

An Intrinsically-Motivated Approach for Learning Highly Exploring and Fast Mixing Policies

Mirco Mutti

Politecnico di Milano, Milan, Italy
Università di Bologna, Bologna, Italy
mirco.mutti@polimi.it

Marcello Restelli

Politecnico di Milano, Milan, Italy
marcello.restelli@polimi.it

Abstract

What is a good exploration strategy for an agent that interacts with an environment in the absence of external rewards? Ideally, we would like to get a policy driving towards a uniform state-action visitation (highly exploring) in a minimum number of steps (fast mixing), in order to ease efficient learning of any goal-conditioned policy later on. Unfortunately, it is remarkably arduous to directly learn an optimal policy of this nature. In this paper, we propose a novel surrogate objective for learning highly exploring and fast mixing policies, which focuses on maximizing a lower bound to the entropy of the steady-state distribution induced by the policy. In particular, we introduce three novel lower bounds, that lead to as many optimization problems, that tradeoff the theoretical guarantees with computational complexity. Then, we present a model-based reinforcement learning algorithm, IDE³AL, to learn an optimal policy according to the introduced objective. Finally, we provide an empirical evaluation of this algorithm on a set of hard-exploration tasks.

1 Introduction

In general, the Reinforcement Learning (RL) framework (Sutton and Barto 2018) assumes the presence of a reward signal coming from a, potentially unknown, environment to a learning agent. When this signal is sufficiently informative about the utility of the agent’s decisions, RL has proved to be rather successful in solving challenging tasks, even at a super-human level (Mnih et al. 2015; Silver et al. 2017). However, in most real-world scenarios, we cannot rely on a well-shaped, complete reward signal. This may prevent the agent from learning anything until, while performing random actions, it eventually stumbles into some sort of external reward. Thus, what is a good objective for a learning agent to pursue, in the absence of an external reward signal, to prepare itself to learn efficiently, eventually, a goal-conditioned policy?

Intrinsic motivation (Chentanez, Barto, and Singh 2005; Oudeyer and Kaplan 2009) traditionally tries to answer this pressing question by designing self-motivated goals that favor exploration. In a curiosity-driven approach, first proposed in (Schmidhuber 1991), the intrinsic objective encourages the agent to explore novel states by rewarding prediction

errors (Stadie, Levine, and Abbeel 2015; Pathak et al. 2017; Burda et al. 2019a; 2019b). On a similar flavor, other works propose to relate an intrinsic reward to some sort of learning progress (Lopes et al. 2012) or information gain (Mohamed and Rezende 2015; Houthoof et al. 2016), stimulating the agent’s empowerment over the environment. Count-based approaches (Bellemare et al. 2016; Tang et al. 2017; Ostrovski et al. 2017) consider exploration bonuses proportional to the state visitation frequencies, assigning high rewards to rarely visited states. Although the mentioned approaches have been relatively effective in solving sparse-rewards, hard-exploration tasks (Pathak et al. 2017; Burda et al. 2019b), they have some common limitations that may affect their ability to methodically explore an environment in the absence of external rewards, as pointed out in (Ecoffet et al. 2019). Especially, due to the consumable nature of their intrinsic bonuses, the learning agent could prematurely lose interest in a frontier of high rewards (*detachment*). Furthermore, the agent may suffer from *derailment* by trying to return to a promising state, previously discovered, if a naïve exploratory mechanism, such as ϵ -greedy, is combined to the intrinsic motivation mechanism (which is often the case). To overcome these limitations, recent works suggest alternative approaches to motivate the agent towards a more systematic exploration of the environment (Hazan et al. 2019; Ecoffet et al. 2019). Especially, in (Hazan et al. 2019) the authors consider an intrinsic objective which is directed to the maximization of an entropic measure over the state distribution induced by a policy. Then, they provide a provably efficient algorithm to learn a mixture of deterministic policies that is overall optimal w.r.t. the maximum-entropy exploration objective. To the best of our knowledge, none of the mentioned approaches explicitly address the related aspect of the mixing time of an exploratory policy, which represents the time it takes for the policy to reach its full capacity in terms of exploration. Nonetheless, in many cases we would like to maximize the probability of reaching any potential target state having a fairly limited number of interactions at hand for exploring the environment. Notably, this context presents some analogies to the problem of maximizing the efficiency of a random walk (Hassibi et al. 2014).

In this paper, we present a novel approach to learn exploratory policies that are, at the same time, highly exploring and fast mixing. In Section 3, we propose a surrogate objec-

tive to address the problem of maximum-entropy exploration over both the state space (Section 3.1) and the action space (Section 3.2). The idea is to search for a policy that maximizes a lower bound to the entropy of the induced steady-state distribution. We introduce three new lower bounds and the corresponding optimization problems, discussing their pros and cons. Furthermore, we discuss how to complement the introduced objective to account for the mixing time of the learned policy (Section 3.3). In Section 4, we present the *Intrinsically-Driven Effective and Efficient Exploration ALgorithm* (IDE³AL), a novel, model-based, reinforcement learning method to learn highly exploring and fast mixing policies through iterative optimizations of the introduced objective. In Section 5, we provide an empirical evaluation to illustrate the merits of our approach on hard-exploration, finite domains, and to show how it fares in comparison to count-based and maximum-entropy approaches. Finally, in Section 6, we discuss the proposed approach and related works. The proofs of the Theorems are reported in Appendix A¹.

2 Preliminaries

A discrete-time Markov Decision Process (MDP) (Puterman 2014) is defined as a tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, P, R, d_0)$, where \mathcal{S} is the state space, \mathcal{A} is the action space, $P(s'|s, a)$ is a Markovian transition model defining the distribution of the next state s' given the current state s and action a , R is the reward function, such that $R(s, a)$ is the expected immediate reward when taking action a from state s , and d_0 is the initial state distribution. A policy $\pi(a|s)$ defines the probability of taking an action a in state s .

In the following we will indifferently turn to scalar or matrix notation, where \mathbf{v} denotes a vector, \mathbf{M} denotes a matrix, and \mathbf{v}^T , \mathbf{M}^T denote their transpose. A matrix is row (column) stochastic if it has non-negative entries and all of its rows (columns) sum to one. A matrix is *doubly stochastic* if it is both row and column stochastic. We denote with \mathbb{P} the space of doubly stochastic matrices. The L_∞ -norm $\|\mathbf{M}\|_\infty$ of a matrix is its maximum absolute row sum, while $\|\mathbf{M}\|_2 = (\max \text{eig } \mathbf{M}^T \mathbf{M})^{\frac{1}{2}}$ and $\|\mathbf{M}\|_F = (\sum_i \sum_j (\mathbf{M}(i, j))^2)^{\frac{1}{2}}$ are its L_2 and Frobenius norms respectively. We denote with $\mathbf{1}_n$ a column vector of n ones and with $\mathbf{1}_{n \times m}$ a matrix of ones with n rows and m columns. Using matrix notation, \mathbf{d}_0 is a column vector of size $|\mathcal{S}|$ having elements $d_0(s)$, \mathbf{P} is a row stochastic matrix of size $(|\mathcal{S}| \times |\mathcal{A}| \times |\mathcal{S}|)$ that describes the transition model $\mathbf{P}((s, a), s') = P(s'|s, a)$, $\mathbf{\Pi}$ is a row stochastic matrix of size $(|\mathcal{S}| \times |\mathcal{S}| \times |\mathcal{A}|)$ that contains the policy $\mathbf{\Pi}(s, (s, a)) = \pi(a|s)$, and $\mathbf{P}^\pi = \mathbf{\Pi} \mathbf{P}$ is a row stochastic matrix of size $(|\mathcal{S}| \times |\mathcal{S}|)$ that represents the state transition matrix under policy π . We denote with \mathbb{I} the space of all the stationary Markovian policies.

In the absence of any reward, i.e., when $R(s, a) = 0$ for every (s, a) , a policy π induces, over the MDP \mathcal{M} , a Markov Chain (MC) (Levin and Peres 2017) defined by $\mathcal{C} = (\mathcal{S}, P^\pi, d_0)$ where $P^\pi(s'|s) = \mathbf{P}^\pi(s, s')$ is the state transition model. Having defined the t -step transition matrix

as $\mathbf{P}_t^\pi = (\mathbf{P}^\pi)^t$, the state distribution of the MC at time step t is $\mathbf{d}_t^\pi = (\mathbf{P}_t^\pi)^T \mathbf{d}_0$, while $\mathbf{d}^\pi = \lim_{t \rightarrow \infty} \mathbf{d}_t^\pi$ is the steady state distribution. If the MC is ergodic, i.e., aperiodic and recurrent, it admits a unique steady-state distribution, such that $\mathbf{d}^\pi = (\mathbf{P}^\pi)^T \mathbf{d}^\pi$. The mixing time t_{mix} of the MC describes how fast the state distribution converges to the steady state:

$$t_{\text{mix}} = \min \{t \in \mathbb{N} : \sup_{d_0} \|\mathbf{d}_t^\pi - \mathbf{d}^\pi\|_\infty \leq \epsilon\}, \quad (1)$$

where ϵ is the mixing threshold. An MC is reversible if the condition $\mathbf{P}^\pi \mathbf{d}^\pi = (\mathbf{P}^\pi)^T \mathbf{d}^\pi$ holds. Let λ_π be the eigenvalues of \mathbf{P}^π . For ergodic reversible MCs the largest eigenvalue is 1 with multiplicity 1. Then, we can define the second largest eigenvalue modulus $\lambda_\pi(2)$ and the spectral gap γ_π as:

$$\lambda_\pi(2) = \max_{\lambda_\pi(i) \neq 1} |\lambda_\pi(i)|, \quad \gamma_\pi = 1 - \lambda_\pi(2). \quad (2)$$

3 Optimization Problems for Highly Exploring and Fast Mixing Policies

In this section, we define a set of optimization problems whose goal is to identify a stationary Markovian policy that effectively explores the state-action space. The optimization problem is introduced in three steps: first we ask for a policy that maximizes some lower bound to the steady-state distribution entropy, then we foster exploration over the action space by adding a constraint on the minimum action probability, and finally we add another constraint to reduce the mixing time of the Markov chain induced by the policy.

3.1 Highly Exploring Policies over the State Space

Intuitively, a good exploration policy should guarantee to visit the state space as uniformly as possible. In this view, a potential objective function is the entropy of the steady-state distribution induced by a policy over the MDP (Hazan et al. 2019). The resulting optimal policy is:

$$\pi^* \in \arg \max_{\pi \in \mathbb{I}} H(\mathbf{d}^\pi), \quad (3)$$

where $H(\mathbf{d}^\pi) = -\mathbb{E}_{s \sim \mathbf{d}^\pi} [\log \mathbf{d}^\pi(s)]$ is the state distribution entropy. Unfortunately, a direct optimization of this objective is particularly arduous since the steady-state distribution entropy is not a concave function of the policy (Hazan et al. 2019). To overcome this issue, a possible solution (Hazan et al. 2019) is to use the conditional gradient method, such that the gradients of the steady-state distribution entropy become the intrinsic reward in a sequence of approximate dynamic programming problems (Bertsekas 1995).

In this paper, we follow an alternative route that consists in maximizing a lower bound to the policy entropy. In particular, in the following we will consider three lower bounds that lead to as many optimization problems (named Infinity, Frobenius, Column Sum) that show different trade-offs between theoretical guarantees and computational complexity.

Infinity From the theory of Markov chains (Levin and Peres 2017), we know a necessary and sufficient condition for a policy to induce a uniform steady-state distribution (i.e., to achieve the maximum possible entropy). We report this result in the following theorem.

¹A complete version of the paper, which includes the Appendix, is available at <https://arxiv.org/abs/1907.04662>

Theorem 3.1. *Let P be the transition matrix of a given MDP. The steady-state distribution d^π induced by a policy π is uniform over \mathcal{S} iff the matrix $P^\pi = \Pi P$ is doubly stochastic.*

Unfortunately, given the constraints specified by the transition matrix P , a stationary Markovian policy that induces a doubly stochastic P^π may not exist. On the other hand, it is possible to lower bound the entropy of the steady-state distribution induced by policy π as a function of the minimum L_∞ -norm between P^π and any doubly stochastic matrix.

Theorem 3.2. *Let P be the transition matrix of a given MDP and \mathbb{P} the space of doubly stochastic matrices. The entropy of the steady-state distribution d^π induced by a policy π is lower bounded by:*

$$H(d^\pi) \geq \log |\mathcal{S}| - |\mathcal{S}| \inf_{P^u \in \mathbb{P}} \|P^u - \Pi P\|_\infty^2.$$

The maximization of this lower bound leads to the following constrained optimization problem:

$$\underset{P^u \in \mathbb{P}, \Pi \in \Pi}{\text{minimize}} \quad \|P^u - \Pi P\|_\infty \quad (4)$$

It is worth noting that this optimization problem can be reformulated as a linear program with $|\mathcal{S}|^2 + |\mathcal{S}||\mathcal{A}| + |\mathcal{S}|$ optimization variables and $2^{|\mathcal{S}|}|\mathcal{S}| + |\mathcal{S}|^2 + |\mathcal{S}||\mathcal{A}|$ inequality constraints and $3|\mathcal{S}|$ equality constraints (the linear program formulation can be found in Appendix B.1. In order to avoid the exponential growth of the number of constraints as a function of the number of states, we are going to introduce alternative optimization problems.

Frobenius It is worth noting that different transition matrices P^π having equal $\|P^u - P^\pi\|_\infty$ might lead to significantly different state distribution entropies $H(d^\pi)$, as the L_∞ -norm only accounts for the state corresponding to the maximum absolute row sum. The Frobenius norm can better capture the distance between P^u and P^π over all the states, as discussed in Appendix C. For this reason, we have derived a lower bound to the policy entropy that replace the L_∞ -norm with the Frobenius one.

Theorem 3.3. *Let P be the transition matrix of a given MDP and \mathbb{P} the space of doubly stochastic matrices. The entropy of the steady-state distribution d^π induced by a policy π is lower bounded by:*

$$H(d^\pi) \geq \log |\mathcal{S}| - |\mathcal{S}|^2 \inf_{P^u \in \mathbb{P}} \|P^u - \Pi P\|_F^2.$$

It can be shown (see Corollary A.1 in Appendix A that the lower bound based on the Frobenius norm cannot be better (i.e., larger) than the one with the Infinite norm. However, we have the advantage that the resulting optimization problem has significantly less constraints than Problem (4):

$$\underset{P^u \in \mathbb{P}, \Pi \in \Pi}{\text{minimize}} \quad \|P^u - \Pi P\|_F. \quad (5)$$

This problem is a (linearly constrained) quadratic problem with $|\mathcal{S}|^2 + |\mathcal{S}||\mathcal{A}|$ optimization variables and $|\mathcal{S}|^2 + |\mathcal{S}||\mathcal{A}|$ inequality constraints and $3|\mathcal{S}|$ equality constraints.

Column Sum Problems (4) and (5) are aiming at finding a policy associated with a state transition matrix that is

doubly stochastic. To achieve this result it is enough to guarantee that the column sums of the matrix P^π are all equal to one (Kirkland 2010). A measure that can be used to evaluate the distance to a doubly stochastic matrix can be the absolute sum of the difference between one and the column sums: $\sum_{s \in \mathcal{S}} |1 - \sum_{s' \in \mathcal{S}} P^\pi(s|s')| = \left\| (I - (\Pi P)^T) \cdot \mathbf{1}_{|\mathcal{S}|} \right\|_1$. The following theorem provides a lower bound to the policy entropy as a function of this measure.

Theorem 3.4. *Let P be the transition matrix of a given MDP. The entropy of the steady-state distribution d^π induced by a policy π is lower bounded by:*

$$H(d^\pi) \geq \log |\mathcal{S}| - |\mathcal{S}| \left\| (I - (\Pi P)^T) \cdot \mathbf{1}_{|\mathcal{S}|} \right\|_1^2.$$

The optimization of this lower bound leads to the following linear program:

$$\underset{\Pi \in \Pi}{\text{minimize}} \quad \left\| (I - (\Pi P)^T) \cdot \mathbf{1}_{|\mathcal{S}|} \right\|_1. \quad (6)$$


Besides being a linear program, unlike the other optimization problems presented, Problem (6) does not require to optimize over the space of all the doubly stochastic matrices, thus significantly reducing the number of optimization variables ($|\mathcal{S}| + |\mathcal{S}||\mathcal{A}|$) and constraints ($2|\mathcal{S}| + |\mathcal{S}||\mathcal{A}|$ inequalities and $|\mathcal{S}|$ equalities). The linear program formulation of Problem (6) can be found in Appendix B.2.

3.2 Highly Exploring Policies over the State and Action Space

Although the policy resulting from the optimization of one of the above problems may lead to the most uniform exploration of the state space, the actual goal of the exploration phase is to collect enough information on the environment to optimize, at some point, a goal-conditioned policy (Pong et al. 2019). To this end, it is essential to have an exploratory policy that adequately covers the action space \mathcal{A} in any visited state. Unfortunately, the optimization of Problems (4), (5), (6) does not guarantee even that the obtained policy is stochastic. Thus, we need to embed in the problem a secondary objective that takes into account the exploration over \mathcal{A} . This can be done by enforcing a minimal entropy over actions in the policy to be learned, adding to (4), (5), (6) the following constraints:

$$\pi(a|s) \geq \xi, \quad \forall s \in \mathcal{S}, \quad \forall a \in \mathcal{A}, \quad (7)$$

where $\xi \in [0, \frac{1}{|\mathcal{A}|}]$. This secondary objective is actually in competition with the objective of uniform exploration over states. Indeed, an overblown incentive in the exploration over actions may limit the state distribution entropy of the optimal policy. Having a low probability of visiting a state decreases the likelihood of sampling an action from that state, hence, also reducing the exploration over actions. To illustrate that, Figure 1a shows state distribution entropies ($H(d^\pi)$) and state-action distribution entropies, i.e., $H(d^\pi \Pi)$, achieved by the optimal policy w.r.t. Problem (5) on the Single Chain domain (Furmston and Barber 2010) for different values of ξ .



(a) State entropy and state-action entropy

(b) State entropy and mixing time

Figure 1: (a) State distribution entropy ($H(d^\pi)$), state-action distribution entropy ($H(d^\pi \Pi)$) for different values of ξ on the Single Chain domain. (b) State distribution entropy ($H(d^\pi)$), spectral gap (γ_π) for different values of ζ on the Single Chain domain (left), color-coded state distribution overlaid on a 4-rooms gridworld for different values of ζ (right).

3.3 An Objective to Make Highly Exploring Policies Mix Faster

In many cases, such as in episodic tasks where the horizon for exploration is capped, we may have interest in trading inferior state entropy for faster convergence of the learned policy. Although the doubly stochastic matrices are equally valid in terms of steady-state distribution, the choice of the target P^u strongly affects the mixing properties of the P^π induced by the policy. Indeed, while an MC with a uniform transition matrix, i.e., transition probabilities $P^u(s, s') = \frac{1}{|S|}$ for any s, s' , mixes in no time, an MC with probability one on the self-loops never converges to a steady state. This is evident considering that the mixing time t_{mix} of an MC is trapped as follows (Levin and Peres 2017, Theorems 12.3 and 12.4):

$$\frac{1 - \gamma_\pi}{\gamma_\pi} \log \frac{1}{2\epsilon} \leq t_{\text{mix}} \leq \frac{1}{\gamma_\pi} \log \frac{1}{d_{\min}^\pi \epsilon}, \quad (8)$$

where ϵ is the mixing threshold, d_{\min}^π is a minorization of d^π , and γ_π is the spectral gap of P^π (2). From the literature of MCs, we know that a variant of the Problems (4), (5) having the uniform transition matrix as target P^u and the L_2 as matrix norm, is equivalent to the problem of finding the fastest mixing transition matrix P^π (Boyd, Diaconis, and Xiao 2004). However, the choice of this target may overly limit the entropy over the state distribution induced by the optimal policy. Instead, we look for a generalization that allows us to prioritize fast exploration at will. Thus, we consider a continuum of relaxations in the fastest mixing objective by embedding in Problems (4) and (5) (but not in Problem (6)) the following constraints:

$$P^u(s, s') \leq \zeta, \quad \forall s, s' \in S, \quad (9)$$

where $\zeta \in [\frac{1}{|S|}, 1]$. By setting $\zeta = \frac{1}{|S|}$, we force the optimization problem to consider the uniform transition matrix as a target, thus aiming to reduce the mixing time, while larger values of ζ relax this objective, allowing us to get a higher steady-state distribution entropy. In Figure 1b we show how the parameter ζ affects the trade-off between high steady-state entropy and low mixing times (i.e., high spectral gaps), reporting the values obtained by optimal policies w.r.t. Problem (5) for different ζ .

4 A Model-Based Algorithm for Highly Exploring and Fast Mixing Policies

In this section, we present an approach to incrementally learn a highly exploring and fast mixing policy through interactions with an unknown environment, developing a novel model-based exploration algorithm called *Intrinsically-Driven Effective and Efficient Exploration ALgorithm* (IDE³AL). Since Problems (4), (5), (6) requires an explicit representation of the matrix P , we need to estimate the transition model from samples before performing an objective optimization (model-based approach). In tabular settings, this can be easily done by adopting the transition frequency as a proxy for the (unknown) transition probabilities, obtaining an estimated transition model $\hat{P}(s'|s, a)$. However, in hard-exploration tasks, it can be arbitrarily arduous to sample transitions from the most difficult-to-reach states by relying on naïve exploration mechanisms, such as a random policy. To address the issue, we lean on an iterative approach in which we alternate model estimation phases with optimization sweeps of the objectives (4), (5) or (6). In this way, we combine the benefit of collecting samples with highly exploring policies to better estimate the transition model and the benefit of having a better-estimated model to learn superior exploratory policies. In order to foster the policy towards (s, a) pairs that have never been sampled, we keep their corresponding distribution $\hat{P}(\cdot|s, a)$ to be uniform over all possible states, thus making the pair (s, a) particularly valuable in the perspective of the optimization problem. The algorithm converges whenever the exploratory policy remains unchanged during consecutive optimization sweeps and, if we know the size of the MDP, when all state-action pairs have been sufficiently explored. In Algorithm 1 we report the pseudo-code of IDE³AL. Finally, in Figure 2 we compare the iterative formulation against a not-iterative one, i.e., an approach that collects samples with a random policy and then optimizes the exploration objective off-line. Considering an exploration task on the Double Chain domain (Furmston and Barber 2010), we show that the iterative form has a clear edge in reducing the model estimation error $\|P - \hat{P}\|_F$. Both the approaches employ a Frobenius formulation.

Algorithm 1 IDE³AL

Input: ξ, ζ , batch size N Initialize π_0 and transition counts $C \in \mathbb{N}^{|\mathcal{S}|^2 \times |\mathcal{A}|}$ **for** $i = 0, 1, 2, \dots$, until convergence **do**Collect N steps with π_i and update C

Estimate the transition model as:

$$\hat{P}_i(s'|s, a) = \begin{cases} \frac{C(s'|s, a)}{\sum_{s'} C(s'|s, a)}, & \text{if } C(\cdot|s, a) > 0 \\ 1/|\mathcal{S}|, & \text{otherwise} \end{cases}$$

 $\pi_{i+1} \leftarrow$ optimal policy for (4) (or (5) or (6)),given the parameters ξ, ζ , and \hat{P}_i **end for****Output:** exploratory policy π_i

5 Experimental Evaluation

In this section, we provide the experimental evaluation of IDE³AL. First, we show a set of experiments on the illustrative *Single Chain* and *Double Chain* domains (Furmston and Barber 2010; Peters, Mulling, and Altun 2010). The Single Chain consists of 10 states having 2 possible actions, one to climb up the chain from state 0 to 9, and the other to directly fall to the initial state 0. The two actions are flipped with a probability $p_{slip} = 0.1$, making the environment stochastic and reducing the probability of visiting the higher states. The Double Chain concatenates two Single Chain into a bigger one sharing the central state 9, which is the initial state. Thus, the chain can be climbed in two directions. These two domains, albeit rather simple from a dimensionality standpoint, are actually hard to explore uniformly, due to the high shares of actions returning to the initial state and preventing the agent to consistently reach the higher states. Then, we present an experiment on the much more complex *Knight Quest* environment (Fruit et al. 2018, Appendix), having $|\mathcal{S}| = 360$ and $|\mathcal{A}| = 8$. This domain takes inspiration from classical arcade games, in which a knight has to rescue a princess in the shortest possible time without being killed by the dragon. To accomplish this feat, the knight has to perform an intricate sequence of actions. In the absence of any reward, it is a fairly challenging environment for exploration. On these domains, we address the task of learning the best exploratory policy in a limited number of samples. Especially, we evaluate these policies in terms of the induced state entropy $H(\mathbf{d}^\pi)$ and state-action entropy $H(\mathbf{d}^\pi \mathbf{\Pi})$.

We compare our approach with *MaxEnt* (Hazan et al. 2019), the model-based algorithm to learn maximum entropy exploration that we have previously discussed in the paper, and a count-based approach inspired by the exploration bonuses of MBIE-EB (Strehl and Littman 2008), which we refer as *CountBased* in the following. The latter shares the same structure of our algorithm, but replace the policy optimization sweeps with approximate value iterations (Bertsekas 1995), where the reward for a given state is inversely proportional to the visit count of that state. It is worth noting that the results reported for the MaxEnt algorithm are related to the mixture policy $\pi_{\text{mix}} = (\mathcal{D}, \alpha)$, where $\mathcal{D} = (\pi_0, \dots, \pi_{k-1})$ is a set of k ϵ -deterministic policies, and $\alpha \in \Delta_k$ is a probability distribution over \mathcal{D} . For the sake of simplicity, we have equipped

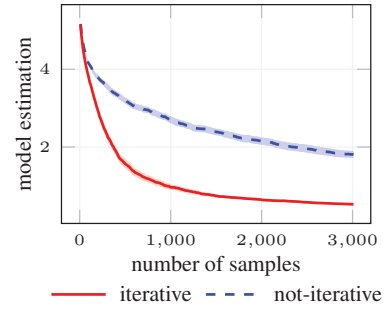


Figure 2: Model estimation error on the Double Chain with $\xi = 0.1, \zeta = 0.7, N = 10$ (100 runs, 95% c.i.).

all the approaches with a little domain knowledge, i.e., the cardinality of \mathcal{S} and \mathcal{A} . However, this can be avoided without a significant impact on the presented results. For every experiment, we will report the batch-size N , and the parameters ξ, ζ of IDE³AL. CountBased and MaxEnt employ ϵ -greedy policies having $\epsilon = \xi$ in all the experiments. In any plot, we will additionally provide the performance of a baseline policy, denoted as *Random*, that randomly selects an action in every state. Detailed information about the presented results, along with an additional experiment, can be found in Appendix D.

First, in Figure 3, we compare the Problems (4), (5), (6) on the Single Chain environment. On one hand, we show the performance achieved by the exact solutions, i.e., computed with a full knowledge of \mathbf{P} . While the plain formulations ($\xi = 0, \zeta = 1$) are remarkably similar, adding a constraint over the action entropy ($\xi = 0.1$) has a significantly different impact. On the other hand, we illustrate the performance of IDE³AL, equipped with the alternative optimization objectives, in learning a good exploratory policy from samples. In this case, the Frobenius clearly achieves a better performance. In the following, we will report the results of IDE³AL considering only the best-performing formulation, which, for all the presented experiments, corresponds to the Frobenius.

In Figure 4a, we show that IDE³AL compares well against the other approaches in exploring the Double Chain domain. It achieves superior state entropy and state-action entropy, and it converges faster to the optimum. It displays also a higher probability of visiting the least favorable state, and it behaves positively in the estimation of $\hat{\mathbf{P}}$. Notably, the CountBased algorithm fails to reach high exploration due to a detachment problem (Ecoffet et al. 2019), since it fluctuates between two exploratory policies that are greedy towards the two directions of the chain. By contrast, in a domain having a clear direction for exploration, such as the simpler Single Chain domain, CountBased ties the explorative performances of IDE³AL (Figure 4b). On the other hand, MaxEnt is effective in the exploration performance, but much more slower to converge, both in the Double Chain and the Single Chain. Note that in Figure 4a, the model estimation error of MaxEnt starts higher than the other, since it employs a different strategy to fill the transition probabilities of never reached states, inspired by (Brafman and Tennenholtz 2002). In Figure 4c, we present an experiment on the higher-dimensional

| | $H(\mathbf{d}^\pi)$ | $\min \mathbf{d}^\pi$ |
|----------------------------|---------------------|-----------------------|
| Frobenius ($\xi = 0$) | 0.98 | $6.4 \cdot 10^{-2}$ |
| Infinity ($\xi = 0$) | 0.98 | $6.4 \cdot 10^{-2}$ |
| Column Sum ($\xi = 0$) | 0.98 | $6 \cdot 10^{-2}$ |
| Frobenius ($\xi = 0.1$) | 0.94 | $4.1 \cdot 10^{-2}$ |
| Infinity ($\xi = 0.1$) | 0.89 | $2.6 \cdot 10^{-2}$ |
| Column Sum ($\xi = 0.1$) | 0.95 | $3.8 \cdot 10^{-2}$ |

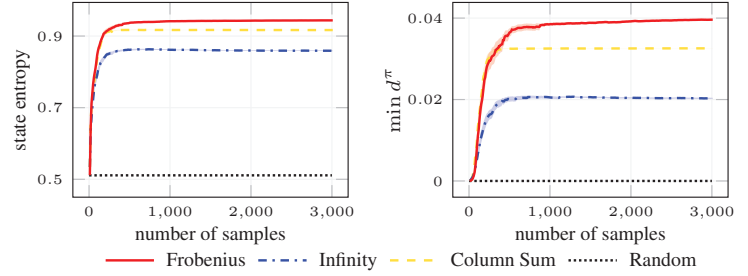


Figure 3: State distribution entropy ($H(\mathbf{d}^\pi)$) and probability of the least favorable state ($\min \mathbf{d}^\pi$) for different objective formulations on the Single Chain domain. We report exact solutions with $\zeta = 0$ (left), and approximate optimizations with $\xi = 0.1$, $\zeta = 0.7$, $N = 10$ (100 runs, 95% c.i.) (right).

Knight Quest environment. IDE³AL achieves a remarkable state entropy, while MaxEnt struggles to converge towards a satisfying exploratory policy. CountBased (not reported in Figure 4c, see Appendix D), fails to explore the environment altogether, oscillating between policies with low entropy.

In Figure 4d, we illustrate how the exploratory policies learned in the Double Chain environment are effective to ease learning of any possible goal-conditioned policy afterwards. To this end, the exploratory policies, learned by the three approaches through 3000 samples (Figure 4a), are employed to collect samples in a fixed horizon (within a range from 10 to 100 steps). Then, a goal-conditioned policy is learned off-line through approximate value iteration (Bertsekas 1995) on this small amount of samples. The goal is to optimize a reward function that is 1 for the hardest state to reach (i.e., the state that is less frequently visited with a random policy), 0 in all the other states. In this setting, all the methods prove to be rather successful w.r.t. the baseline, though IDE³AL compares positively against the other strategies.

6 Discussion

In this section, we first discuss how the proposed approach might be extended beyond tabular settings and an alternative formulation for the policy entropy optimization. Then, we consider some relevant work related to this paper.

6.1 Potential Extension to Continuous

We believe that the proposed approach has potential to be extended to more general, continuous, settings, by exploiting the core idea of avoiding a probability concentration on a subset of outgoing transitions from a state. Indeed, a compelling feature of the presented lower bounds is that they characterize an infinite-step property, the entropy of the steady-state distribution, relying only on one-step quantities, i.e., without requiring to unroll several times the state transition matrix P^π . In addition to this, the lower bounds provide an evaluation for the current policy, and they can be computed for any policy. Thus, we could potentially operate a direct search in the policy space through the gradient of an approximation of these lower bounds. To perform the approximation we could use a kernel for a soft aggregation over regions of the, now continuous, state space.

6.2 A Dual Formulation

A potential alternative to deal with the optimization of the objective (3) is to consider its dual formulation. This is rather similar to the approach proposed in (Tarbouriech and Lazaric 2019) to address the different problem of active exploration in an MDP. The basic idea is to directly maximize the entropy over the *state-action* stationary distribution and then to recover the policy afterwards. In this setting, we define the state-action stationary distribution induced by a policy π as $\omega^\pi = \mathbf{d}^\pi \mathbf{\Pi}$, where ω^π is a vector of size $|\mathcal{S}||\mathcal{A}|$ having elements $\omega^\pi(s, a)$. Since not all the distribution over the state-action space can be actually induced by a policy over the MDP, we characterize the set of feasible distributions:

$$\Omega = \left\{ \omega \in \Delta(\mathcal{S} \times \mathcal{A}) : \forall s \in \mathcal{S}, \sum_{a \in \mathcal{A}} \omega(s, a) = \sum_{s' \in \mathcal{S}, a' \in \mathcal{A}} P(s|s', a') \omega(s', a') \right\}.$$

Then, we can formulate the Dual Problem as:

$$\underset{\omega \in \Omega}{\text{maximize}} \quad H(\omega) \quad (10)$$

Finally, let ω^* denotes the solution of Problem (10), we can recover the policy inducing the optimal state-action entropy as $\pi_{\omega^*}(a|s) = \omega^*(s, a) / \sum_{a' \in \mathcal{A}} \omega^*(s, a')$, $\forall s \in \mathcal{S}, \forall a \in \mathcal{A}$.

The Dual Problem displays some appealing features. Especially, the objective in (10) is already convex, so that it can be optimized right away, and it allows to explicitly maximize the entropy over the state-action space. Nonetheless, we think that this alternative formulation has three major shortcomings. First, the optimization of the convex program (10) could be way slower than the optimization of the linear programs Column Sum and Infinity (Grötschel, Lovász, and Schrijver 1993). Secondly, it does not allow to control the mixing time of the learned policy, which can be extremely relevant. Lastly, the applicability of the Dual Problem to continuous environments seems far-fetched. It is worth noting that, from an empirical evaluation, the dual formulation does not provide any significant benefit in the entropy of the learned policy w.r.t. the lower bounds formulations (see Appendix D). Figure 4e shows how the solve time of the Column Sum scales better with the number of variables ($|\mathcal{S}||\mathcal{A}|$) in incrementally large Knight Quest domains.

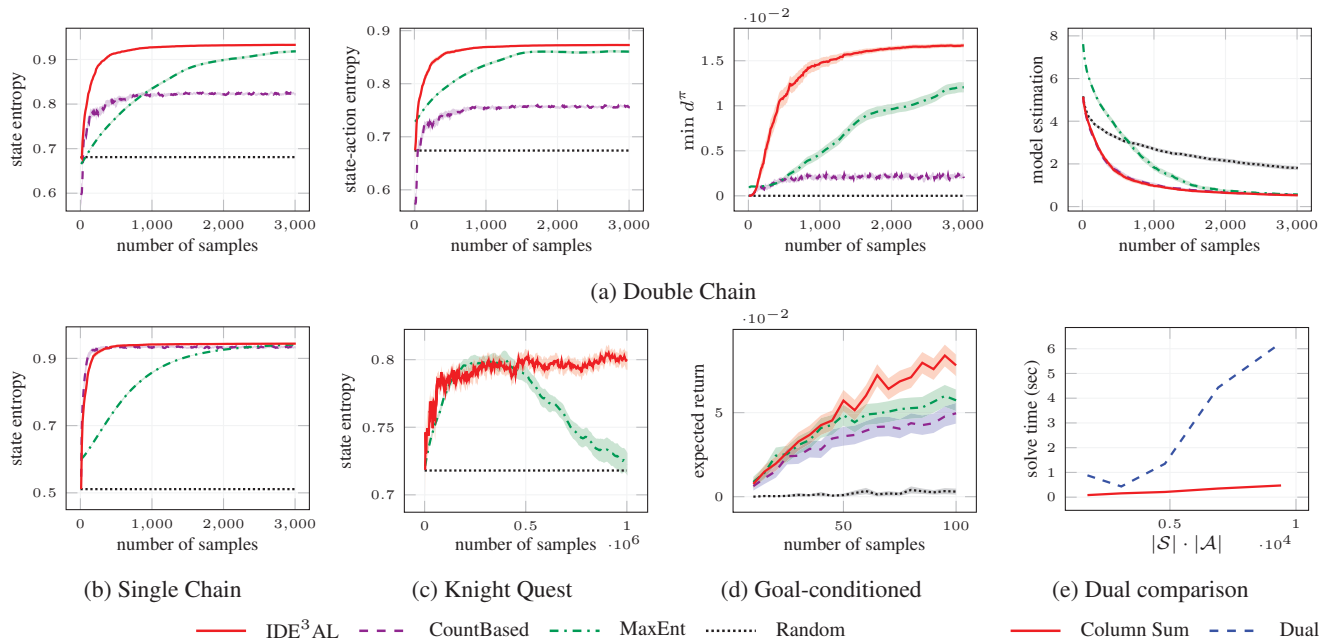


Figure 4: Comparison of the algorithms on exploration tasks (a, b, c) and goal-conditioned learning (d), with parameters $\xi = 0.1$, $\zeta = 0.7$, $N = 10$ (a, b, d) and $\xi = 0.01$, $\zeta = 1$, $N = 2500$ (c). (95% c.i. over 100 runs (a, b), 40 runs (c), 500 runs (d)). Comparison of the solve time (e) achieved by Column Sum and Dual formulations as a function of the number of variables.

6.3 Related Work

As discussed in the previous sections, Hazan et al. (2019) consider an objective not that dissimilar to the one presented in this paper, even if they propose a fairly different solution to the problem. Their method learns a mixture of deterministic policies instead of a single stochastic policy. In a similar flavor, Tarbouriech and Lazaric (2019) develop an approach, based on a dual formulation of the objective, to learn a mixture of stochastic policies for active exploration.

Other propose to intrinsically motivate the agent towards learning to reach all possible states in the environment (Lim and Auer 2012). To extend this same idea from the tabular setting to the context of a continuous, high-dimensional state space, Pong et al. (2019) employ a generative model to seek for a maximum-entropy goal distribution. Ecoffet et al. (2019) propose a method, called Go-Explore, to methodically reach any state by keeping an archive of any visited state and the best trajectory that brought the agent there. At each iteration, the agent draws a promising state from the archive, returns there replicating the stored trajectory (Go), then explores from this state trying to discover new states (Explore).

Another promising intrinsic objective is to make value out of the exploration phase by acquiring a set of reusable skills, typically formulated by means of the option framework (Sutton, Precup, and Singh 1999). In (Barto, Singh, and Chentanez 2004), a set of options is learned by maximizing an intrinsic reward that is generated at the occurrence of some, user-defined, salient event. The approach proposed by Bonarini, Lazaric, and Restelli (2006), which presents some similarities with the work in (Ecoffet et al. 2019), is based on learning a set of options to return with high proba-

bility to promising states. In their context, a promising state presents high unbalance between the probabilities of the input and output transitions (Bonarini et al. 2006), so that it is both a hard state to reach, and a doorway to reach many other states. In this way, the learned options heuristically favor an even exploration of the state space.

7 Conclusions

In this paper, we proposed a new model-based algorithm, IDE³AL, to learn highly exploring and fast mixing policies. The algorithm outputs a policy that maximizes a lower bound to the entropy of the steady-state distribution. We presented three formulations of the lower bound that differently tradeoff tightness with computational complexity of the optimization. The experimental evaluation showed that IDE³AL is able to achieve superior performance than other approaches striving for uniform exploration of the environment, while it avoids the risk of detachment and derailment (Ecoffet et al. 2019). Future works could focus on extending the applicability of the presented approach to non-tabular environments, following the blueprint in Section 6.1. We believe that this work provides a valuable contribution in view of solving the conundrum on what should a reinforcement learning agent learn in the absence of any reward coming from the environment.

Acknowledgments

This work has been partially supported by the Italian MIUR PRIN 2017 Project ALGADIMAR “Algorithms, Games, and Digital Market”.

References

- Barto, A. G.; Singh, S.; and Chentanez, N. 2004. Intrinsically motivated learning of hierarchical collections of skills. In *International Conference on Development and Learning*.
- Bellemare, M.; Srinivasan, S.; Ostrovski, G.; Schaul, T.; Saxton, D.; and Munos, R. 2016. Unifying count-based exploration and intrinsic motivation. In *Neural Information Processing Systems*.
- Bertsekas, D. P. 1995. *Dynamic programming and optimal control*. Athena Scientific.
- Bonarini, A.; Lazaric, A.; Restelli, M.; and Vitali, P. 2006. Self-development framework for reinforcement learning agents. In *International Conference on Development and Learning*.
- Bonarini, A.; Lazaric, A.; and Restelli, M. 2006. Incremental skill acquisition for self-motivated learning animats. In *International Conference on Simulation of Adaptive Behavior*.
- Boyd, S.; Diaconis, P.; and Xiao, L. 2004. Fastest mixing markov chain on a graph. *SIAM review*.
- Brafman, R. I., and Tennenholtz, M. 2002. R-max-a general polynomial time algorithm for near-optimal reinforcement learning. *Journal of Machine Learning Research*.
- Burda, Y.; Edwards, H.; Pathak, D.; Storkey, A.; Darrell, T.; and Efros, A. A. 2019a. Large-scale study of curiosity-driven learning. In *International Conference on Learning Representations*.
- Burda, Y.; Edwards, H.; Storkey, A.; and Klimov, O. 2019b. Exploration by random network distillation. In *International Conference on Learning Representations*.
- Chentanez, N.; Barto, A. G.; and Singh, S. P. 2005. Intrinsically motivated reinforcement learning. In *Neural Information Processing Systems*.
- Csiszár, I., and Talata, Z. 2006. Context tree estimation for not necessarily finite memory processes, via bic and mdl. *IEEE Transactions on Information Theory*.
- Diamond, S., and Boyd, S. 2016. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*.
- Ecoffet, A.; Huizinga, J.; Lehman, J.; Stanley, K. O.; and Clune, J. 2019. Go-explore: a new approach for hard-exploration problems. *arXiv preprint arXiv:1901.10995*.
- Fruit, R.; Pirota, M.; Lazaric, A.; and Ortner, R. 2018. Efficient bias-span-constrained exploration-exploitation in reinforcement learning. In *International Conference on Machine Learning*.
- Furnston, T., and Barber, D. 2010. Variational methods for reinforcement learning. In *International Conference on Artificial Intelligence and Statistics*.
- Grötschel, M.; Lovász, L.; and Schrijver, A. 1993. The ellipsoid method. In *Geometric Algorithms and Combinatorial Optimization*.
- Hassibi, B.; Hansen, M.; Dimakis, A. G.; Alshamary, H. A. J.; and Xu, W. 2014. Optimized markov chain monte carlo for signal detection in mimo systems: An analysis of the stationary distribution and mixing time. *IEEE Transactions on Signal Processing*.
- Hazan, E.; Kakade, S.; Singh, K.; and Van Soest, A. 2019. Provably efficient maximum entropy exploration. In *International Conference on Machine Learning*.
- Houthoofd, R.; Chen, X.; Duan, Y.; Schulman, J.; De Turck, F.; and Abbeel, P. 2016. Vime: Variational information maximizing exploration. In *Neural Information Processing Systems*.
- Hunter, J. J. 2010. Some stochastic properties of “semi-magic” and “magic” markov chains. *Linear Algebra and its Applications*.
- Kirkland, S. 2010. Column sums and the conditioning of the stationary distribution for a stochastic matrix. *Operators and Matrices*.
- Levin, D. A., and Peres, Y. 2017. *Markov chains and mixing times*. American Mathematical Soc.
- Lim, S. H., and Auer, P. 2012. Autonomous exploration for navigating in mdps. In *Conference on Learning Theory*.
- Lopes, M.; Lang, T.; Toussaint, M.; and Oudeyer, P.-Y. 2012. Exploration in model-based reinforcement learning by empirically estimating learning progress. In *Advances in Neural Information Processing Systems*.
- 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. 2015. Human-level control through deep reinforcement learning. *Nature*.
- Mohamed, S., and Rezende, D. J. 2015. Variational information maximisation for intrinsically motivated reinforcement learning. In *Neural Information Processing Systems*.
- Ostrovski, G.; Bellemare, M. G.; van den Oord, A.; and Munos, R. 2017. Count-based exploration with neural density models. In *International Conference on Machine Learning*.
- Oudeyer, P.-Y., and Kaplan, F. 2009. What is intrinsic motivation? a typology of computational approaches. *Frontiers in neurorobotics*.
- Pathak, D.; Agrawal, P.; Efros, A. A.; and Darrell, T. 2017. Curiosity-driven exploration by self-supervised prediction. In *International Conference on Machine Learning*.
- Peters, J.; Mulling, K.; and Altun, Y. 2010. Relative entropy policy search. In *24th AAAI Conference on Artificial Intelligence*.
- Petersen, K. B.; Pedersen, M. S.; et al. 2008. *The matrix cookbook*. Technical University of Denmark.
- Pong, V. H.; Dalal, M.; Lin, S.; Nair, A.; Bahl, S.; and Levine, S. 2019. Skew-fit: state-covering self-supervised reinforcement learning. *arXiv preprint arXiv:1903.03698*.
- Puterman, M. L. 2014. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons.
- Schmidhuber, J. 1991. A possibility for implementing curiosity and boredom in model-building neural controllers. In *International Conference on Simulation of Adaptive Behavior*.
- Schweitzer, P. J. 1968. Perturbation theory and finite markov chains. *Journal of Applied Probability*.
- Silver, D.; Schrittwieser, J.; Simonyan, K.; Antonoglou, I.; Huang, A.; Guez, A.; Hubert, T.; Baker, L.; Lai, M.; Bolton, A.; et al. 2017. Mastering the game of go without human knowledge. *Nature*.
- Stadie, B. C.; Levine, S.; and Abbeel, P. 2015. Incentivizing exploration in reinforcement learning with deep predictive models. *arXiv preprint arXiv:1507.00814*.
- Strehl, A. L., and Littman, M. L. 2008. An analysis of model-based interval estimation for markov decision processes. *Journal of Computer and System Sciences*.
- Sutton, R. S., and Barto, A. G. 2018. *Reinforcement learning: An introduction*. MIT press.
- Sutton, R. S.; Precup, D.; and Singh, S. 1999. Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence*.
- Tang, H.; Houthoofd, R.; Foote, D.; Stooke, A.; Chen, O. X.; Duan, Y.; Schulman, J.; DeTurck, F.; and Abbeel, P. 2017. # exploration: A study of count-based exploration for deep reinforcement learning. In *Neural Information Processing Systems*.
- Tarbouriech, J., and Lazaric, A. 2019. Active exploration in markov decision processes. In *International Conference on Artificial Intelligence and Statistics*.