB with SoTA offline RL algorithms (Fujimoto et al., 2019; Kumar et al., 2019; 2020; Wu et al., 2019) and include CPC (Oord et al., 2018) as a representation learning baseline. Besides the state-action representation learned in CTRL-LCB, one major difference between CPC and CTRL-LCB is the additional reward penalty introduced by CTRL-LCB. The average normalized scores are presented in Table 3. CTRL-LCB performs well on the medium-expert set of tasks. The medium set of tasks also benefit from our method. However, CTRL-LCB is unable to improve the performance of medium-replay, potentially due to low-quality data induced representations being incapable of recovering the good policy.

7. Conclusion

In this paper, we tackle the computational difficulty in linear MDP and make linear MDP practical with SoTA performance, while still preserve the desirable statistical benefits.

Acknowledgements

We would like to thank the anonymous reviewers for their detailed and constructive feedback. This project occurred under the Google-BAIR Commons at UC Berkeley.

References

- Agarwal, A., Kakade, S., Krishnamurthy, A., and Sun, W. Flambe: Structural complexity and representation learning of low rank mdps. *arXiv preprint arXiv:2006.10814*, 2020.
- Antos, A., Munos, R., and Szepesvári, C. Fitted q-iteration in continuous action-space mdps. 2007.
- Arbel, M., Zhou, L., and Gretton, A. Generalized energy based models. *arXiv preprint arXiv:2003.05033*, 2020.
- Azar, M. G., Osband, I., and Munos, R. Minimax regret bounds for reinforcement learning. In *International Conference on Machine Learning*, pp. 263–272. PMLR, 2017.
- Barreto, A., Dabney, W., Munos, R., Hunt, J. J., Schaul, T., Van Hasselt, H., and Silver, D. Successor features for transfer in reinforcement learning. *arXiv preprint arXiv:1606.05312*, 2016.
- Bartlett, P. L. and Mendelson, S. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(Nov):463–482, 2002.
- Bellman, R. Dynamic programming. *Science*, 153(3731): 34–37, 1966.
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. Openai gym, 2016.
- Castro, P. S. Scalable methods for computing state similarity in deterministic markov decision processes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 10069–10076, 2020.
- Chua, K., Calandra, R., McAllister, R., and Levine, S. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. *arXiv preprint arXiv:1805.12114*, 2018.
- Dai, B., Shaw, A., Li, L., Xiao, L., He, N., Liu, Z., Chen, J., and Song, L. Sbeed: Convergent reinforcement learning with nonlinear function approximation. In *International Conference on Machine Learning*, pp. 1125–1134. PMLR, 2018.
- Dai, B., Liu, Z., Dai, H., He, N., Gretton, A., Song, L., and Schuurmans, D. Exponential family estimation via adversarial dynamics embedding. *arXiv preprint arXiv:1904.12083*, 2019.
- Dann, C. and Brunskill, E. Sample complexity of episodic fixed-horizon reinforcement learning. *arXiv preprint arXiv:1510.08906*, 2015.
- Dayan, P. Improving generalization for temporal difference learning: The successor representation. *Neural Computation*, 5(4):613–624, 1993.
- Du, S., Krishnamurthy, A., Jiang, N., Agarwal, A., Dudik, M., and Langford, J. Provably efficient rl with rich observations via latent state decoding. In *International Conference on Machine Learning*, pp. 1665–1674. PMLR, 2019.
- Duan, Y., Ke, Z. T., and Wang, M. State aggregation learning from markov transition data. *arXiv preprint arXiv:1811.02619*, 2018.
- Ferns, N., Panangaden, P., and Precup, D. Metrics for finite markov decision processes. In *UAI*, volume 4, pp. 162–169, 2004.
- Foster, D. P., Kakade, S. M., and Zhang, T. Multi-view dimensionality reduction via canonical correlation analysis. 2008.
- Fu, J., Kumar, A., Nachum, O., Tucker, G., and Levine, S. D4rl: Datasets for deep data-driven reinforcement learning. *arXiv preprint arXiv:2004.07219*, 2020.
- Fujimoto, S., Meger, D., and Precup, D. Off-policy deep reinforcement learning without exploration. In *International Conference on Machine Learning*, pp. 2052–2062. PMLR, 2019.
- Gelada, C., Kumar, S., Buckman, J., Nachum, O., and Bellemare, M. G. Deepmdp: Learning continuous latent space models for representation learning. In *International Conference on Machine Learning*, pp. 2170–2179. PMLR, 2019.
- Gutmann, M. and Hyvärinen, A. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pp. 297–304. JMLR Workshop and Conference Proceedings, 2010.
- Gutmann, M. U. and Hyvärinen, A. Noise-contrastive estimation of unnormalized statistical models, with applications to natural image statistics. *Journal of Machine Learning Research*, 13(2), 2012.
- Haarnoja, T., Zhou, A., Abbeel, P., and Levine, S. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, pp. 1861–1870. PMLR, 2018.
- Hafner, D., Lillicrap, T., Fischer, I., Villegas, R., Ha, D., Lee, H., and Davidson, J. Learning latent dynamics for planning from pixels. In *International Conference on Machine Learning*, pp. 2555–2565. PMLR, 2019.

- Hinton, G. E. Training products of experts by minimizing contrastive divergence. *Neural computation*, 14(8):1771–1800, 2002.
- Hyvärinen, A. and Dayan, P. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(4), 2005.
- Jaksch, T., Ortner, R., and Auer, P. Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research*, 11(4), 2010.
- Jin, C., Allen-Zhu, Z., Bubeck, S., and Jordan, M. I. Is q-learning provably efficient? *arXiv preprint arXiv:1807.03765*, 2018.
- Jin, C., Yang, Z., Wang, Z., and Jordan, M. I. Provably efficient reinforcement learning with linear function approximation. In *Conference on Learning Theory*, pp. 2137–2143. PMLR, 2020.
- Kulkarni, T. D., Saedi, A., Gautam, S., and Gershman, S. J. Deep successor reinforcement learning. *arXiv preprint arXiv:1606.02396*, 2016.
- Kumar, A., Fu, J., Tucker, G., and Levine, S. Stabilizing off-policy q-learning via bootstrapping error reduction. *arXiv preprint arXiv:1906.00949*, 2019.
- Kumar, A., Zhou, A., Tucker, G., and Levine, S. Conservative q-learning for offline reinforcement learning. *arXiv preprint arXiv:2006.04779*, 2020.
- Kurutach, T., Clavera, I., Duan, Y., Tamar, A., and Abbeel, P. Model-ensemble trust-region policy optimization. *arXiv preprint arXiv:1802.10592*, 2018.
- LeCun, Y., Chopra, S., Hadsell, R., Ranzato, M., and Huang, F. A tutorial on energy-based learning. *Predicting structured data*, 1(0), 2006.
- Levine, S., Finn, C., Darrell, T., and Abbeel, P. End-to-end training of deep visuomotor policies. *The Journal of Machine Learning Research*, 17(1):1334–1373, 2016.
- Li, Y., He, H., Wu, J., Katabi, D., and Torralba, A. Learning compositional koopman operators for model-based control. *arXiv preprint arXiv:1910.08264*, 2019.
- Ma, Z. and Collins, M. Noise contrastive estimation and negative sampling for conditional models: Consistency and statistical efficiency. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 3698–3707, 2018.
- Mahadevan, S. and Maggioni, M. Proto-value functions: A laplacian framework for learning representation and control in markov decision processes. *Journal of Machine Learning Research*, 8(10), 2007.
- Misra, D., Henaff, M., Krishnamurthy, A., and Langford, J. Kinematic state abstraction and provably efficient rich-observation reinforcement learning. In *International conference on machine learning*, pp. 6961–6971. PMLR, 2020.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *nature*, 518(7540): 529–533, 2015.
- Modi, A., Chen, J., Krishnamurthy, A., Jiang, N., and Agarwal, A. Model-free representation learning and exploration in low-rank mdps. *arXiv preprint arXiv:2102.07035*, 2021.
- Munos, R. and Szepesvári, C. Finite-time bounds for fitted value iteration. *Journal of Machine Learning Research*, 9 (5), 2008.
- Nachum, O. and Dai, B. Reinforcement learning via fenchel-rockafellar duality. *arXiv preprint arXiv:2001.01866*, 2020.
- Nachum, O. and Yang, M. Provable representation learning for imitation with contrastive fourier features. *arXiv preprint arXiv:2105.12272*, 2021.
- Nachum, O., Norouzi, M., Xu, K., and Schuurmans, D. Bridging the gap between value and policy based reinforcement learning. *arXiv preprint arXiv:1702.08892*, 2017.
- Oord, A. v. d., Li, Y., and Vinyals, O. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- Osband, I. and Van Roy, B. On lower bounds for regret in reinforcement learning. *arXiv preprint arXiv:1608.02732*, 2016.
- Puterman, M. L. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- Qiu, S., Wang, L., Bai, C., Yang, Z., and Wang, Z. Contrastive ucb: Provably efficient contrastive self-supervised learning in online reinforcement learning, 2022.
- Ren, T., Zhang, T., Szepesvári, C., and Dai, B. A free lunch from the noise: Provable and practical exploration for representation learning. *arXiv preprint arXiv:2111.11485*, 2021.
- Riou-Durand, L. and Chopin, N. Noise contrastive estimation: Asymptotic properties, formal comparison with mcml. *Electronic Journal of Statistics*, 12(2):3473–3518, 2018.

- Sridharan, K. and Kakade, S. M. An information theoretic framework for multi-view learning. 2008.
- Sugiyama, M., Suzuki, T., and Kanamori, T. *Density ratio estimation in machine learning*. Cambridge University Press, 2012.
- Tassa, Y., Doron, Y., Muldal, A., Erez, T., Li, Y., Casas, D. d. L., Budden, D., Abdolmaleki, A., Merel, J., Lefrancq, A., et al. Deepmind control suite. *arXiv preprint arXiv:1801.00690*, 2018.
- Tosh, C., Krishnamurthy, A., and Hsu, D. Contrastive learning, multi-view redundancy, and linear models. In *Algorithmic Learning Theory*, pp. 1179–1206. PMLR, 2021.
- Uehara, M., Zhang, X., and Sun, W. Representation learning for online and offline rl in low-rank mdps. *arXiv preprint arXiv:2110.04652*, 2021.
- Wang, T., Bao, X., Clavera, I., Hoang, J., Wen, Y., Langlois, E., Zhang, S., Zhang, G., Abbeel, P., and Ba, J. Benchmarking model-based reinforcement learning. *arXiv preprint arXiv:1907.02057*, 2019.
- Watter, M., Springenberg, J. T., Boedecker, J., and Riedmiller, M. Embed to control: A locally linear latent dynamics model for control from raw images. *arXiv preprint arXiv:1506.07365*, 2015.
- Wu, Y., Tucker, G., and Nachum, O. The laplacian in rl: Learning representations with efficient approximations. *arXiv preprint arXiv:1810.04586*, 2018.
- Wu, Y., Tucker, G., and Nachum, O. Behavior regularized offline reinforcement learning. *arXiv preprint arXiv:1911.11361*, 2019.
- Yang, G., Zhang, A., Morcos, A., Pineau, J., Abbeel, P., and Calandra, R. Plan2vec: Unsupervised representation learning by latent plans. In *Learning for Dynamics and Control*, pp. 935–946. PMLR, 2020.
- Yang, L. and Wang, M. Reinforcement learning in feature space: Matrix bandit, kernels, and regret bound. In *International Conference on Machine Learning*, pp. 10746–10756. PMLR, 2020.
- Yang, M., Levine, S., and Nachum, O. Trail: Near-optimal imitation learning with suboptimal data. *arXiv preprint arXiv:2110.14770*, 2021.
- Zhang, A., Lyle, C., Sodhani, S., Filos, A., Kwiatkowska, M., Pineau, J., Gal, Y., and Precup, D. Invariant causal prediction for block mdps. In *International Conference on Machine Learning*, pp. 11214–11224. PMLR, 2020a.
- Zhang, A., McAllister, R., Calandra, R., Gal, Y., and Levine, S. Learning invariant representations for reinforcement learning without reconstruction. *arXiv preprint arXiv:2006.10742*, 2020b.
- Zhang, T. Effective dimension and generalization of kernel learning. In *NIPs*, volume 4, pp. 454–461. Citeseer, 2002.

A. Details of Proofs

A.1. Consistency of NCE towards MLE

Lemma 5.3. Consider the conditional model $p(x|u)$ defined as $\frac{f(x,u)}{Z_f(u)}$ where $Z_f(u) = \int f(x,u)dx$. Define

$$\hat{p}_n^{\text{MLE}} = \arg \max_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n \log f(x_i, u_i) / Z_f(u_i), \quad (16)$$

and $\hat{p}_n^{\text{NCE}}(x|u) = \frac{\hat{f}_n^{\text{NCE}}(x,u)}{\hat{Z}_n^{\text{NCE}}(u)}$. The estimator obtained with the ranking objective (10) satisfies

$$\lim_{K \rightarrow \infty} \hat{p}_n^{\text{NCE,ranking}} = \hat{p}_n^{\text{MLE}} \quad (17)$$

almost surely. Moreover, if $Z_f(u)$ is a constant function that does not depend on u for all $f \in \mathcal{F}$, then the estimator obtained by (8) satisfies

$$\lim_{K \rightarrow \infty} \hat{p}_n^{\text{NCE,binary}} = \hat{p}_n^{\text{MLE}} \quad (18)$$

almost surely.

In the following part, we will show the proof for the binary objective and ranking objective accordingly.

Binary Objective The consistency of the binary objective has been discussed in (Riou-Durand & Chopin, 2018) under general regularity conditions. We include such result for completeness under more practical assumptions.

Proof. First, recall that

$$\left(\hat{f}_n^{\text{NCE}}, \hat{\gamma}_n^{\text{NCE}} \right) = \arg \max_{f \in \mathcal{F}, \gamma \in \mathbb{R}} \frac{1}{n} \sum_{i=1}^n \left[\log(h(x_i, u_i; \gamma)) + \sum_{j=1}^K \log(1 - h(y_{i,j}, u_i; \gamma)) \right], \quad (19)$$

where $h(x, u; \gamma)$ is defined via

$$h(x, u; \gamma) = \frac{f(x, u) \exp(-\gamma) / q(x)}{f(x, u) \exp(-\gamma) / q(x) + K}. \quad (20)$$

Notice that

$$\begin{aligned} \lim_{K \rightarrow \infty} \sum_{j=1}^K \log(1 - h(y_{i,j}, u_i; \gamma)) &= - \lim_{K \rightarrow \infty} \frac{\sum_{j=1}^K \log \left(1 + \frac{f(y_{i,j}, u_i) \exp(-\gamma)}{K q(y_{i,j})} \right)^K}{K} \\ &= - \lim_{K \rightarrow \infty} \frac{\sum_{j=1}^K \left(\frac{f(y_{i,j}, u_i) \exp(-\gamma)}{q(y_{i,j})} \right)}{K} = - \mathbb{E}_{y_i \sim q(y)} \left[\frac{f(y_i, u_i) \exp(-\gamma)}{q(y_i)} \right] = - \exp(-\gamma) \int f(y, u_i) dy. \end{aligned}$$

Meanwhile, we have

$$\begin{aligned} \lim_{K \rightarrow \infty} [\log(h(x_i, u_i; \gamma)) + \log K] &= \log \frac{f(x_i, u_i)}{q(x_i)} - \gamma - \lim_{K \rightarrow \infty} \log \left(1 + \frac{f(y_i, u_i) \exp(-\gamma)}{K q(y_i)} \right) \\ &= \log \frac{f(x_i, u_i)}{q(x_i)} - \gamma. \end{aligned}$$

Hence,

$$\lim_{K \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \left[\log(h(x_i, u_i; \gamma)) + \sum_{j=1}^K \log(1 - h(y_{i,j}, u_i; \gamma)) \right] + \log K = \frac{1}{n} \sum_{i=1}^n \log \frac{f(x_i, u_i)}{q(x_i)} - \gamma - \exp(-\gamma) \int f(y, u_i) dy.$$

Take the derivative w.r.t γ , we know the maximum is achieved when $\hat{\gamma}_n^{\text{NCE}} = \log \int f(y, u_i) dy$. As q is a fixed noise distribution, when $K \rightarrow \infty$, the maximizer \hat{f}_n^{NCE} should also maximize $\frac{1}{n} \log f(x_i, u_i) - \log \int f(y, u_i) dy$, which means it's also the solution of MLE, hence finishes the proof. \square

Ranking Objective We then show the consistency of the ranking objective.

Proof. Notice that, with the law of large numbers and continuous mapping theorem, we have

$$\begin{aligned} \lim_{K \rightarrow \infty} \log \left(\frac{1}{K} \left(f(x_i, u_i)/q(x_i) + \sum_{j=1}^K f(y_{i,j}, u_i)/q(y_{i,j}) \right) \right) &= \lim_{K \rightarrow \infty} \log \left(\frac{1}{K} \left(\sum_{j=1}^K f(y_{i,j}, u_i)/q(y_{i,j}) \right) \right) \\ &= \log \mathbb{E}_{y \sim q(y)} f(y, u_i)/q(y) = \log \int p(y|u_i) dy. \end{aligned}$$

Hence, we have

$$\lim_{K \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \log \frac{f(x_i, u_i)/q(x_i)}{f(x_i, u_i)/q(x_i) + \sum_{j=1}^K f(y_{i,j}, u_i)/q(y_{i,j})} + \log K = \frac{1}{n} \sum_{i=1}^n \left[\log f(x_i, u_i)/q(x_i) - \log \int f(y, u_i) dy \right].$$

With the similar reasoning, we can conclude the proof for the ranking objective. \square

Remark A.1. Notice that, the ranking objective does not require explicit parameter γ to approximate the normalizing constant. Hence, for the conditional model $p(y|x)$ where the normalizing constant $\int p(y|x) dy$ can depend on the conditional variable x , the ranking objective can still provide consistent estimation, while the binary objective cannot provide consistent estimation, as γ is only a scalar that does not depend on x by definition.

B. Planning Module Implementation

As introduced in Section 4.3, with the learned $\phi(s, a)$ fixed, we parametrize the Q -function with

$$Q_W(s, a) = w_1^\top \phi(s, a) + w_2^\top \sigma(w_3^\top \phi(s, a)), \quad (21)$$

where $W = (w_1, w_2, w_3)$ denotes the parameters to be learned. We consider policy $\pi_v(a|s)$, where v denotes the policy parameters.

For practical consideration, we also introduce entropy-regularization in MDP, which will lead to an extra $\log \pi(a|s)$ in the Bellman equation (Haarnoja et al., 2018; Nachum et al., 2017; Dai et al., 2018), *i.e.*,

$$Q^\pi(s, a) = r(s, a) + \hat{b}_n(s, a) + \gamma \mathbb{E}_{P^\pi} [Q^\pi(s', a')] - \log \pi(a|s). \quad (22)$$

In fact, for function completeness, we should add one more nonlinear component for capturing $\log \pi(a|s)$ in (21). Empirically, we noticed that the nonlinear part in parametrization (21) is enough.

We apply fitted Q -Evaluation with target network, which optimizes the loss functions for Q -function in t -th iteration,

$$\ell(Q_W; \mathcal{D}) = \mathbb{E}_{\mathcal{D}} \left[(Q_W(s, a) - \tilde{y}_{s,a} + \log \pi(a|s))^2 \right] \quad (23)$$

where $\tilde{y}_{s,a} := r(s, a) + \hat{b}_n(s, a) + \tilde{Q}(s', a')$, with \tilde{Q} denoting the target network. Then, the policy is updated by policy gradient w.r.t. v as $\mathbb{E}_{\mathcal{D}} [\nabla_v \log \pi_v(a|s) Q_W(s, a)]$ in online case.

For the offline RL case, we keep everything the same for learning Q except using $-\hat{b}_n(s, a)$ as bonus, and we use BRAC (Wu et al., 2019) for policy update with the learned Q .

We emphasize that we only learn the parameters W in Q , and the $\phi(s, a)$ is fixed without back-propagation in the training stage. This is significantly different from the vanilla model-free Q -learning algorithms. In fact, this can largely alleviate the sample complexity since the size of function space is reduced effectively.

Meanwhile, comparing to the vanilla model-based algorithms, where the learned model is used as simulator for generating samples for optimal policy seeking, the proposed planner extracts the representation function $\phi(s, a)$ rather than samples only. This exploits more information from the learned model directly, accelerating the policy learning.

C. Linear and MLP Representations of CTRL-UCB.

As we mentioned in the Section 4, the bonus term and entropy term doesn't necessarily lie in the linear space of the representation thus using a linear layer on top of ϕ might not be sufficient. To deal with this issue, we build two versions of CTRL-UCB: CTRL-UCBL, which adds an additional loss to try to predict bonus from $\phi(s, a)$ using a linear layer; and CTRL-UCBM, which builds a 2-layer MLP for predicting the Q -value from $\phi(s, a)$. We study both algorithms in the MuJoCo setting and report their performance separately. To distinguish from the baseline SAC and illustrate the effectiveness of representation learning, in DeepMind Control Suite, we use the same neural network architecture for both SAC and CTRL-UCBM (both use a 2-layer MLP). The detailed network architecture can be found in Appendix E. Therefore, the




Figure 2. **Environments:** Visualization all the environments we used in our experiments. (left) Hand-designed four-room maze environment. (middle) Standard MuJoCo Ant environment. (right) DM Control Suite Walker environment.

only difference is that the Q -function SAC takes raw state-action pairs as input while CTRL-UCB takes $\phi(s, a)$. We show that CTRL-UCBL has performance comparable to SAC, while CTRL-UCBM significantly outperforms SAC using both a 2-layer and a 3-layer MLP for the Q -function.

D. Environment Visualization

We visualize all the environments used for our experiments for a better understanding. In Figure 2, the left one is the hand-designed four-room maze environment. We need to move the agent from the bottom right start state to the upper left target state. Walls shown in black will block all the moves for the agent. The middle one is the Ant environment used in MuJoCo experiments. The task for this is to control the ant so it can crawl. The right one is the Walker task in the DeepMind Control Suite. We need to control its joints and let it walk/run fast.

E. Experiment Details

In this section, we list all the hyperparameter and network architecture we use for our experiments. For online MuJoCo and DM Control tasks, the hyperparameters can be found at Table 5. For bonus scaling coefficient in MuJoCo, we found that since the environment has super dense reward and it is well manually designed, we don't need any bonus for exploration. Therefore, we set this term to 0 for MuJoCo tasks. However, this bonus is critical to the success of DM Control Suite (especially sparse reward environments) and we put a thorough ablation in Section 6.3. Note that we use exactly the same actor and critic network architecture for all the algorithms in the DM Control Suite experiment.

Table 5. Hyperparameters used for CTRL-UCB in all the environments in MuJoCo and DM Control Suite.

	Hyperparameter Value
NCE temperature	0.2
Bonus Coefficient (MuJoCo)	0.0
Bonus Coefficient (DM Control)	5.0
Actor lr	0.0003
Model lr	0.0003
Actor Network Size (MuJoCo)	(256, 256)
Actor Network Size (DM Control)	(1024, 1024)
NCE Embedding Network Size (MuJoCo)	(1024, 1024, 2048)
NCE Embedding Network Size (DM Control)	(1024, 1024, 1024)
Critic Network Size (MuJoCo)	(2048, 1)
Critic Network Size (DM Control)	(1024, 1)
Discount	0.99
Target Update Tau	0.005
Model Update Tau	0.005
Batch Size	256

The hyperparameters for D4RL benchmark is listed in Table 6.

Table 6. Hyperparameters used for CTRL-LCB in D4RL evaluation.

	Hyperparameter Value
NCE temperature	0.2
Bonus Coefficient	1.0
Actor lr	0.0003
Critic lr	0.001
Representation lr	0.0003
Behavior regularizer	1.
Actor Network Size	(256, 256)
NCE Embedding Network Size	(512, 512, 512)
Critic Network Size	(512, 1)
Discount	0.99
Target Update Tau	0.005
Batch Size	256

F. Additional Experiments

In this section, we provide two additional ablation studies on the feature dimension and classification vs. ranking based NCE objective. We also provide the learning curves for all the ablations and DM Control suite experiments. Note that we didn't compare with SPEDE in this setting since the original paper only contains MuJoCo experiments.

F.1. Additional Ablations

In addition to the bonus scaling factor, we are also interested in other two components of CTRL: (1) Whether there is a preference for using classification-based NCE versus ranking-based NCE? (2) Whether there is a significant effect from the dimension of representation ϕ ? We study these two component separately in Table 7 and Table 8.

According to the table, We found that larger size of feature dimension in general helps the performance. Interestingly, ranking based objective yields a slightly more stable performance than the classification based objective.

Table 7. Ablation study on feature dimension. Larger feature dimension will help the performance in all the tasks. And sometimes it is critical to the success of the algorithm.

	cheetah_run	walker_run	hopper_hop
CTRL-UCBM ($h = 256$)	471.4±85.7	661.2±40.7	99.1±84.8
CTRL-UCBM ($h = 512$)	580.1±50.9	772.5±17.3	116.4±62.4
CTRL-UCBM ($h = 1024$)	679.0±40.8	743.1±53.0	161.9±76.1

Table 8. Ablation study on classification-based and ranking-based NCE objective. We found that in general ranking based objective yields a more stable performance.

	cheetah_run	walker_run	hopper_hop
CTRL-UCBM (classify)	455.2±2.7	630.5±39.1	62.6±34.2
CTRL-UCBM (rank)	679.0±40.8	743.1±53.0	161.9±76.1

F.2. Policy Sampling with Randomness.

One difference between our method and (Uehara et al., 2021) is that we don't uniformly sampling from the policy at each time step to collect data for the representation learning objective. We use mixture of learned policy and uniform policy for action sampling. We here include an ablation on the effect for uniform sampling part in Tab. 9. We can see the uniform action sampling indeed helps in some environments, but the improvement is not consistent across different tasks. We conjecture the entropy-regularized policy may also introduce randomness for uniform action sampling.

Table 9. Ablation study on the sampling protocol. At t -th epoch, the action is sampled from $(1 - \epsilon)\pi(a|s) + \epsilon\pi_0(a)$.

	cheetah_run	walker_run	walker_run_sparse
CTRL-UCBM ($\epsilon = 0$)	679.0±40.8	743.1±53.0	697.0±44.8
CTRL-UCBM ($\epsilon = 0.01$)	614.6±57.2	742.3±22.9	734.9±24.3
CTRL-UCBM ($\epsilon = 0.05$)	632.0±65.7	777.5±15.7	662.2±166.3

F.3. Performance Curves for DM Control

We provide all the performance curves for DM Control experiments. In Figure 3, we clearly see that CTRL-UCBM surpasses all the baseline methods and the gain is significantly in sparse-reward settings (e.g., in walker-run-sparse and hopper-hop, CTRL-UCBM surpasses SAC with a good margin). We also provide all the performance curves for ablation studies. In Figure 4, we see that the bonus is critical to the success of the algorithm in sparse-reward settings. In Figure 5, we can conclude that larger feature dimension induces better performance. In Figure 7, we can see that ranking-based objective has slightly better performance than classification-based objective.

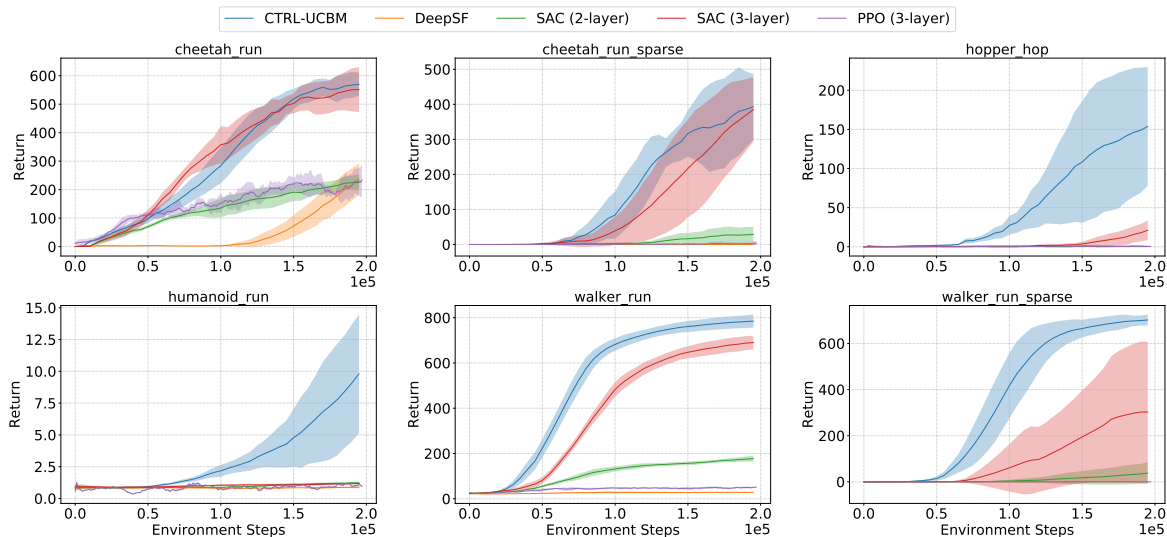


Figure 3. Performance for the DM Control Suite Experiments: We plot the performance curve for DM Control Suite experiments. We can see that CTRL-UCBM achieves the best performance in all the tasks, even comparing to SAC with 3-layer MLP.

F.4. Performance Curves for D4RL Benchmarks.

We also provide learning curves for all the algorithms in D4RL benchmarks. The benefit of CTRL-LCB is most noticeable in medium-expert tasks.




Figure 4. Performance Curve for Bonus Scaling Coefficient: We study the effect of bonus with changing the bonus scaling factor. We see that with 0 bonus, the algorithm hardly learns in cheetah-run-sparse and walker-run-sparse. This shows the effectiveness of bonus in online exploration settings.

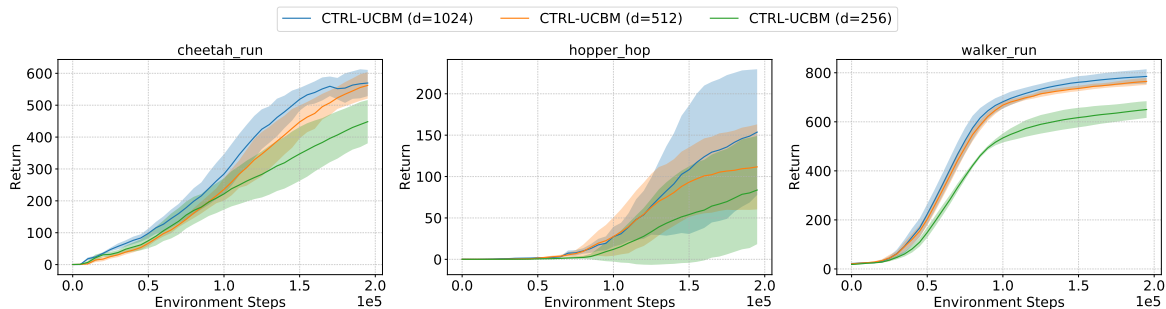


Figure 5. Performance Curve for Varying Feature Dimension: We study the effect of feature dimension in the linear MDP settings. In general the trend is pretty clear: larger feature dimension induces better performance.




Figure 6. Performance Curve for Ranking-Based Objective versus Classification-Based Objective: We study the effect of using a ranking-based objective for NCE versus classification-based objective for NCE. Interestingly, ranking-based objective results in a slightly stabler performance.

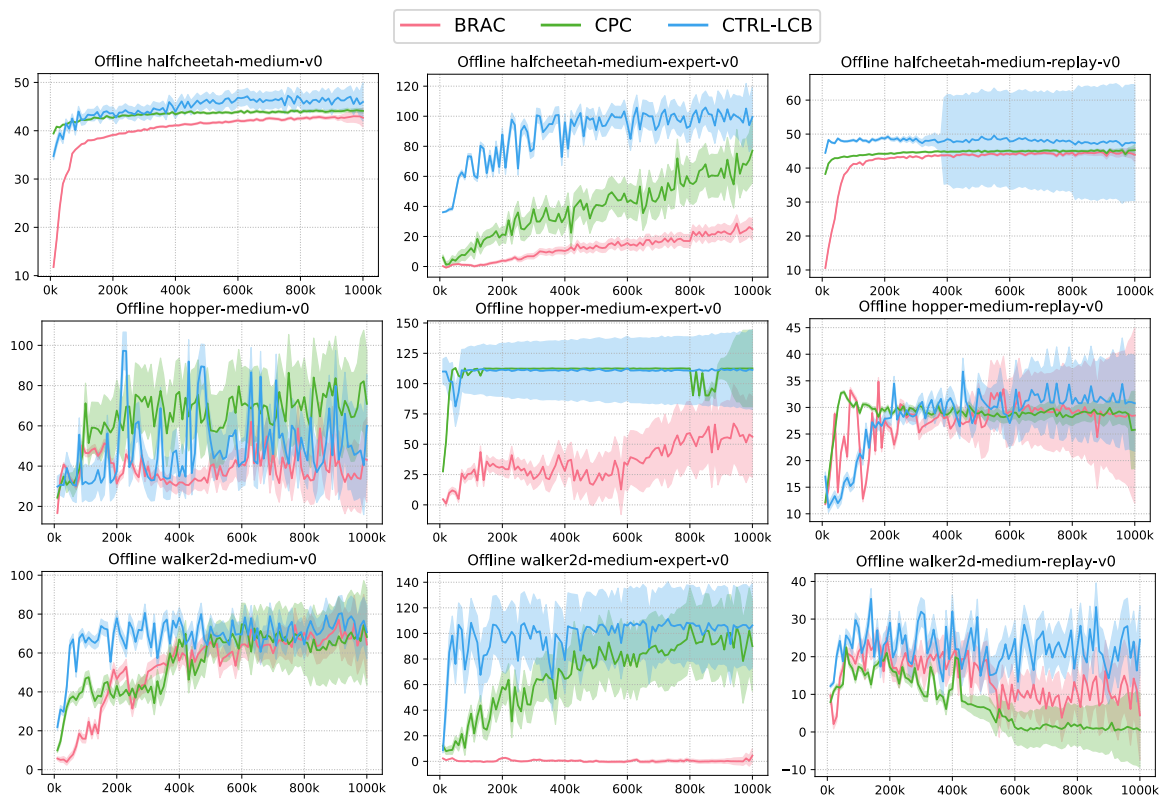


Figure 7. Performance Curve for CTRL-LCB on D4RL: The benefit of CTRL-LCB is the most evident on medium-expert tasks, when the dataset covers some amount of expert demonstration.