
MADE: Exploration via Maximizing Deviation from Explored Regions

Tianjun Zhang^{*1}

Paria Rashidinejad^{*1}

Jiantao Jiao¹

Yuandong Tian²

Joseph E. Gonzalez¹

Stuart Russell¹

¹University of California, Berkeley
{tianjunz, paria.rashidinejad, jiantao, russell}@berkeley.edu

²Facebook AI Research
yuandong@fb.com

Abstract

In online reinforcement learning (RL), efficient exploration remains particularly challenging in high-dimensional environments with sparse rewards. In low-dimensional environments, where tabular parameterization is possible, count-based upper confidence bound (UCB) exploration methods achieve minimax near-optimal rates. However, it remains unclear how to efficiently implement UCB in realistic RL tasks that involve nonlinear function approximation. To address this, we propose a new exploration approach via *maximizing* the deviation of the occupancy of the next policy from the explored regions. We add this term as an adaptive regularizer to the standard RL objective to trade off between exploration and exploitation. We pair the new objective with a provably convergent algorithm, giving rise to a new intrinsic reward that adjusts existing bonuses. The proposed intrinsic reward is easy to implement and combine with other existing RL algorithms to conduct exploration. As a proof of concept, we evaluate the new intrinsic reward on tabular examples across a variety of model-based and model-free algorithms, showing improvements over count-only exploration strategies. When tested on navigation and locomotion tasks from MiniGrid and DeepMind Control Suite benchmarks, our approach significantly improves sample efficiency over state-of-the-art methods.²

1 Introduction

Online RL is a useful tool for an agent to learn how to perform tasks, particularly when expert demonstrations are unavailable and reward information needs to be used instead [92]. To learn a satisfactory policy, an RL agent needs to effectively balance between exploration and exploitation, which remains a central question in RL [23, 15]. Exploration is particularly challenging in environments with sparse rewards. One popular approach to exploration is based on *intrinsic motivation*, often applied by adding an intrinsic reward (or bonus) to the extrinsic reward provided by the environment. In provable exploration methods, bonus often captures the value estimate uncertainty and the agent takes an action that maximizes the upper confidence bound (UCB) [5, 8, 41, 48, 44]. In tabular setting, UCB bonuses are often constructed based on either Hoeffding’s inequality, which only uses visitation counts, or Bernstein’s inequality, which uses value function variance in addition to visitation counts.

^{*}Equal Contribution.

²Our code is available at <https://github.com/tianjunz/MADE>.




Figure 1: Normalized samples use of different methods with respect to MADE (smaller values are better). MADE consistency achieves a better sample efficiency compared to all other baselines. Infinity means the method fails to achieve maximum reward in given steps.

The latter is proved to be minimax near-optimal in environments with bounded rewards [44, 65] as well as bounded total reward [112] and reward-free settings [64, 49, 45, 113]. It remains an open question how one can efficiently compute confidence bounds to construct UCB bonus in non-linear function approximation. Furthermore, Bernstein-style bonuses are often hard to compute in practice beyond tabular setting, due to difficulties in computing value function variance.

In practice, various approaches are proposed to design intrinsic rewards: visitation pseudo-count bonuses estimate count-based UCB bonuses using function approximation [10, 15], curiosity-based bonuses seek states where model prediction error is high, uncertainty-based bonuses [77, 87] adopt ensembles of networks for estimating variance of the Q-function, empowerment-based approaches [52, 34, 84, 68] lead the agent to states over which the agent has control, and information gain bonuses [51] reward the agent based on the information gain between state-action pairs and next states.

Although the performance of practical intrinsic rewards is good in certain domains, empirically they are observed to suffer from issues such as detachment, derailment, and catastrophic forgetting [3, 23]. Moreover, these methods usually lack a clear objective and can get stuck in local optimum [3]. Indeed, the impressive performance currently achieved by some deep RL algorithms often revolves around manually designing dense rewards [13], complicated exploration strategies utilizing a significant amount of domain knowledge [23], or operating in the known environment regime [88, 69].

Motivated by current practical challenges and the gap between theory and practice, we propose a new algorithm for exploration by maximizing deviation from explored regions. This yields a practical algorithm with strong empirical performance. To be specific, we make the following contributions:

1. Exploration via maximizing deviation Our approach is based on modifying the standard RL objective (i.e. the cumulative reward) by adding a regularizer that adaptively changes across iterations. The regularizer can be a general function depending on the state-action visitation density and previous state-action coverage. We show that the regularized objective naturally admits several common existing exploration methods. We then choose a particular regularizer that **MA**ximizes the **DE**viation (MADE) of the next policy visitation d^π from the regions covered by prior policies ρ_{cov}^k :

$$L_k(d^\pi) = J(d^\pi) + \tau_k \sum_{s,a} \sqrt{\frac{d^\pi(s,a)}{\rho_{\text{cov}}^k(s,a)}}. \quad (1)$$

Here, k is the iteration number, $J(d^\pi)$ is the standard RL objective, and the regularizer encourages $d^\pi(s,a)$ to be large when $\rho_{\text{cov}}^k(s,a)$ is small. We give an algorithm for solving the regularized objective and prove that with access to an approximate planning oracle, it converges to the global optimum. We show that objective (1) results in an intrinsic reward that can be easily added to any RL algorithm to improve performance, as suggested by our empirical studies. Furthermore, the intrinsic reward applies a simple modification to the UCB-style bonus that considers prior visitation counts. This simple modification can also be added to existing bonuses in practice.

2. Tabular studies In the special case of tabular parameterization, we show that MADE only applies some simple adjustments to the Hoeffding-style count-based bonus. We compare the performance of MADE to Hoeffding and Bernstein bonuses in three different RL algorithms, for the exploration task in the stochastic diabolical bidirectional lock [3, 66], which has sparse rewards and local optima. Our results show that MADE robustly improves over the Hoeffding bonus and is competitive to the Bernstein bonus, across all three RL algorithms. Interestingly, MADE bonus and exploration strategy

appear to be very close to the Bernstein bonus, *without computing or estimating variance*, suggesting that MADE potentially captures some environmental structures. Additionally, we empirically show that MADE regularizer can improve the optimization rate in policy gradient methods.

3. Experiments on MiniGrid and DeepMind Control Suite We empirically show that MADE works well when combined with model-free (IMAPLA [25], RAD [55]) and model-based (Dreamer [35]) RL algorithms, greatly improving the sample efficiency over existing baselines. When tested in the procedurally-generated MiniGrid environments, MADE manages to converge with two to five times fewer samples compared to state-of-the-art method BeBold [111]. In DeepMind Control Suite [95], we build upon the model-free method RAD [55] and the model-based method Dreamer [35], improving the return up to 150 in 500K steps compared to baselines. Figure 1 shows normalized sample size to achieve maximum reward with respect to our algorithm.

2 Background

Markov decision processes. An infinite-horizon discounted MDP is described by a tuple $M = (\mathcal{S}, \mathcal{A}, P, r, \rho, \gamma)$, where \mathcal{S} is the state space, \mathcal{A} is the action space, $P : \mathcal{S} \times \mathcal{A} \mapsto \Delta(\mathcal{S})$ is the transition kernel, $r : \mathcal{S} \times \mathcal{A} \mapsto [0, 1]$ is the (extrinsic) reward function, $\rho : \mathcal{S} \mapsto \Delta(\mathcal{S})$ is the initial distribution, and $\gamma \in [0, 1]$ is the discount factor. A stationary (stochastic) policy $\pi \in \Delta(\mathcal{A} | \mathcal{S})$ specifies a distribution over actions in each state. Each policy π induces a visitation density over state-action pairs $d^\pi : \mathcal{S} \times \mathcal{A} \mapsto [0, 1]$ defined as $d^\pi_\rho(s, a) := (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t \mathbb{P}_t(s_t = s, a_t = a; \pi)$, where $\mathbb{P}_t(s_t = s, a_t = a; \pi)$ denotes (s, a) visitation probability at step t , starting at $s_0 \sim \rho(\cdot)$ and following π . An important quantity is the value a policy π , which is the discounted sum of rewards $V^\pi(s) := \mathbb{E}[\sum_{t=0}^{\infty} \gamma^t r_t | s_0 = s, a_t \sim \pi(\cdot | s_t) \text{ for all } t \geq 0]$ starting at state $s \in \mathcal{S}$.

Policy mixture. For a sequence of policies $\mathcal{C}^k = (\pi_1, \dots, \pi_k)$ with corresponding mixture distribution $w^k \in \Delta_{k-1}$, the policy mixture $\pi_{\text{mix}, k} = (\mathcal{C}^k, w^k)$ is obtained by first sampling a policy from w^k and then following that policy over subsequent steps [36]. The mixture policy induces a state-action visitation density according to $d^{\pi_{\text{mix}}}(s, a) = \sum_{i=1}^k w_i^k d^{\pi_i}(s, a)$. While the π_{mix} may not be stationary in general, there exists a stationary policy π' such that $d^{\pi'} = d^{\pi_{\text{mix}}}$ Puterman [80].

Online reinforcement learning. Online RL is the problem of finding a policy with a maximum value from an unknown MDP, using samples collected during exploration. Oftentimes, the following objective is considered, which is a scalar summary of the performance of policy π :

$$J_M(\pi) := \mathbb{E}_{s \sim \rho}[V^\pi(s)] = (1 - \gamma)^{-1} \mathbb{E}_{(s,a) \sim d^\pi(\cdot, \cdot)}[r(s, a)] = J(d^\pi). \quad (2)$$

We drop index M when it is clear from context. We denote an optimal policy by $\pi^* \in \arg \max_\pi J(\pi)$ and use the shorthand $V^* := V^{\pi^*}$ to denote the optimal value function. It is straightforward to check that $J(\pi)$ can equivalently be represented by the expectation of the reward over the visitation measure of π . We slightly abuse the notation and sometimes write $J(d^\pi)$ to denote the RL objective.

3 Adaptive regularization of the RL objective

3.1 Regularization to guide exploration

In online RL, the agent faces a dilemma in each step: whether it should select a seemingly optimal policy (exploit) or it should explore different regions of the MDP. To allow flexibility in this choice and trade off between exploration and exploitation, we propose to add a regularizer to the standard RL objective that changes throughout iterations of an online RL algorithm:

$$L_k(d^\pi) = \underbrace{J(d^\pi)}_{\text{exploitation}} + \tau_k \underbrace{R(d^\pi; \{d^{\pi_i}\}_{i=1}^k)}_{\text{exploration}}. \quad (3)$$

Here, $R(d^\pi; \{d^{\pi_i}\}_{i=1}^k)$ is a function of state-action visitation of π as well as the visitation of prior policies π_1, \dots, π_k . The temperature parameter τ_k determines the strength of regularization. Objective (3) is a *population* objective in the sense that it does not involve empirical estimations affected by the randomness in sample collection. In the following section, we give our particular choice of regularizer and discuss how this objective can describe some popular exploration bonuses. We then provide a convergence guarantee for the regularized objective in Section 3.2.

3.2 Exploration via maximizing deviation from policy cover

We develop our exploration strategy MADE based on a simple intuition: maximizing the deviation from the explored regions, i.e. all states and actions visited by prior policies. We define *policy cover* at iteration k to be the density over regions explored by policies π_1, \dots, π_k , i.e. $\rho_{\text{cov}}^k(s, a) := \frac{1}{k} \sum_{i=1}^k d^{\pi_i}(s, a)$. We then design our regularizer to encourage d^π to be different from ρ_{cov}^k :

$$R_k(d^\pi; \{d^{\pi_i}\}_{i=1}^k) = \sum_{s,a} \sqrt{\frac{d^\pi(s, a)}{\rho_{\text{cov}}^k(s, a)}}. \quad (4)$$

It is easy to check that the maximizer of above function is $d^\pi(s, a) \propto \frac{1}{\rho_{\text{cov}}^k(s, a)}$. Our motivation behind this particular deviation is that it results in a simple modification of UCB bonus in tabular case.

We now compute the reward yielded by the new objective. First, define a policy mixture $\pi_{\text{mix},k}$ with policy sequence (π_1, \dots, π_k) and weights $((1-\eta)^{k-1}, (1-\eta)^{k-2}\eta, (1-\eta)^{k-3}\eta, \dots, \eta)$ for $\eta > 0$. Let $d^{\pi_{\text{mix},k}}$ be the visitation density of $\pi_{\text{mix},k}$. We compute the total reward at iteration k by taking the gradient of the new objective with respect to d^π at $d^{\pi_{\text{mix},k}}$:

$$r_k(s, a) = (1-\gamma) \nabla_d L_k(d) \Big|_{d=d^{\pi_{\text{mix},k}}} = r(s, a) + (1-\gamma) \tau_k \nabla_d R_k(d; \{d^{\pi_i}\}_{i=1}^k) \Big|_{d=d^{\pi_{\text{mix},k}}}, \quad (5)$$

which gives the following reward

$$r_k(s, a) = r(s, a) + \frac{(1-\gamma) \tau_k / 2}{\sqrt{d^{\pi_{\text{mix},k}}(s, a) \rho_{\text{cov}}^k(s, a)}}. \quad (6)$$

The intrinsic reward is constructed based on two densities: ρ_{cov}^k a uniform combination of past visitation densities and $\hat{d}^{\pi_{\text{mix},k}}$ a weighted mixture of the past visitation densities. As we will discuss shortly, policy cover $\rho_{\text{cov}}^k(s, a)$ is related to the visitation count of (s, a) pair in previous iterations and resembles count-based bonuses [10, 44] or their approximates such as RND [15]. Therefore, for an appropriate choice of τ_k , MADE intrinsic reward decreases as the number of visitations increases.

MADE intrinsic reward is also proportional to $1/\sqrt{\hat{d}^{\pi_{\text{mix},k}}(s, a)}$, which can be viewed as a correction applied to the count-based bonus. In effect, due to the decay of weights in $\pi_{\text{mix},k}$, the above construction gives a higher reward to (s, a) pairs visited earlier. Experimental results suggest that this correction may alleviate major difficulties in sparse reward exploration, namely detachment and catastrophic forgetting, by encouraging the agent to revisit forgotten states and actions.

Empirically, MADE's intrinsic reward is computed based on estimates $\hat{d}^{\pi_{\text{mix},k}}$ and $\hat{\rho}_{\text{cov}}^k$ from data collected by iteration k . Furthermore, practically we consider a smoothed version of the above regularizer by adding $\lambda > 0$ to both numerator and denominator; see Equation (7).

MADE intrinsic reward in tabular case. In tabular empirical setting, the empirical estimation of policy cover is simply $\hat{\rho}_{\text{cov}}^k(s, a) = \frac{N_k(s, a)}{N_k}$, where $N_k(s, a)$ is (s, a) pair's visitation count and N_k is the total count, by iteration k . Setting $\tau_k = 1/\sqrt{N_k}$, MADE simply modifies the Hoeffding-type bonus via the mixture density and is proportional to $1/\sqrt{\hat{d}^{\pi_{\text{mix},k}}(s, a) N_k(s, a)}$.

Bernstein bonus is another tabular UCB bonus that modifies Hoeffding bonus via an empirical estimate of the value function variance. Bernstein bonus is shown to improve over Hoeffding count-only bonus by exploiting additional environment structure [106] and close the gap between algorithmic upper bounds and information-theoretic limits up to logarithmic factors [112, 113]. However, a practical and efficient implementation of a bonus that exploits variance information in non-linear function approximation parameterization still remains an open question; see Section 6 for further discussion. On the other hand, our proposed modification based on the mixture density can be easily and efficiently incorporated with non-linear parameterization.

Deriving some popular bonuses from regularization. The regularization in (3) can describe some popular bonuses. Exploration bonuses that only depend on state-action visitation counts can be expressed in the form (3) by setting the regularizer a linear function of d^π and the exploration bonus $r_i(s, a)$, i.e., $R_k(d^\pi; \{d^{\pi_i}\}_{i=1}^k) = \sum_{s,a} d^\pi(s, a) r_i(s, a)$. One can check that taking the gradient of the regularizer with respect to d^π recovers $r_i(s, a)$. As another example, one can set the regularizer to Shannon entropy $R_k(d^\pi; \{d^{\pi_i}\}_{i=1}^k) = -\sum_{s,a} d^\pi(s, a) \log d^\pi(s, a)$, which gives the intrinsic reward $-\log d^\pi(s, a)$ (up to an additive constant) and recovers the result in Zhang et al. [109].

Algorithm 1 Policy computation for adaptively regularized objective

- 1: **Inputs:** Iteration count K , planning error ϵ_p , visitation density error ϵ_d .
 - 2: Initialize policy mixture $\pi_{\text{mix},1}$ with $\mathcal{C}_1 = (\pi_1)$ and $w^1 = (1)$
 - 3: **for** $k = 1, \dots, K$ **do**
 - 4: Estimate the visitation density $\hat{d}^{\pi_{\text{mix},k}}$ of $\pi_{\text{mix},k}$ via a visitation density oracle.
 - 5: Compute reward $r_k(s, a) = r(s, a) + (1 - \gamma)\tau_k \nabla_d R_k(d; \{\pi_i\}_{i=1}^k)|_{d=\hat{d}^{\pi_{\text{mix},k}}}$.
 - 6: Run approximate planning on modified MDP $M^k = (\mathcal{S}, \mathcal{A}, P, r_k, \gamma)$ and return π_{k+1} .
 - 7: Update policy mixture $\mathcal{C}^{k+1} = (\mathcal{C}_k, \pi_{k+1})$ and $w^{k+1} = ((1 - \eta)w^k, \eta)$.
 - 8: **Return:** $\pi_{\text{mix},K} = (\mathcal{C}^K, w^K)$.
-

3.3 Solving the regularized objective

Recall that our goal is to find a policy that maximizes the regularized objective $\pi^* \in \arg \max_{\pi} L_k(d^{\pi}) = J(d^{\pi}) + \tau_k R(d^{\pi}; \{d^{\pi_i}\}_{i=1}^k)$. Despite being nonconcave in π , this objective is concave in d^{π} . Therefore, one can solve the following constrained concave optimization problem instead: $\max_{d^{\pi} \in \mathcal{V}} L_k(d^{\pi})$, where \mathcal{V} is the set of all valid visitation densities.

To solve this constrained optimization problem, we use the conditional gradient method or the Frank-Wolfe algorithm [31], which has been used in the context of RL in works such as [36, 1, 105]. The conditional gradient method involves iteratively solving $d^{\pi_{k+1}} \in \arg \max_d \langle d, \nabla_{d^{\pi}} L_k(d^{\pi})|_{d^{\pi}=d^{\pi_{\text{mix},k}}} \rangle$ and obtaining $d^{\pi_{\text{mix},k+1}}$ from a weighted combination of $d^{\pi_{\text{mix},k}}$ and $d^{\pi_{k+1}}$. Note that the first step is equivalent to planning via a reward function proportional to $\nabla_{d^{\pi}} L_k(d^{\pi})|_{d^{\pi}=d^{\pi_{\text{mix},k}}}$, which justifies they way we computed the total reward in Equation [5].

We provide convergence guarantee for Algorithm [1] in the following theorem whose proof is given in Appendix A.

Theorem 1. Consider the following regularizer for [3] with $\lambda > 0$ and a valid visitation density d

$$R_{\lambda}(d; \{d^{\pi_i}\}_{i=1}^k) = \sum_{s,a} \sqrt{\frac{d(s,a) + \lambda}{\rho_{\text{cov}}^k(s,a) + \lambda}}, \quad (7)$$

Set $\tau_k = \tau/k^c$, where $0 < \tau < 1$ and $c > 0$. For any $\epsilon > 0$ and $\eta \leq 0.4\epsilon\lambda^2$, there exists $\epsilon_p, \epsilon_d, c, B$ such that $\pi_{\text{mix},K}$ returned by Algorithm [1] after $K \geq \eta^{-1} \log(10B\epsilon^{-1})$ iterations satisfies $L_k(d^{\pi_{\text{mix},K}}) \geq \max_{\pi} L_k(d^{\pi}) - \epsilon$.

Remark 1. One does not need to maintain the functional forms of past policies to estimate $\hat{d}^{\pi_{\text{mix},k}}$. Practically, one may truncate the dataset to a (prioritized) buffer and estimate the density over that.

4 A tabular study

We first study the performance of MADE in tabular toy examples. In the Bidirectional Lock experiment, we compare MADE to theoretically guaranteed Hoeffding-style and Bernstein-style bonuses in a sparse reward exploration task. In the Chain MDP, we investigate whether MADE's regularizer [4] provides any benefits in improving optimization rate in policy gradient methods.




Figure 2: In a bidirectional lock, the agent starts at s_0 and enters one of the chains based on the selected action. Each chain has a positive reward at the end, H good states, and H dead states. Both actions available to the agent lead it to the dead state, one with probability one and the other with probability $p < 1$. MADE performs better than Hoeffding-style bonus and comparable to Bernstein-style bonus across all three RL algorithms.

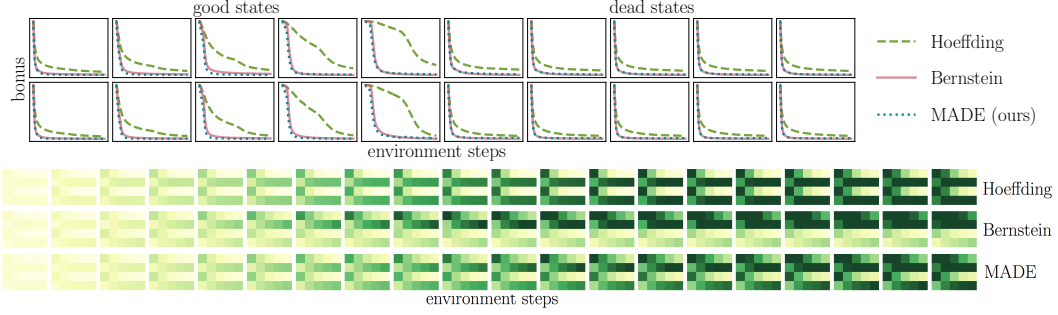


Figure 3: Hoeffding, Bernstein, and MADE intrinsic rewards over iterations in a bidirectional lock.

4.1 Exploration in bidirectional lock

We consider a stochastic version of the bidirectional combination lock (Figure 2), which is considered a particularly difficult exploration task in tabular setting [66, 3]. This environment is challenging because: (1) positive rewards are sparse, (2) a small negative reward is given when transitioning to a good state and thus, moving to a dead state is locally optimal, and (3) the agent may forget to explore one chain and get stuck in local minima upon receiving an end reward in one lock [3].

RL algorithms and exploration strategies. We compare the performance of Hoeffding and Bernstein bonuses [44] to MADE in three different RL algorithms. To implement MADE in tabular setting, we simply use two buffers: one that stores all past state-action pairs to estimate ρ_{cov} and another one that only maintains the most recent B pairs to estimate d_{μ}^{π} . We use empirical counts to estimate both densities, which give a bonus $\propto 1/\sqrt{N_k(s, a)B_k(s, a)}$, where $N_k(s, a)$ is the total count and $B_k(s, a)$ is the recent buffer count of (s, a) pair. We combine three bonuses with three RL algorithms: (1) value iteration with bonus [37], (2) proximal policy optimization (PPO) with a model [16], and (3) Q-learning with bonus [44].

Results. Figure 2 summarizes our results showing MADE improves over the Hoeffding bonus and is competitive to the Bernstein bonus in all three algorithms. Unlike Bernstein bonus that is hard to compute beyond tabular setting, MADE bonus design is simple and can be effectively combined with any deep RL algorithm. The experimental results suggest several interesting properties for MADE. First, MADE applies a simple modification to the Hoeffding bonus which improves the performance. Second, as illustrated in Figure 3, bonus values and exploration pattern of MADE is somewhat similar to the Bernstein bonus. This suggests that MADE may capture some structural information of the environment, similar to Bernstein bonus, which captures certain environmental properties such as the degree of stochasticity [106].

4.2 Policy gradient in chain MDP

We consider the chain MDP (Figure 4) presented in [2], which suffers from vanishing gradients with policy gradient approach [93] as a positive reward is only achieved if the agent always takes action a_1 . This leads to an exponential iteration complexity lower bound on the convergence of vanilla policy

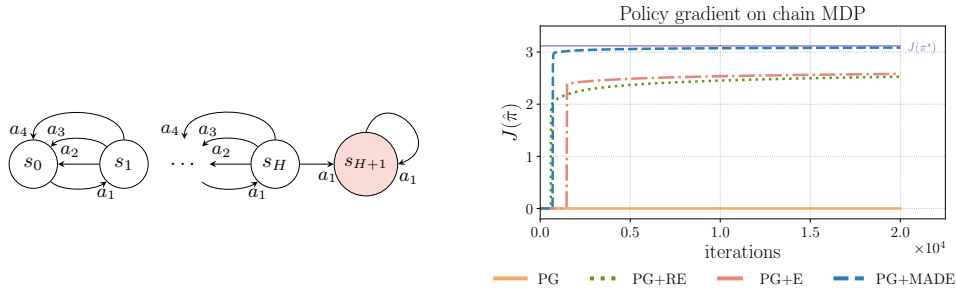


Figure 4: A deterministic chain MDP that suffers from vanishing gradients [2]. We consider a constrained tabular policy parameterization with $\pi(a|s) = \theta_{s,a}$ and $\sum_a \theta_{s,a} = 1$. The agent always starts from s_0 and the only non-zero reward is $r(s_{H+1}, a_1) = 1$.

gradient approach even with access to exact gradients [2]. In this environment the agent always starts at state s_0 and recent guarantees on the global convergence of exact policy gradients are vacuous [11, 2, 62]. This is because the rates depend on the ratio between the optimal and learned visitation densities, known as *concentrability coefficient* [47, 85, 32, 82], or the ratio between the optimal visitation density and initial distribution [2].

RL algorithms. Since our goal in this experiment is to investigate the optimization effects and not exploration, we assume access to exact gradients. In this setting, we consider MADE regularizer with the form $\sum_{s,a} \sqrt{d^\pi(s,a)}$. Note that policy gradients take gradient of the objective with respect to the policy parameters θ and not d^π . We compare optimizing the policy gradient objective with four methods: vanilla version PG (e.g. uses policy gradient theorem [99, 93, 53]), relative policy entropy regularization PG+RE [2], policy entropy regularization PG+E [67, 62], and MADE regularization.

Results. Figure 4 illustrates our results on policy gradient methods. As expected [2], the vanilla version has a very slow convergence rate. Both entropy and relative entropy regularization methods are proved to achieve a linear convergence rate of $\exp(-t)$ in the iteration count t [62, 2]. Interestingly, MADE seems to outperform the policy entropy regularizers, quickly converging to a globally optimal policy.

5 Experiments on MiniGrid and DeepMind Control Suite

In addition to the tabular setting, MADE can also be integrated with various model-free and model-based deep RL algorithms such as IMPALA [25], RAD [57], and Dreamer [35]. As we will see shortly, MADE exploration strategy on MiniGrid [19] and DeepMind Control Suite [95] tasks achieves state-of-the-art sample efficiency.

For a practical estimation of ρ_{cov}^k and $d^{\pi_{\text{mix},k}}$, we adopt the two buffer idea described in the tabular setting. However, since now the state space is high-dimensional, we use RND [15] to estimate $N_k(s,a)$ (and thus ρ_{cov}^k) and use a variational auto-encoder (VAE) to estimate $d^{\pi_{\text{mix},k}}$. Specifically, for RND, we minimize the difference between a predictor network $\phi'(s,a)$ and a randomly initialized target network $\phi(s,a)$ and train it in an online manner as the agent collects data. We sample data from the recent buffer \mathcal{B} to train a VAE. The length of \mathcal{B} is a design choice for which we do an ablation study. Thus, the intrinsic reward in deep RL setting takes the following form

$$(1 - \gamma)\tau_k \frac{\|\phi(s,a) - \phi'(s,a)\|}{\sqrt{d^{\pi_{\text{mix},k}}(s,a)}}.$$

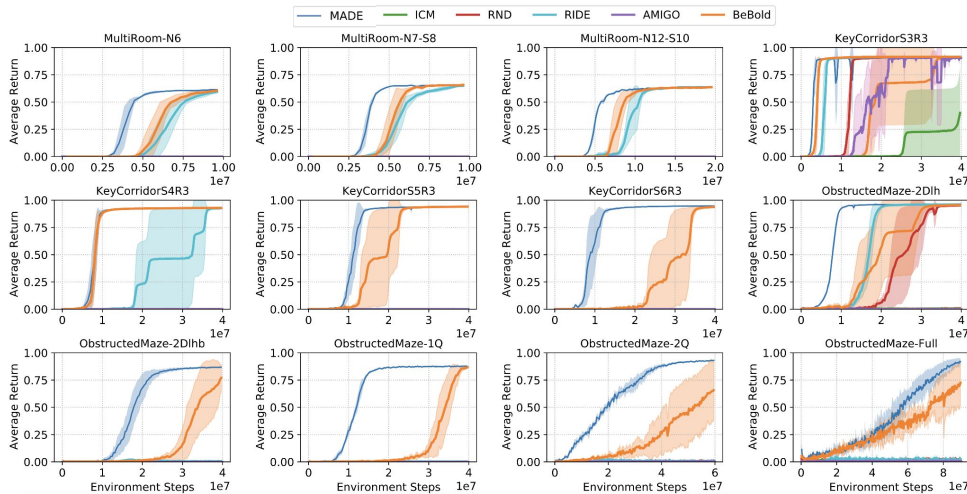


Figure 5: Results for various hard exploration tasks from MiniGrid. MADE successfully solves all the environments while other algorithms (except for BeBold) fail to solve several environments. MADE finds the optimal solution with 2-5 times fewer samples, yielding a much better sample efficiency.

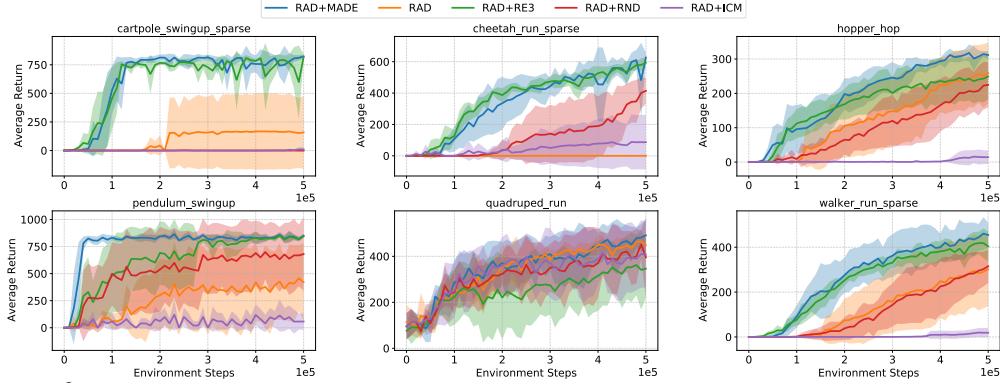


Figure 6: Results for several DeepMind control suite locomotion tasks. Comparing to all baselines, the performance of MADE is consistently better. Sometimes baseline methods even fail to solve the task.

Model-free RL baselines. We consider several baselines in MiniGrid: **IMPALA** [25] is a variant of policy gradient algorithms which we use as the training baseline; **ICM** [76] learns a forward and reverse model for predicting state transition and uses the forward model prediction error as intrinsic reward; **RND** [15] trains a predictor network to mimic a randomly initialized target network as discussed above; **RIDE** [81] learns a representation similar to ICM and uses the difference of learned representations along a trajectory as intrinsic reward; **AMIGO** [17] learns a teacher agent to assign intrinsic reward; **BeBold** [111] adopts a regulated difference of novelty measure using RND. In DeepMind Control Suite, we consider **RE3** [86] as a baseline which uses a random encoder for state embedding followed by a k -nearest neighbour bonus for a maximum state coverage objective.

Model-based RL baselines. MADE can be combined with model-based RL algorithms to improve sample efficiency. For baselines, we consider **Dreamer**, which is a well-known model-based RL algorithm for DeepMind Control Suite, as well as **Dreamer+RE3**, which includes RE3 bonus on top of Dreamer.

MADE achieves state-of-the-art results on both navigation and locomotion tasks by a substantial margin, greatly improving the sample efficiency of the RL exploration in both model-free and model-based methods. Details on experiments and hyperparameters are provided in Appendix B.

5.1 Model-free RL on MiniGrid

MiniGrid [19] is a widely used benchmark for exploration in RL. Despite having symbolic states and a discrete action space, MiniGrid tasks are quite challenging. The easiest task is **MultiRoom** (MR) in which the agent needs to navigate to the goal by going to different rooms connected by the doors. In **KeyCorridor** (KC), the agent needs to search around different rooms to find the key and then use it to open the door. **ObstructedMaze** (OM) is a harder version of KC where the key is hidden in a box and sometimes the door is blocked by an obstruct. In addition to that, the entire environment is procedurally-generated. This adds another layer of difficulty to the problem.

From Figure 5 we can see that MADE manages to solve all the challenging tasks within 90M steps while all other baselines (except BeBold) only solve up to 50% of them. Compared to BeBold, MADE uses significantly (2-5 times) fewer samples.

5.2 Model-free RL on DeepMind Control

We also test MADE on image-based continuous control tasks of DeepMind Control Suite [95], which is a collection of diverse control tasks such as Pendulum, Hopper, and Acrobot with realistic simulations. Compared to MiniGrid, these tasks are more realistic and complex as they involve stochastic transitions, high-dimensional states, and continuous actions. For baselines, we build our algorithm on top of RAD [57], a strong model-free RL algorithm with a competitive sample efficiency. We compare our approach with ICM, RND, as well as RE3, which is the SOTA algorithm³. Note that we compare MADE to very strong baselines. Other algorithms such as DrQ [54], CURL [89], ProtoRL [104], SAC+AE [103]) perform worse based on the results reported in the original papers.

³As we were not provided with the source code, we implemented ICM and RND ourselves. The performance for ICM is slightly worse than what the author reported, but the performance of RND and RE3 is similar.




Figure 7: Ablation study on buffer size in MADE. The optimal buffer size varies in different tasks. We found that a buffer size of 10000 empirically works consistently reasonable.

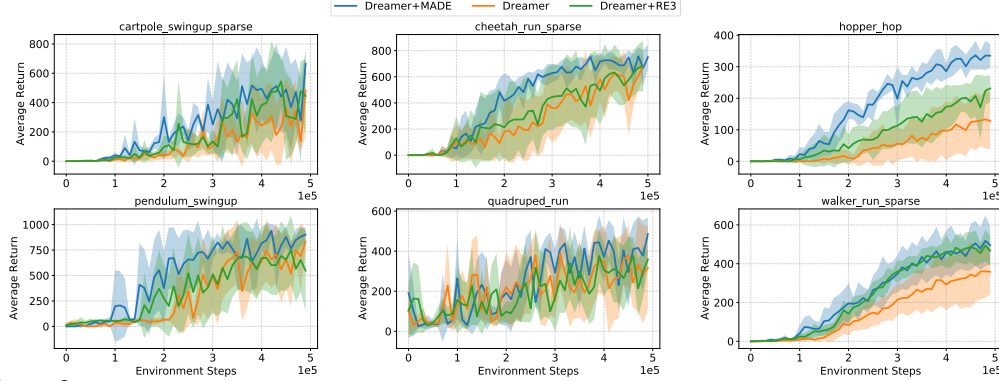


Figure 8: Results for DeepMind control suite locomotion tasks in model-based RL setting. Comparing to all baselines, the performance of MADE is consistently better. Some baseline methods even fail to solve the task.

MADE show consistent improvement in sample efficiency: 2.6 times over RAD+RE3, 3.3 times over RAD+RND, 19.7 times over CURL, 15.0 times over DrQ and 3.8 times over RAD.

From Figure 6, we can see that MADE consistently improves sample efficiency compared to all baselines. For these tasks, RND and ICM do not perform well and even fail on Cartpole-Swingup. RE3 achieves a comparable performance in two tasks, however, its performance on Pendulum-Swingup, Quadruped-Run, Hopper-Hop and Walker-Run is significantly worse than MADE. For example, in Pendulum-Swingup, MADE achieves a reward of around 800 in only 30K steps while RE3 requires 300k samples. In Quadruped-Run, there is a 150 reward gap between MADE and RE3, which seems to be still enlarging. These tasks show the strong performance of MADE in model-free RL.

Ablation study. We study how the buffer length affects the performance of our algorithm in some DeepMind Control tasks. Results show that for different tasks the optimal length is slightly different. We empirically found that using a buffer length of 1000 consistently works well across different tasks.

5.3 Model-based RL on DeepMind Control

We also empirically verify the performance of MADE combined with the SOTA model-based RL algorithm Dreamer [35]. We compare MADE with Dreamer and Dreamer combined with RE3 in Figure 8. Results show that MADE has great sample efficiency in maps like Cheetah-Run-Sparse, Hopper-Hop and Pendulum-Swingup. For example, in Hopper-Hop, MADE achieves more than 100 higher return than RE3 and 250 higher than Dreamer, achieving a new SOTA result.

6 Related work

Provable optimistic exploration. Most provable exploration strategies are based on optimism in the face of uncertainty (OFU) principle. In tabular setting, model-based exploration algorithms include variants of UCB [50, 12], UCRL [56, 41, 106, 49, 64], and Thompson sampling [101, 5, 83] and value-based methods include optimistic Q-learning [44, 98, 90, 60, 65] and value-iteration with UCB [8, 112, 113, 45]. These methods are recently extended to linear MDP setting leading to a variety of model-based [114, 7, 42, 115], value-based [97, 46], and policy-based algorithms [16, 108, 3]. Going beyond linear function approximation, systematic exploration strategies are developed based on structural assumptions on MDP such as low Bellman rank [43] and block MDP [22]. These methods are either computationally intractable [43, 91, 7, 107, 102, 21, 96] or are only oracle efficient

[26, 4]. The recent work [27] provides a sample efficient approach with non-linear policies, however, the algorithm requires maintaining the functional form of all prior policies.

Practical exploration via intrinsic reward. Apart from previously-discussed methods, other works give intrinsic reward based on the difference in (abstraction of) consecutive states [110, 61, 81]. However, this approach is inconsistent: the intrinsic reward does not converge to zero and thus, even with infinite samples, the final policy does not maximize the RL objective. Other intrinsic rewards try to estimate pseudo-counts [10, 94, 15, 14, 75, 9], inspired by provable count-based methods. Though favoring novel states, practically these methods might suffer from *detachment and derailment* [23, 24], and *forgetting* [3]. More recent works propose a combination of different criteria. RIDE [81] learns a representation using curiosity criterion and uses the difference of consecutive states along the trajectory as the bonus. AMIGO [17] learns a teacher agent for assigning rewards for exploration. Go-Explore [23] explicitly decouples the exploration and exploitation stage, yields a more sophisticated algorithm with many hand-tuned hyperparameters. Prior work also tries to add a bonus on maximizing the KL-divergence between the current policy and the previous policies [39]. However, we emphasize that this is indeed an entirely different principle compared with ours, simply because difference of visitation density of the policies is not equivalent to the KL-divergence between policies. In another word, two significantly different policies do not necessarily induce significantly different visitation densities and vice versa.

Maximum entropy exploration. Another line of work encourages exploration via maximizing some type of entropy. One category maximizes policy entropy [67] or relative entropy [2] in addition to the RL objective. The work [28] modifies the RL objective by introducing an adversarial policy which results in the next policy to move away from prior policies while staying close to the current policy. In contrast, our approach focuses on the regions explored by prior policies as opposed to the prior policies themselves. Recently, effects of policy entropy regularization have been studied theoretically [72, 33]. In policy gradient methods with access to exact gradients, policy entropy regularization results in faster convergence by improving the optimization landscape [62, 63, 6, 18]. Another category considers maximizing the entropy of state or state-action visitation densities such as Shannon entropy [36, 40, 58, 86] or Rényi entropy [109]. Empirically, our approach achieves better performance over entropy-based methods.

Other exploration strategies. Besides intrinsic motivation, other strategies are also fruitful in encouraging the RL agent to visit a wide range of states. One example is exploration by injecting noise to the action space [59, 73, 38, 74] or parameter space [30, 78]. Another example is the reward-shaping category, in which diverse goals are set to guide exploration [20, 29, 71, 79].

7 Discussion

We introduce a new exploration strategy MADE based on maximizing deviation from explored regions. We show that by simply adding a regularizer to the original RL objective, we get an easy-to-implement intrinsic reward which can be incorporated with any RL algorithm. We provide a policy computation algorithm for this objective and prove that it converges to a global optimum, provided that we have access to an approximate planner. In tabular setting, MADE consistently improves over the Hoeffding bonus and shows competitive performance to the Bernstein bonus, while the latter is impractical to compute beyond tabular. We conduct extensive experiments on MiniGrid, showing a significant (over 5 times) reduction of the required sample size. MADE also performs well in DeepMind Control Suite when combined with both model-free and model-based RL algorithms, achieving SOTA sample efficiency results. One limitation of the current work is that it only uses the naive representations of states (e.g., one-hot representation in tabular case). In fact, exploration could be conducted much more efficiently if MADE is implemented with a more compact representation of states. We leave this direction to future work.

8 Acknowledgements

The authors are grateful to Andrea Zanette for helpful discussions. The authors thank Alekh Agarwal, Michael Henaff, Sham Kakade, and Wen Sun for providing their code. Paria Rashidinejad is partially supported by the Open Philanthropy Foundation, Intel, and the Leverhulme Trust. Jiantao Jiao is partially supported by NSF grants IIS-1901252, CCF-1909499, and DMS-2023505. Tianjun Zhang is supported by the BAIR Commons at UC-Berkeley and thanks Commons sponsors for their support. In

addition to NSF CISE Expeditions Award CCF-1730628, UC Berkeley research is supported by gifts from Alibaba, Amazon Web Services, Ant Financial, CapitalOne, Ericsson, Facebook, Futurewei, Google, Intel, Microsoft, Nvidia, Scotiabank, Splunk and VMware.

References

- [1] Pieter Abbeel and Andrew Y Ng. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*, page 1, 2004.
- [2] Alekh Agarwal, Sham M Kakade, Jason D Lee, and Gaurav Mahajan. On the theory of policy gradient methods: Optimality, approximation, and distribution shift. *arXiv preprint arXiv:1908.00261*, 2019.
- [3] Alekh Agarwal, Mikael Henaff, Sham Kakade, and Wen Sun. PC-PG: Policy cover directed exploration for provable policy gradient learning. *arXiv preprint arXiv:2007.08459*, 2020.
- [4] Alekh Agarwal, Sham Kakade, Akshay Krishnamurthy, and Wen Sun. Flambe: Structural complexity and representation learning of low rank MDPs. *arXiv preprint arXiv:2006.10814*, 2020.
- [5] Shipra Agrawal and Randy Jia. Optimistic posterior sampling for reinforcement learning: Worst-case regret bounds. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 1184–1194, 2017.
- [6] Zafarali Ahmed, Nicolas Le Roux, Mohammad Norouzi, and Dale Schuurmans. Understanding the impact of entropy on policy optimization. In *International Conference on Machine Learning*, pages 151–160. PMLR, 2019.
- [7] Alex Ayoub, Zeyu Jia, Csaba Szepesvari, Mengdi Wang, and Lin Yang. Model-based reinforcement learning with value-targeted regression. In *International Conference on Machine Learning*, pages 463–474. PMLR, 2020.
- [8] Mohammad Gheshlaghi Azar, Ian Osband, and Rémi Munos. Minimax regret bounds for reinforcement learning. In *International Conference on Machine Learning*, pages 263–272. PMLR, 2017.
- [9] Adrià Puigdomènech Badia, Pablo Sprechmann, Alex Vitvitskyi, Daniel Guo, Bilal Piot, Steven Kapturowski, Olivier Tieleman, Martín Arjovsky, Alexander Pritzel, Andrew Bolt, et al. Never give up: Learning directed exploration strategies. *arXiv preprint arXiv:2002.06038*, 2020.
- [10] Marc Bellemare, Sriram Srinivasan, Georg Ostrovski, Tom Schaul, David Saxton, and Remi Munos. Unifying count-based exploration and intrinsic motivation. In *Advances in neural information processing systems*, pages 1471–1479, 2016.
- [11] Jalaj Bhandari and Daniel Russo. Global optimality guarantees for policy gradient methods. *arXiv preprint arXiv:1906.01786*, 2019.
- [12] Ronen I Brafman and Moshe Tennenholtz. R-max: A general polynomial time algorithm for near-optimal reinforcement learning. *Journal of Machine Learning Research*, 3(Oct):213–231, 2002.
- [13] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- [14] Yuri Burda, Harri Edwards, Deepak Pathak, Amos Storkey, Trevor Darrell, and Alexei A Efros. Large-scale study of curiosity-driven learning. *arXiv preprint arXiv:1808.04355*, 2018.
- [15] Yuri Burda, Harrison Edwards, Amos Storkey, and Oleg Klimov. Exploration by random network distillation. *arXiv preprint arXiv:1810.12894*, 2018.
- [16] Qi Cai, Zhuoran Yang, Chi Jin, and Zhaoran Wang. Provably efficient exploration in policy optimization. In *International Conference on Machine Learning*, pages 1283–1294. PMLR, 2020.

- [17] Andres Campero, Roberta Raileanu, Heinrich Küttler, Joshua B Tenenbaum, Tim Rocktäschel, and Edward Grefenstette. Learning with AMIGO: Adversarially motivated intrinsic goals. *arXiv preprint arXiv:2006.12122*, 2020.
- [18] Shicong Cen, Chen Cheng, Yuxin Chen, Yuting Wei, and Yuejie Chi. Fast global convergence of natural policy gradient methods with entropy regularization. *arXiv preprint arXiv:2007.06558*, 2020.
- [19] Maxime Chevalier-Boisvert, Lucas Willems, and Suman Pal. Minimalistic gridworld environment for OpenAI Gym. <https://github.com/maximecb/gym-minigrid>, 2018.
- [20] Cédric Colas, Pierre Fournier, Mohamed Chetouani, Olivier Sigaud, and Pierre-Yves Oudeyer. CURIOUS: Intrinsically motivated modular multi-goal reinforcement learning. In *International conference on machine learning*, pages 1331–1340. PMLR, 2019.
- [21] Kefan Dong, Jiaqi Yang, and Tengyu Ma. Provable model-based nonlinear bandit and reinforcement learning: Shelve optimism, embrace virtual curvature. *arXiv preprint arXiv:2102.04168*, 2021.
- [22] Simon Du, Akshay Krishnamurthy, Nan Jiang, Alekh Agarwal, Miroslav Dudik, and John Langford. Provably efficient RL with rich observations via latent state decoding. In *International Conference on Machine Learning*, pages 1665–1674. PMLR, 2019.
- [23] Adrien Ecoffet, Joost Huizinga, Joel Lehman, Kenneth O Stanley, and Jeff Clune. Go-explore: A new approach for hard-exploration problems. *arXiv preprint arXiv:1901.10995*, 2019.
- [24] Adrien Ecoffet, Joost Huizinga, Joel Lehman, Kenneth O Stanley, and Jeff Clune. First return then explore. *arXiv preprint arXiv:2004.12919*, 2020.
- [25] Lasse Espeholt, Hubert Soyer, Remi Munos, Karen Simonyan, Volodymyr Mnih, Tom Ward, Yotam Doron, Vlad Firoiu, Tim Harley, Iain Dunning, et al. Impala: Scalable distributed deep-RL with importance weighted actor-learner architectures. *arXiv preprint arXiv:1802.01561*, 2018.
- [26] Fei Feng, Ruosong Wang, Wotao Yin, Simon S Du, and Lin Yang. Provably efficient exploration for reinforcement learning using unsupervised learning. *Advances in Neural Information Processing Systems*, 33, 2020.
- [27] Fei Feng, Wotao Yin, Alekh Agarwal, and Lin F Yang. Provably correct optimization and exploration with non-linear policies. *arXiv preprint arXiv:2103.11559*, 2021.
- [28] Yannis Flet-Berliac, Johan Ferret, Olivier Pietquin, Philippe Preux, and Matthieu Geist. Adversarially guided actor-critic. In *International Conference on Learning Representations*, 2021.
- [29] Carlos Florensa, David Held, Xinyang Geng, and Pieter Abbeel. Automatic goal generation for reinforcement learning agents. In *International conference on machine learning*, pages 1515–1528. PMLR, 2018.
- [30] Meire Fortunato, Mohammad Gheshlaghi Azar, Bilal Piot, Jacob Menick, Ian Osband, Alex Graves, Vlad Mnih, Remi Munos, Demis Hassabis, Olivier Pietquin, et al. Noisy networks for exploration. *International Conference on Learning Representations*, 2018.
- [31] Marguerite Frank, Philip Wolfe, et al. An algorithm for quadratic programming. *Naval research logistics quarterly*, 3(1-2):95–110, 1956.
- [32] Matthieu Geist, Bilal Piot, and Olivier Pietquin. Is the Bellman residual a bad proxy? In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 3208–3217, 2017.
- [33] Matthieu Geist, Bruno Scherrer, and Olivier Pietquin. A theory of regularized Markov decision processes. In *International Conference on Machine Learning*, pages 2160–2169. PMLR, 2019.
- [34] Karol Gregor, Danilo Jimenez Rezende, and Daan Wierstra. Variational intrinsic control. *arXiv preprint arXiv:1611.07507*, 2016.

- [35] Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. Dream to control: Learning behaviors by latent imagination. *arXiv preprint arXiv:1912.01603*, 2019.
- [36] Elad Hazan, Sham Kakade, Karan Singh, and Abby Van Soest. Provably efficient maximum entropy exploration. In *International Conference on Machine Learning*, pages 2681–2691, 2019.
- [37] Jiafan He, Dongruo Zhou, and Quanquan Gu. Nearly minimax optimal reinforcement learning for discounted MDPs. *arXiv preprint arXiv:2010.00587*, 2020.
- [38] Matteo Hessel, Joseph Modayil, Hado Van Hasselt, Tom Schaul, Georg Ostrovski, Will Dabney, Dan Horgan, Bilal Piot, Mohammad Azar, and David Silver. Rainbow: Combining improvements in deep reinforcement learning. *arXiv preprint arXiv:1710.02298*, 2017.
- [39] Zhang-Wei Hong, Tzu-Yun Shann, Shih-Yang Su, Yi-Hsiang Chang, and Chun-Yi Lee. Diversity-driven exploration strategy for deep reinforcement learning. *arXiv preprint arXiv:1802.04564*, 2018.
- [40] Riashat Islam, Zafarali Ahmed, and Doina Precup. Marginalized state distribution entropy regularization in policy optimization. *arXiv preprint arXiv:1912.05128*, 2019.
- [41] Thomas Jaksch, Ronald Ortner, and Peter Auer. Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research*, 11(4), 2010.
- [42] Zeyu Jia, Lin Yang, Csaba Szepesvari, and Mengdi Wang. Model-based reinforcement learning with value-targeted regression. In *Learning for Dynamics and Control*, pages 666–686. PMLR, 2020.
- [43] Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, John Langford, and Robert E Schapire. Contextual decision processes with low Bellman rank are PAC-learnable. In *International Conference on Machine Learning*, pages 1704–1713. PMLR, 2017.
- [44] Chi Jin, Zeyuan Allen-Zhu, Sebastien Bubeck, and Michael I Jordan. Is Q-learning provably efficient? In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 4868–4878, 2018.
- [45] Chi Jin, Akshay Krishnamurthy, Max Simchowitz, and Tiancheng Yu. Reward-free exploration for reinforcement learning. In *International Conference on Machine Learning*, pages 4870–4879. PMLR, 2020.
- [46] Chi Jin, Zhuoran Yang, Zhaoran Wang, and Michael I Jordan. Provably efficient reinforcement learning with linear function approximation. In *Conference on Learning Theory*, pages 2137–2143. PMLR, 2020.
- [47] Sham Kakade and John Langford. Approximately optimal approximate reinforcement learning. In *In Proc. 19th International Conference on Machine Learning*. Citeseer, 2002.
- [48] Sham Kakade, Mengdi Wang, and Lin F Yang. Variance reduction methods for sublinear reinforcement learning. *arXiv preprint arXiv:1802.09184*, 2018.
- [49] Emilie Kaufmann, Pierre Ménard, Omar Darwiche Domingues, Anders Jonsson, Edouard Leurent, and Michal Valko. Adaptive reward-free exploration. In *Algorithmic Learning Theory*, pages 865–891. PMLR, 2021.
- [50] Michael Kearns and Satinder Singh. Near-optimal reinforcement learning in polynomial time. *Machine learning*, 49(2):209–232, 2002.
- [51] Hyoungeok Kim, Jaekyeom Kim, Yeonwoo Jeong, Sergey Levine, and Hyun Oh Song. Emi: Exploration with mutual information. *arXiv preprint arXiv:1810.01176*, 2018.
- [52] Alexander S Klyubin, Daniel Polani, and Chrystopher L Nehaniv. All else being equal be empowered. In *European Conference on Artificial Life*, pages 744–753. Springer, 2005.
- [53] Vijay R Konda and John N Tsitsiklis. Actor-critic algorithms. In *Advances in neural information processing systems*, pages 1008–1014. Citeseer, 2000.

- [54] Ilya Kostrikov, Denis Yarats, and Rob Fergus. Image augmentation is all you need: Regularizing deep reinforcement learning from pixels. *arXiv preprint arXiv:2004.13649*, 2020.
- [55] Michael Laskin, Kimin Lee, Adam Stooke, Lerrel Pinto, Pieter Abbeel, and Aravind Srinivas. Reinforcement learning with augmented data. *arXiv preprint arXiv:2004.14990*, 2020.
- [56] Tor Lattimore and Marcus Hutter. PAC bounds for discounted MDPs. In *International Conference on Algorithmic Learning Theory*, pages 320–334. Springer, 2012.
- [57] Kimin Lee, Kibok Lee, Jinwoo Shin, and Honglak Lee. Network randomization: A simple technique for generalization in deep reinforcement learning. *arXiv preprint arXiv:1910.05396*, 2019.
- [58] Lisa Lee, Benjamin Eysenbach, Emilio Parisotto, Eric Xing, Sergey Levine, and Ruslan Salakhutdinov. Efficient exploration via state marginal matching. *arXiv preprint arXiv:1906.05274*, 2019.
- [59] Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- [60] Shuang Liu and Hao Su. Regret bounds for discounted MDPs. *arXiv preprint arXiv:2002.05138*, 2020.
- [61] Kenneth Marino, Abhinav Gupta, Rob Fergus, and Arthur Szlam. Hierarchical RL using an ensemble of proprioceptive periodic policies. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=SJz1x20cFQ>.
- [62] Jincheng Mei, Chenjun Xiao, Csaba Szepesvari, and Dale Schuurmans. On the global convergence rates of softmax policy gradient methods. In *International Conference on Machine Learning*, pages 6820–6829. PMLR, 2020.
- [63] Jincheng Mei, Yue Gao, Bo Dai, Csaba Szepesvari, and Dale Schuurmans. Leveraging non-uniformity in first-order non-convex optimization. *arXiv preprint arXiv:2105.06072*, 2021.
- [64] Pierre Ménard, Omar Darwiche Domingues, Anders Jonsson, Emilie Kaufmann, Edouard Leurent, and Michal Valko. Fast active learning for pure exploration in reinforcement learning. *arXiv preprint arXiv:2007.13442*, 2020.
- [65] Pierre Menard, Omar Darwiche Domingues, Xuedong Shang, and Michal Valko. UCB momentum Q-learning: Correcting the bias without forgetting. *arXiv preprint arXiv:2103.01312*, 2021.
- [66] Dipendra Misra, Mikael Henaff, Akshay Krishnamurthy, and John Langford. Kinematic state abstraction and provably efficient rich-observation reinforcement learning. In *International conference on machine learning*, pages 6961–6971. PMLR, 2020.
- [67] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pages 1928–1937, 2016.
- [68] Shakir Mohamed and Danilo Jimenez Rezende. Variational information maximisation for intrinsically motivated reinforcement learning. In *Advances in neural information processing systems*, pages 2125–2133, 2015.
- [69] Matej Moravčík, Martin Schmid, Neil Burch, Viliam Lisý, Dustin Morrill, Nolan Bard, Trevor Davis, Kevin Waugh, Michael Johanson, and Michael Bowling. Deepstack: Expert-level artificial intelligence in heads-up no-limit poker. *Science*, 356(6337):508–513, 2017.
- [70] Ofir Nachum, Mohammad Norouzi, Kelvin Xu, and Dale Schuurmans. Bridging the gap between value and policy based reinforcement learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 2772–2782, 2017.

- [71] Ashvin V Nair, Vitchyr Pong, Murtaza Dalal, Shikhar Bahl, Steven Lin, and Sergey Levine. Visual reinforcement learning with imagined goals. *Advances in Neural Information Processing Systems*, 31:9191–9200, 2018.
- [72] Gergely Neu, Anders Jonsson, and Vicens Gómez. A unified view of entropy-regularized Markov decision processes. *arXiv preprint arXiv:1705.07798*, 2017.
- [73] Ian Osband, Charles Blundell, Alexander Pritzel, and Benjamin Van Roy. Deep exploration via bootstrapped DQN. In *Advances in neural information processing systems*, pages 4026–4034, 2016.
- [74] Ian Osband, Benjamin Van Roy, Daniel J Russo, and Zheng Wen. Deep exploration via randomized value functions. *Journal of Machine Learning Research*, 20(124):1–62, 2019.
- [75] Georg Ostrovski, Marc G Bellemare, Aaron van den Oord, and Rémi Munos. Count-based exploration with neural density models. *arXiv preprint arXiv:1703.01310*, 2017.
- [76] Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. Curiosity-driven exploration by self-supervised prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 16–17, 2017.
- [77] Deepak Pathak, Dhiraj Gandhi, and Abhinav Gupta. Self-supervised exploration via disagreement. *arXiv preprint arXiv:1906.04161*, 2019.
- [78] Matthias Plappert, Rein Houthoofd, Prafulla Dhariwal, Szymon Sidor, Richard Y Chen, Xi Chen, Tamim Asfour, Pieter Abbeel, and Marcin Andrychowicz. Parameter space noise for exploration. *International Conference on Learning Representations*, 2018.
- [79] Vitchyr Pong, Murtaza Dalal, Steven Lin, Ashvin Nair, Shikhar Bahl, and Sergey Levine. Skew-Fit: State-covering self-supervised reinforcement learning. In *International Conference on Machine Learning*, pages 7783–7792. PMLR, 2020.
- [80] Martin L Puterman. Markov decision processes. *Handbooks in operations research and management science*, 2:331–434, 1990.
- [81] Roberta Raileanu and Tim Rocktäschel. Ride: Rewarding impact-driven exploration for procedurally-generated environments. *arXiv preprint arXiv:2002.12292*, 2020.
- [82] Paria Rashidinejad, Banghua Zhu, Cong Ma, Jiantao Jiao, and Stuart Russell. Bridging offline reinforcement learning and imitation learning: A tale of pessimism. *arXiv preprint arXiv:2103.12021*, 2021.
- [83] Daniel Russo. Worst-case regret bounds for exploration via randomized value functions. *arXiv preprint arXiv:1906.02870*, 2019.
- [84] Christoph Salge, Cornelius Glackin, and Daniel Polani. Empowerment: An introduction. In *Guided Self-Organization: Inception*, pages 67–114. Springer, 2014.
- [85] Bruno Scherrer. Approximate policy iteration schemes: A comparison. In *International Conference on Machine Learning*, pages 1314–1322. PMLR, 2014.
- [86] Younggyo Seo, Lili Chen, Jinwoo Shin, Honglak Lee, Pieter Abbeel, and Kimin Lee. State entropy maximization with random encoders for efficient exploration. *arXiv preprint arXiv:2102.09430*, 2021.
- [87] Pranav Shyam, Wojciech Jaśkowski, and Faustino Gomez. Model-based active exploration. In *International Conference on Machine Learning*, pages 5779–5788, 2019.
- [88] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of Go without human knowledge. *nature*, 550(7676):354–359, 2017.
- [89] Aravind Srinivas, Michael Laskin, and Pieter Abbeel. Curl: Contrastive unsupervised representations for reinforcement learning. *arXiv preprint arXiv:2004.04136*, 2020.

- [90] Alexander L Strehl, Lihong Li, Eric Wiewiora, John Langford, and Michael L Littman. PAC model-free reinforcement learning. In *Proceedings of the 23rd international conference on Machine learning*, pages 881–888, 2006.
- [91] Wen Sun, Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, and John Langford. Model-based RL in contextual decision processes: PAC bounds and exponential improvements over model-free approaches. In *Conference on Learning Theory*, pages 2898–2933. PMLR, 2019.
- [92] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- [93] Richard S Sutton, David McAllester, Satinder Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. In *Proceedings of the 12th International Conference on Neural Information Processing Systems*, pages 1057–1063, 1999.
- [94] Haoran Tang, Rein Houthooft, Davis Foote, Adam Stooke, OpenAI Xi Chen, Yan Duan, John Schulman, Filip DeTurck, and Pieter Abbeel. #Exploration: A study of count-based exploration for deep reinforcement learning. In *Advances in neural information processing systems*, pages 2753–2762, 2017.
- [95] Yuval Tassa, Saran Tunyasuvunakool, Alistair Muldal, Yotam Doron, Piotr Trochim, Siqi Liu, Steven Bohez, Josh Merel, Tom Erez, Timothy Lillicrap, et al. dm_control: Software and tasks for continuous control. *arXiv preprint arXiv:2006.12983*, 2020.
- [96] Ruosong Wang, Russ R Salakhutdinov, and Lin Yang. Reinforcement learning with general value function approximation: Provably efficient approach via bounded Eluder dimension. *Advances in Neural Information Processing Systems*, 33, 2020.
- [97] Yining Wang, Ruosong Wang, Simon S Du, and Akshay Krishnamurthy. Optimism in reinforcement learning with generalized linear function approximation. *arXiv preprint arXiv:1912.04136*, 2019.
- [98] Yuanhao Wang, Kefan Dong, Xiaoyu Chen, and Liwei Wang. Q-learning with UCB exploration is sample efficient for infinite-horizon MDP. In *International Conference on Learning Representations*, 2019.
- [99] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.
- [100] Ronald J Williams and Jing Peng. Function optimization using connectionist reinforcement learning algorithms. *Connection Science*, 3(3):241–268, 1991.
- [101] Zhihan Xiong, Ruoqi Shen, and Simon S Du. Randomized exploration is near-optimal for tabular MDP. *arXiv preprint arXiv:2102.09703*, 2021.
- [102] Zhuoran Yang, Chi Jin, Zhaoran Wang, Mengdi Wang, and Michael I Jordan. Bridging exploration and general function approximation in reinforcement learning: Provably efficient kernel and neural value iterations. *arXiv preprint arXiv:2011.04622*, 2020.
- [103] Denis Yarats, Amy Zhang, Ilya Kostrikov, Brandon Amos, Joelle Pineau, and Rob Fergus. Improving sample efficiency in model-free reinforcement learning from images. *arXiv preprint arXiv:1910.01741*, 2019.
- [104] Denis Yarats, Rob Fergus, Alessandro Lazaric, and Lerrel Pinto. Reinforcement learning with prototypical representations. *arXiv preprint arXiv:2102.11271*, 2021.
- [105] Tom Zahavy, Alon Cohen, Haim Kaplan, and Yishay Mansour. Apprenticeship learning via frank-wolfe. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 6720–6728, 2020.
- [106] Andrea Zanette and Emma Brunskill. Tighter problem-dependent regret bounds in reinforcement learning without domain knowledge using value function bounds. In *International Conference on Machine Learning*, pages 7304–7312. PMLR, 2019.

- [107] Andrea Zanette, Alessandro Lazaric, Mykel Kochenderfer, and Emma Brunskill. Learning near optimal policies with low inherent Bellman error. In *International Conference on Machine Learning*, pages 10978–10989. PMLR, 2020.
- [108] Andrea Zanette, Ching-An Cheng, and Alekh Agarwal. Cautiously optimistic policy optimization and exploration with linear function approximation. *arXiv preprint arXiv:2103.12923*, 2021.
- [109] Chuheng Zhang, Yuanying Cai, and Longbo Huang Jian Li. Exploration by maximizing Rényi entropy for reward-free RL framework. *Proceedings of the AAAI Conference on Artificial Intelligence*, 2021.
- [110] Jingwei Zhang, Niklas Wetzels, Nicolai Dorka, Joschka Boedecker, and Wolfram Burgard. Scheduled intrinsic drive: A hierarchical take on intrinsically motivated exploration. *arXiv preprint arXiv:1903.07400*, 2019.
- [111] Tianjun Zhang, Huazhe Xu, Xiaolong Wang, Yi Wu, Kurt Keutzer, Joseph E Gonzalez, and Yuandong Tian. BeBold: Exploration beyond the boundary of explored regions. *arXiv preprint arXiv:2012.08621*, 2020.
- [112] Zihan Zhang, Xiangyang Ji, and Simon S Du. Is reinforcement learning more difficult than bandits? a near-optimal algorithm escaping the curse of horizon. *arXiv preprint arXiv:2009.13503*, 2020.
- [113] Zihan Zhang, Yuan Zhou, and Xiangyang Ji. Almost optimal model-free reinforcement learning via reference-advantage decomposition. *Advances in Neural Information Processing Systems*, 33, 2020.
- [114] Dongruo Zhou, Quanquan Gu, and Csaba Szepesvari. Nearly minimax optimal reinforcement learning for linear mixture Markov decision processes. *arXiv preprint arXiv:2012.08507*, 2020.
- [115] Dongruo Zhou, Jiafan He, and Quanquan Gu. Provably efficient reinforcement learning for discounted MDPs with feature mapping. *arXiv preprint arXiv:2006.13165*, 2020.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? **[Yes]** We claim to propose efficient exploration criterion that achieves strong results. The experiments are conducted in MiniGrid, and DeepMind Control.
 - (b) Did you describe the limitations of your work? **[Yes]** We discuss limitations in the Discussion section.
 - (c) Did you discuss any potential negative societal impacts of your work? **[No]** We do not foresee any obvious negative societal impacts from our work, which contributes a general-purpose exploration algorithm.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? **[Yes]**
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? **[Yes]** We state the full set of assumptions
 - (b) Did you include complete proofs of all theoretical results? **[Yes]** The complete proof is in Appendix.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **[Yes]** We upload the code in supplemental material.

- (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [\[Yes\]](#) We include the training detail in supplemental material.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [\[Yes\]](#) We include the error bars for experiments.
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [\[No\]](#)
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- (a) If your work uses existing assets, did you cite the creators? [\[Yes\]](#)
 - (b) Did you mention the license of the assets? [\[No\]](#)
 - (c) Did you include any new assets either in the supplemental material or as a URL? [\[N/A\]](#)
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [\[No\]](#) We use publicly available datasets/tasks.
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [\[No\]](#) We don't use data of this sort.
5. If you used crowdsourcing or conducted research with human subjects...
- (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [\[N/A\]](#)
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [\[N/A\]](#)
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [\[N/A\]](#)