

---

# Path Consistency Learning in Tsallis Entropy Regularized MDPs

---

Ofir Nachum<sup>\*1</sup> Yinlam Chow<sup>\*2</sup> Mohamamd Ghavamzadeh<sup>\*2</sup>

## Abstract

We study the sparse entropy-regularized reinforcement learning (ERL) problem in which the entropy term is a special form of the *Tsallis* entropy. The optimal policy of this formulation is sparse, i.e., at each state, it has non-zero probability for only a small number of actions. This addresses the main drawback of the standard Shannon entropy-regularized RL (soft ERL) formulation, in which the optimal policy is *softmax*, and thus, may assign a non-negligible probability mass to non-optimal actions. This problem is aggravated as the number of actions is increased. In this paper, we follow the work of Nachum et al. (2017) in the soft ERL setting, and propose a class of novel path consistency learning (PCL) algorithms, called *sparse PCL*, for the sparse ERL problem that can work with both on-policy and off-policy data. We first derive a *sparse consistency* equation that specifies a relationship between the optimal value function and policy of the sparse ERL along any system trajectory. Crucially, a weak form of the converse is also true, and we quantify the sub-optimality of a policy which satisfies sparse consistency, and show that as we increase the number of actions, this sub-optimality is better than that of the soft ERL optimal policy. We then use this result to derive the sparse PCL algorithms. We empirically compare sparse PCL with its soft counterpart, and show its advantage, especially in problems with a large number of actions.

## 1. Introduction

In reinforcement learning (RL), the goal is to find a policy with maximum long-term performance, defined as the sum of discounted rewards generated by following the policy (Bertsekas & Tsitsiklis, 1996; Sutton & Barto, 1998). In case the number of states and actions are small, and the

model is known, the optimal policy is the solution of the *non-linear* Bellman optimality equations (Bellman, 1957). When the system is large or the model is unknown, greedily solving the Bellman equations often results in policies that are far from optimal. A principled way of dealing with this issue is *regularization*. Among different forms of regularization, such as  $\ell_2$  (e.g., Farahmand et al. 2008; 2009) and  $\ell_1$  (e.g., Kolter & Ng 2009; Johns et al. 2010; Ghavamzadeh et al. 2011), *entropy regularization* is among the most studied in both value-based (e.g., Kappen 2005; Todorov 2006; Ziebart 2010; Azar et al. 2012; Fox et al. 2016; O’Donoghue et al. 2017; Asadi & Littman 2017) and policy-based (e.g., Peters et al. 2010; Todorov 2010) RL formulations. In particular, two of the most popular deep RL algorithms, TRPO (Schulman et al., 2015) and A3C (Mnih et al., 2016), are based on entropy-regularized policy search. We refer the interested readers to Neu et al. (2017), for an insightful discussion on entropy-regularized RL algorithms and their connection to online learning.

In entropy-regularized RL (ERL), an entropy term is added to the Bellman equation. This formulation has four main advantages: **1**) it softens the non-linearity of the Bellman equations and makes it possible to solve them more easily, **2**) the solution of the softened problem is quantifiably not much worse than the optimal solution in terms of accumulated return, **3**) the addition of the entropy term brings nice properties, such as encouraging exploration (Shannon entropy) (e.g., Fox et al. 2016; Nachum et al. 2017) and maintaining a close distance to a baseline policy (relative entropy) (e.g., Schulman et al. 2015; Belousov & Peters 2017; Nachum et al. 2018), and **4**) unlike the original problem that has a deterministic solution, the solution to the softened problem is stochastic, which is preferable in problems in which exploration or dealing with unexpected situations is important. However, in the most common form of ERL, in which a Shannon (or relative) entropy term is added to the Bellman equations, the optimal policy is of the form of *softmax*. Despite the advantages of a softmax policy in terms of exploration, its main drawback is that at each step, it assigns a non-negligible probability mass to non-optimal actions, a problem that is aggravated as the number of actions is increased. This may result in policies that may not be safe to execute. To address this issue, Lee et al. (2018) proposed to add a special form of a general no-

---

<sup>\*</sup>Equal contribution <sup>1</sup>Google Brain <sup>2</sup>DeepMind. Correspondence to: Yinlam Chow <yinlamchow@google.com>.

tion of entropy, called Tsallis entropy (Tsallis, 1988), to the Bellman equations. This formulation has the property that its solution has sparse distributions, i.e., at each state, only a small number of actions have non-zero probability. Lee et al. (2018) studied the properties of this ERL formulation, proposed value-based algorithms (fitted Q-iteration and Q-learning) to solve it, and showed that although it is harder to solve than its soft counterpart, it potentially has a solution closer to that of the original problem.

In this paper, we propose novel path consistency learning (PCL) algorithms for the Tsallis ERL problem, called *sparse PCL*. PCL is a class of actor-critic type algorithms developed by Nachum et al. (2017) for the soft (Shannon entropy) ERL problem. It uses a nice property of soft ERL, namely the equivalence of consistency and optimality, and learns parameterized policy and value functions by minimizing a loss that is based on the consistency equation of soft ERL. The most notable feature of soft PCL is that it can work with both on-policy (sub-trajectories generated by the current policy) and off-policy (sub-trajectories generated by a policy different than the current one, including any sub-trajectory from the replay buffer) data. We first derive a multi-step consistency equation for the Tsallis ERL problem, called *sparse consistency*. We then prove that in this setting, while optimality implies consistency (similar to the soft case), unlike the soft case, consistency only implies sub-optimality. We then use the sparse consistency equation and derive PCL algorithms that use both on-policy and off-policy data to solve the Tsallis ERL problem. We empirically compare sparse PCL with its soft counterpart. As expected, we gain from using the sparse formulation when the number of actions is large, both in algorithmic tasks and in discretized continuous control problems.

## 2. Markov Decision Processes (MDPs)

We consider the reinforcement learning (RL) problem in which the agent’s interaction with the system is modeled as a MDP. A MDP is a tuple  $\mathcal{M} = (\mathcal{X}, \mathcal{A}, r, P, P_0, \gamma)$ , where  $\mathcal{X}$  and  $\mathcal{A}$  are state and action spaces;  $r : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}$  and  $P : \mathcal{X} \times \mathcal{A} \rightarrow \Delta_{\mathcal{X}}$  are the reward function and transition probability distribution, with  $r(x, a) \in [0, R_{\max}]$  and  $P(\cdot|x, a)$  being the reward and the next state probability of taking action  $a$  in state  $x$ ;  $P_0 : \mathcal{X} \rightarrow \Delta_{\mathcal{X}}$  is the initial state distribution; and  $\gamma \in [0, 1)$  is a discounting factor. In this paper, we assume that the action space is finite, but can be large. The goal in RL is to find a stationary Markovian policy, i.e., a mapping from state and action spaces to a simplex over the actions  $\mu : \mathcal{X} \times \mathcal{A} \rightarrow \Delta_{\mathcal{A}}$ , that maximizes the expected discounted sum of rewards, i.e.,

$$\begin{aligned} \max_{\mu} \mathbb{E} \left[ \sum_{t=0}^{\infty} \gamma^t r(x_t, a_t) \right] \\ \text{s.t. } \forall x \sum_{a \in \mathcal{A}} \mu(a|x) = 1, \quad \forall x, a \mu(a|x) \geq 0, \end{aligned} \quad (1)$$

where  $x_0 \sim P_0$ ,  $a_t \sim \mu(\cdot|x_t)$ , and  $x_{t+1} \sim P(\cdot|x_t, a_t)$ . For a given policy  $\mu$ , we define its value and action-value functions as

$$\begin{aligned} V^{\mu}(x) &= \mathbb{E} \left[ \sum_{t=0}^{\infty} \gamma^t r(x_t, a_t) | x_0 = x, \mu, P \right], \\ Q^{\mu}(x, a) &= \mathbb{E} \left[ \sum_{t=0}^{\infty} \gamma^t r(x_t, a_t) | x_0 = x, a_0 = a, \mu, P \right]. \end{aligned}$$

Any solution of the optimization problem (1) is called an *optimal* policy and is denoted by  $\mu^*$ . Note that while a MDP may have several optimal policies, it only has a single optimal value function  $V^* = V^{\mu^*}$ . It has been proven that (1) has a solution in the space of *deterministic* policies, i.e.,  $\Pi_d = \{\mu : \mu : \mathcal{X} \rightarrow \mathcal{A}\}$ , which can be obtained as the *greedy* action w.r.t. the optimal action-value function, i.e.,  $\mu^*(x) \in \arg \max_a Q^*(x, a)$  (Puterman, 1994; Bertsekas & Tsitsiklis, 1996). The optimal action-value function  $Q^*$  is the *unique* solution of the non-linear Bellman optimality equations, i.e., for all  $x \in \mathcal{X}$  and  $a \in \mathcal{A}$ ,

$$Q(x, a) = r(x, a) + \gamma \sum_{x' \in \mathcal{X}} P(x'|x, a) \max_{a' \in \mathcal{A}} Q(x', a'). \quad (2)$$

Any optimal policy  $\mu^*$  and the optimal state and state-action value functions,  $V^*$  and  $Q^*$ , satisfy the following equations for all states and action,

$$\begin{aligned} Q^*(x, a) &= r(x, a) + \gamma \sum_{x' \in \mathcal{X}} P(x'|x, a) V^*(x'), \\ V^*(x) &= \max_{a \in \mathcal{A}} Q^*(x, a), \quad \mu^*(x) \in \arg \max_{a \in \mathcal{A}} Q^*(x, a). \end{aligned}$$

## 3. Entropy Regularized MDPs

As discussed in Section 2, finding an optimal policy for a MDP involves solving a non-linear system of equations (see Eq. 2), which is often complicated. Moreover, the optimal policy may be deterministic, always selecting the same optimal action at a state even when there are several optimal actions in that state. This is undesirable when it is important to explore and to deal with unexpected situations. In such cases, one might be interested in multimodal policies that still have good performance. This is why many researchers have proposed to add a regularizer in the form of an *entropy* term to the objective function (1) and solve the following *entropy-regularized* optimization problem

$$\begin{aligned} \max_{\mu} \mathbb{E} \left[ \sum_{t=0}^{\infty} \gamma^t (r(x_t, a_t) + \alpha H^{\mu}(x_t, a_t)) \right] \\ \text{s.t. } \forall x \sum_{a \in \mathcal{A}} \mu(a|x) = 1, \quad \forall x, a \mu(a|x) \geq 0, \end{aligned} \quad (3)$$

where  $H^{\mu}(x, a)$  is an entropy-related term and  $\alpha$  is the regularization parameter. The entropy term smoothens the objective function (1) such that the resulting problem (3) is often easier to solve than the original one (1). This is another reason for the popularity of entropy-regularized MDPs.

### 3.1. Entropy Regularized MDP with Shannon Entropy

It is common to use  $H_{\text{sf}}^\mu(x_t, a_t) \triangleq -\log \mu(a_t|x_t)$  in entropy-regularized MDPs (e.g., Fox et al. 2016; Nachum et al. 2017). Note that  $H_{\text{sf}}(\mu) = \mathbb{E}_\mu[H_{\text{sf}}^\mu(x, a)]$  is the *Shannon entropy*. Problem (3) with  $H_{\text{sf}}^\mu(x, a)$  can be seen as a RL problem in which the reward function is the sum of the original reward function  $r(x, a)$  and a term that encourages *exploration*.<sup>1</sup> Unlike (1), the optimization problem (3) with  $H_{\text{sf}}^\mu$  has a unique optimal policy  $\mu_{\text{sf}}^*$  and a unique optimal value  $V_{\text{sf}}^*$  (action-value  $Q_{\text{sf}}^*$ ) function that satisfy the following equations:

$$\begin{aligned} Q_{\text{sf}}^*(x, a) &= r(x, a) + \gamma \sum_{x' \in \mathcal{X}} P(x'|x, a) V_{\text{sf}}^*(x'), \\ V_{\text{sf}}^*(x) &= \alpha \cdot \text{sfxmax}(Q_{\text{sf}}^*(x, \cdot)/\alpha), \\ \mu_{\text{sf}}^*(a|x) &= \frac{\exp(Q_{\text{sf}}^*(x, a)/\alpha)}{\sum_{a' \in \mathcal{A}} \exp(Q_{\text{sf}}^*(x, a')/\alpha)}, \end{aligned} \quad (4)$$

where for any function  $f: \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}$ , the *sfxmax* operator is defined as  $\text{sfxmax}(f(x, \cdot)) = \log(\sum_a \exp(f(x, a)))$ . Note that the equations in (4) are derived from the KKT conditions of (3) with  $H_{\text{sf}}^\mu$ . In this case, the optimal policy is *soft-max*, with the regularization parameter  $\alpha$  playing the role of its temperature (see Eq. 4). This is why (3) with  $H_{\text{sf}}^\mu$  is called the *soft MDP* problem. In soft MDPs, the optimal value function  $V_{\text{sf}}^*$  is the unique solution of the *soft Bellman optimality* equations, i.e.,  $\forall x \in \mathcal{X}, \forall a \in \mathcal{A}$ ,

$$V(x) = \alpha \cdot \text{sfxmax}\left([r(x, \cdot) + \gamma \sum_{x'} P(x'|x, \cdot) V(x')]/\alpha\right). \quad (5)$$

Note that the *sfxmax* operator is a smoother function of its inputs than the *max* operator associated with the Bellman optimality equation (2). This means that solving the soft MDP problem is easier than the original one, with the cost that its optimal policy  $\mu_{\text{sf}}^*$  performs worse than the optimal policy of the original MDP  $\mu^*$ . This difference can be quantified as

$$\forall x \in \mathcal{X} \quad V^*(x) - \frac{\alpha}{1-\gamma} \log(|\mathcal{A}|) \leq V^{\mu_{\text{sf}}^*}(x) \leq V^*(x), \quad (6)$$

where we discriminate between the value function of a policy  $\mu$  in the soft  $V_{\text{sf}}^\mu$  and original  $V^\mu$  MDPs. Note that the sub-optimality of  $\mu_{\text{sf}}^*$  is unbounded as  $|\mathcal{A}| \rightarrow \infty$ . This is the main drawback of using softmax policies; in large action space problems, at each step, the policy assigns a non-negligible probability mass to non-optimal actions, which in aggregate can be detrimental to its reward performance.

<sup>1</sup>Another entropy term that has been studied in the literature is  $H_{\text{rel}}^\mu(x_t, a_t) \triangleq -\log \frac{\mu(a_t|x_t)}{\pi_b(a_t|x_t)}$ , where  $\pi_b$  is a baseline policy. Note that  $H_{\text{rel}}(\mu) = \mathbb{E}_\mu[H_{\text{rel}}^\mu(x, a)]$  is the *relative entropy*. Problem (3) with  $H_{\text{rel}}^\mu(x, a)$  can be seen as a RL problem in which the reward function is the sum of the original reward function  $r(x, a)$  and a term that penalizes deviation from the baseline policy  $\pi_b$ .

### 3.2. Entropy Regularized MDP with Tsallis Entropy

To address the issues with the softmax policy, Lee et al. (2018) proposed to use  $H_{\text{sp}}^\mu(x_t, a_t) \triangleq \frac{1}{2}(1 - \mu(a_t|x_t))$  in entropy-regularized MDPs. Note that  $H_{\text{sp}}(\mu) = \mathbb{E}_\mu[H_{\text{sp}}^\mu(x, a)]$  is a special case of a general notion of entropy, called *Tsallis entropy* (Tsallis, 1988), i.e.,  $S_{q,k}(p) = \frac{k}{q-1}(1 - \sum_i p_i^q)$ , for the parameters  $q = 2$  and  $k = \frac{1}{2}$ .<sup>2</sup> Similar to the soft MDP problem, the optimization problem (3) with  $H_{\text{sp}}^\mu$  has a unique optimal policy  $\mu_{\text{sp}}^*$  and a unique optimal value  $V_{\text{sp}}^*$  (action-value  $Q_{\text{sp}}^*$ ) function that satisfy the following equations (Lee et al., 2018):

$$\begin{aligned} Q_{\text{sp}}^*(x, a) &= r(x, a) + \gamma \sum_{x' \in \mathcal{X}} P(x'|x, a) V_{\text{sp}}^*(x'), \\ V_{\text{sp}}^*(x) &= \alpha \cdot \text{spmax}(Q_{\text{sp}}^*(x, \cdot)/\alpha), \\ \mu_{\text{sp}}^*(a|x) &= \left(Q_{\text{sp}}^*(x, a)/\alpha - \mathcal{G}(Q_{\text{sp}}^*(x, \cdot)/\alpha)\right)^+, \end{aligned} \quad (7)$$

where  $(\cdot)^+ = \max(\cdot, 0)$ , and for any function  $f: \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}$ , the *spmax* operator is defined as

$$\text{spmax}(f(x, \cdot)) = \frac{1}{2} \left[ 1 + \sum_{a \in \mathcal{S}(x)} \left( \left( \frac{f(x, a)}{\alpha} \right)^2 - \mathcal{G} \left( \frac{f(x, \cdot)}{\alpha} \right)^2 \right) \right],$$

in which  $\mathcal{G}(f(x, \cdot)) = \frac{\sum_{a \in \mathcal{S}(x)} f(x, a) - 1}{|\mathcal{S}(x)|}$  and  $\mathcal{S}(x)$  is the set of actions satisfying  $1 + i \frac{f(x, a_{(i)})}{\alpha} > \sum_{j=0}^i \frac{f(x, a_{(j)})}{\alpha}$ , where  $a_{(i)}$  indicates the action with the  $i$ 'th largest value of  $f(x, a)$ . Note that the equations in (7) are derived from the KKT conditions of (3) with  $H_{\text{sp}}^\mu$ . In this case, the optimal policy may have zero probability for several actions (see Eq. 7). This is why (3) with  $H_{\text{sp}}^\mu$  is called the *sparse MDP* problem. The regularization parameter  $\alpha$  controls the sparsity of the resulted policy. The policy would be more sparse for smaller values of  $\alpha$ . In sparse MDPs, the optimal value function  $V_{\text{sp}}^*$  is the unique fixed-point of the *sparse Bellman optimality* operator  $\mathcal{T}_{\text{sp}}$  (Lee et al., 2018) that for any function  $f: \mathcal{X} \rightarrow \mathbb{R}$  is defined as

$$(\mathcal{T}_{\text{sp}}f)(x) = \alpha \cdot \text{spmax}\left([r(x, \cdot) + \gamma \sum_{x'} P(x'|x, \cdot) f(x')]/\alpha\right). \quad (8)$$

Similar to (5), the *spmax* operator is a smoother function of its inputs than the *max*, and thus, solving the sparse MDP problem is easier than the original one, with the cost that its optimal policy  $\mu_{\text{sp}}^*$  performs worse than the optimal policy of the original MDP  $\mu^*$ . This difference can be quantified as (Lee et al., 2018),

$$\forall x \in \mathcal{X} \quad V^*(x) - \frac{\alpha}{1-\gamma} \cdot \frac{|\mathcal{A}| - 1}{2|\mathcal{A}|} \leq V^{\mu_{\text{sp}}^*}(x) \leq V^*(x). \quad (9)$$

<sup>2</sup>Note that the Shannon entropy is a special case of the Tsallis entropy for the parameters  $q = k = 1$  (Tsallis, 1988).

On the other hand, the spmax operator is more complex than sfmax, and thus, it is slightly more complicated to solve the sparse MDP problem than its soft counterpart. However, as can be seen from Eqs. 6 and 9, the optimal policy of the sparse MDP,  $\mu_{\text{sp}}^*$ , can have a better performance than its soft counterpart,  $\mu_{\text{sf}}^*$ , and this difference becomes more apparent as the number of actions  $|\mathcal{A}|$  grows. For large action size, the term  $(|\mathcal{A}| - 1)/(2|\mathcal{A}|)$  in (9) turns to a constant, while  $\log |\mathcal{A}|$  in (6) grows unbounded.<sup>3</sup>

#### 4. Path Consistency Learning in Soft MDPs

A nice property of soft MDPs that was elegantly used by Nachum et al. (2017) is that any policy  $\mu$  and function  $V : \mathcal{X} \rightarrow \mathbb{R}$  that satisfy the (one-step) *consistency* equation, i.e., for all  $x \in \mathcal{X}$  and for all  $a \in \mathcal{A}$ ,

$$V(x) = r(x, a) - \alpha \log \mu(a|x) + \gamma \sum_{x' \in \mathcal{X}} P(x'|x, a) V(x'), \quad (10)$$

are optimal, i.e.,  $\mu = \mu_{\text{sf}}^*$  and  $V = V_{\text{sf}}^*$  (*consistency implies optimality*). Due to the uniqueness of the optimal policy in soft MDPs, the reverse is also true, i.e., the optimal policy  $\mu_{\text{sp}}^*$  and the value function  $V_{\text{sp}}^*$  satisfy the *consistency* equation (*optimality implies consistency*).

As shown in Nachum et al. (2017), the (one-step) consistency equation (10) can be easily extended to multi-step, i.e., any policy  $\mu$  and function  $V : \mathcal{X} \rightarrow \mathbb{R}$  that for any state  $x_0$  and sequence of actions  $a_0, \dots, a_{d-1}$ , satisfy the *multi-step consistency* equation

$$V(x_0) = \mathbb{E}_{x_{1:d}|x_0, a_{0:d-1}} \left[ \gamma^d V(x_d) + \sum_{t=0}^{d-1} \gamma^t (r(x_t, a_t) - \alpha \log \mu(a_t|x_t)) \right] \quad (11)$$

are optimal, i.e.,  $\mu = \mu_{\text{sf}}^*$  and  $V = V_{\text{sf}}^*$ .

The property that both single and multiple step consistency equations imply optimality (Eqs. 10 and 11) was the motivation of a RL algorithm by Nachum et al. (2017), *path consistency learning* (PCL). The main idea of (soft) PCL is to learn a parameterized policy and value function by minimizing the following objective function:  $\mathcal{J}(\theta, \phi) = \frac{1}{2} \sum_{\xi} J(\xi_i, \theta, \phi)^2$ , where  $\xi = (x_0, a_0, r_0, \dots, x_{d-1}, a_{d-1}, r_{d-1}, x_d)$  is any  $d$ -length sub-trajectory,  $\theta$  and  $\phi$  are the policy and value function parameters, respectively, and

$$J(\xi, \theta, \phi) = -V_{\phi}(x_0) + \gamma^d V_{\phi}(x_d) + \sum_{t=0}^{d-1} \gamma^t (r(x_t, a_t) - \alpha \log \mu_{\theta}(a_t|x_t)). \quad (12)$$

<sup>3</sup>When  $\alpha = O(1/|\mathcal{A}|)$  in soft MDP, theoretically it has a bounded performance when  $|\mathcal{A}| \gg 1$ , but this would potentially lead to unstable learning as the magnitude of  $Q(\cdot, \cdot)/\alpha$  blows up.

An important property of the soft PCL algorithm is that since the multi-step consistency (11) holds for any  $d$ -length sub-trajectory, it can use both on-policy ( $\xi$ 's generated by the current policy  $\mu_{\theta}$ ) and off-policy data, i.e.,  $\xi$ 's generated by a policy different than the current one, including any  $d$ -length sub-trajectory from the replay buffer.

Note that since both optimal policy  $\mu_{\text{sf}}^*$  and value function  $V_{\text{sf}}^*$  can be written based on the optimal action-value function  $Q_{\text{sf}}^*$  (see Eq. 4), we may write the objective function (12) based on  $Q_{\psi}$ , and optimize only one set of parameters  $\psi$ , instead of separate  $\theta$  and  $\phi$ .

#### 5. Consistency between Optimal Value & Policy in Sparse MDPs

This section begins the main contributions of our work. We first identify a (one-step) consistency equation for the sparse MDPs defined by (3). We then prove the relationship between the *sparse consistency* equation and the optimal policy and value function of the sparse MDP, and highlight its similarities and differences with that in soft MDPs, discussed in Section 4. We then extend the sparse consistency equation to multiple steps and prove results that allow us to use the *multi-step sparse consistency* equation to derive on-policy and off-policy algorithms to solve sparse MDPs, which we fully describe in Section 6. The significance of the sparse consistency equation is in providing an efficient tool for computing a *near-optimal* policy for sparse MDPs, which only involves solving a set of linear equations and linear complementary constraints, as opposed to (iteratively) solving the fixed-point of the non-linear sparse Bellman operator (8). We report the proofs of all the theorems of this section in Appendix A.

For any policy  $\mu$  and value function  $V : \mathcal{X} \rightarrow \mathbb{R}$ , we define the (one-step) consistency equation of the sparse MDPs as, for all state  $x \in \mathcal{X}$  and for all actions  $a \in \mathcal{A}$ ,

$$V(x) = r(x, a) + \frac{\alpha}{2} - \alpha \mu(a|x) + \lambda(a|x) - \Lambda(x) + \gamma \sum_{x'} P(x'|x, a) V(x'), \quad (13)$$

where  $\lambda : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}_+$  and  $\Lambda : \mathcal{X} \rightarrow \mathbb{R}_-$  are Lagrange multipliers, such that  $\lambda(a|x) \cdot \mu(a|x) = 0$  and  $-\frac{\alpha}{2} \leq \Lambda(x) \leq 0$ . We call this the *one-step sparse consistency* equation and is equivalent to Eq. 10 in soft MDPs.

We now present a theorem which states that, similar to soft MDPs, optimality in sparse MDPs is a necessary condition for consistency, i.e., *optimality implies consistency*.

**Theorem 1.** *The optimal policy  $\mu_{\text{sp}}^*$  and value function  $V_{\text{sp}}^*$  of the sparse MDP (3) satisfy the consistency equation (13).*

Theorem 2 shows that in the sparse MDPs, *consistency only implies near-optimality*, as opposed to optimality in the case of soft MDPs.

**Theorem 2.** Any policy  $\mu$  that satisfies the consistency equation (13) is  $\alpha/(1-\gamma)$ -optimal in the sparse MDP (3), i.e., for each state  $x \in \mathcal{X}$ , we have  $V_{sp}^\mu(x) \geq V_{sp}^*(x) - \frac{\alpha}{1-\gamma}$ .

This result indicates that for a fixed  $\gamma$ , as  $\alpha$  decreases, a policy  $\mu$  satisfying the one-step consistency equations approaches the true optimal  $\mu_{sp}^*$ . To connect the performance of  $\mu$  to the original goal of maximizing expected return, we present the following corollary, which is a direct consequence of Theorem 2 and the results reported in Section 3.2 on the performance of  $\mu_{sp}^*$  in the original MDP.

**Corollary 1.** Any policy  $\mu$  that satisfies the consistency equation (13) is  $(\frac{3}{2} - \frac{1}{2|\mathcal{A}|}) \cdot \frac{\alpha}{1-\gamma}$ -optimal in the original MDP (1), i.e., for each state  $x \in \mathcal{X}$ , we have  $V^*(x) - (\frac{3}{2} - \frac{1}{2|\mathcal{A}|}) \cdot \frac{\alpha}{1-\gamma} \leq V^\mu(x) \leq V^*(x)$ .

We now extend the (one-step) sparse consistency equation (13) to multiple steps. For any state  $x_0 \in \mathcal{X}$  and sequence of actions  $a_0, \dots, a_{d-1}$ , define the multi-step consistency equation for sparse MDPs as

$$V(x_0) = \mathbb{E}_{x_{1:d}|x_0, a_{0:d-1}} \left[ \gamma^d V(x_d) + \sum_{t=0}^{d-1} \gamma^t \left( r(x_t, a_t) + \frac{\alpha}{2} - \alpha \mu(a_t|x_t) + \lambda(a_t|x_t) - \Lambda(x_t) \right) \right], \quad (14)$$

where  $\lambda : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}_+$  and  $\Lambda : \mathcal{X} \rightarrow \mathbb{R}_-$  are Lagrange multipliers, such that  $\lambda(a|x) \cdot \mu(a|x) = 0$  and  $-\frac{\alpha}{2} \leq \Lambda(x) \leq 0$ . We call this *multi-step sparse consistency equation*, the equivalent of Eq. 11 in soft MDPs.

From Theorem 1, we can immediately show that multi-step sparse consistency is a necessary condition of optimality.

**Corollary 2.** The optimal policy  $\mu_{sp}^*$  and value function  $V_{sp}^*$  of the sparse MDP (3) satisfy the multi-step consistency equation (14).

Conversely, followed from Theorem 2, we prove the following result on the performance of any policy satisfying the multi-step consistency equation. This is a novel result showing that solving the multi-step consistency equation is indeed *sufficient* to guarantee near-optimality.

**Corollary 3.** Any policy  $\mu$  that satisfies the consistency equation (14) is  $\alpha/(1-\gamma)$ -optimal in sparse MDP (3).

Equipped with the above results on the relationship between (near)-optimality and multi-step consistency in sparse MDPs, we are now ready to present our off-policy RL algorithms to solve the sparse MDP (3).

## 6. Path Consistency Learning in Sparse MDPs

Similar to the PCL algorithm for soft MDPs, in sparse MDPs the multi-step consistency equation (14) naturally leads to a path-wise algorithm for training a policy  $\mu_\theta$  and

value function  $V_\phi$  parameterized by  $\theta$  and  $\phi$ , as well as Lagrange multipliers  $\Lambda_\rho$  and  $\lambda_{\theta,\rho}$  parameterized by the auxiliary parameter  $\rho$ . To characterize the objective function of this algorithm, we first define the *soft consistency error* for the  $d$ -step sub-trajectory  $\xi$  as a function of  $\theta$ ,  $\rho$ , and  $\phi$ ,

$$J(\xi; \theta, \rho, \phi) = -V_\phi(x_0) + \gamma^d V_\phi(x_d) + \sum_{t=0}^{d-1} \gamma^t \left( r(x_t, a_t) + \frac{\alpha}{2} - \alpha \mu_\theta(a_t|x_t) + \lambda_{\theta,\rho}(a_t|x_t) - \Lambda_\rho(x_t) \right).$$

The goal of our algorithm is to learn  $V_\phi$ ,  $\mu_\theta$ ,  $\lambda_{\theta,\rho}$ , and  $\Lambda_\rho$ , such that the expectation of  $J(\xi; \theta, \rho, \phi)$  for any initial state  $x_0$  and action sequence  $a_{0:d-1}$  is as close to 0 as possible. Our sparse PCL algorithm minimizes the empirical objective function  $\mathcal{J}_n(\theta, \rho, \phi) = \frac{1}{2} \sum_{\xi_i} J(\xi_i; \theta, \rho, \phi)^2$ , which converges to  $\mathcal{J}(\theta, \rho, \phi) = \mathbb{E}_{x_0, a_{0:d-1}} [\mathbb{E}[J(\xi; \theta, \rho, \phi)^2 | x_0, a_{0:d-1}]]$  as  $i \rightarrow \infty$ . By the Cauchy-Schwarz inequality,  $\mathcal{J}(\theta, \rho, \phi)$  is a conservative surrogate of  $\mathbb{E}[J(\xi; \theta, \rho, \phi)]^2$ , which represents the error of the multi-step consistency equation. This relationship justifies that the solution policy of the sparse PCL algorithm is near-optimal (see Corollary 5). Moreover, the gradient of  $J(\xi)$  w.r.t. the parameters is as follows:

$$\frac{\partial J(\xi)}{\partial \theta} = J(\xi; \theta, \rho, \phi) \sum_{t=0}^{d-1} \gamma^t \nabla_\theta (\lambda_{\theta,\rho}(a_t|x_t) - \alpha \mu_\theta(a_t|x_t)),$$

$$\frac{\partial J(\xi)}{\partial \rho} = J(\xi; \theta, \rho, \phi) \sum_{t=0}^{d-1} \gamma^t \nabla_\rho (\lambda_{\theta,\rho}(a_t|x_t) - \Lambda_\rho(x_t)),$$

$$\frac{\partial J(\xi)}{\partial \phi} = J(\xi; \theta, \rho, \phi) \nabla_\phi (V_\phi(x_0) - \gamma^d V_\phi(x_d)).$$

We may relate the sparse PCL algorithm to the standard actor-critic (AC) algorithm (Konda & Tsitsiklis, 2000; Sutton et al., 2000), where  $\partial J(\xi)/\partial \theta$  and  $\partial J(\xi)/\partial \phi$  correspond to the actor and critic updates, respectively. An advantage of sparse PCL over the standard AC is that it does not need the multi-time-scale update for convergence.

While optimizing  $J(\theta, \rho, \phi)$  minimizes the mean square of the soft consistency error, in order to satisfy the multi-step consistency in (14), one still needs to impose the following constraints on Lagrange multipliers into the optimization problem: **(i)**  $-\frac{\alpha}{2} \leq \Lambda_\rho(x) \leq 0$ , and **(ii)**  $\lambda_{\theta,\rho}(a|x) \cdot \mu_\theta(a|x) = 0, \forall x \in \mathcal{X}, \forall a \in \mathcal{A}$ . One standard approach is to replace the above constraints with adding penalty functions (Bertsekas, 1999) to the original objective function  $\mathcal{J}_n$ . Note that each penalty function is associated with a penalty parameter and there are  $|\mathcal{X}| \cdot |\mathcal{A}| + 2|\mathcal{X}|$  constraints. When  $|\mathcal{X}|$  and  $|\mathcal{A}|$  are large, tuning all the parameters becomes computationally expensive. Another approach is to update the penalty parameters using gradient ascent methods (Bertsekas, 2014). This is equivalent to finding the saddle point of the Lagrangian function in the constrained optimization problem. However, the challenge is to balance the primal and dual updates in practice.

We hereby describe an alternative and a much simpler methodology to parameterize the Lagrange multipliers  $\lambda_{\theta,\rho}(a|x)$  and  $\Lambda_\rho(x)$ , such that the aforementioned constraints are immediately satisfied. Although this method may impose extra restrictions to the representations of their function approximations, it avoids the difficulties of directly solving a constrained optimization problem. Specifically, to satisfy the constraint (i), one can parameterize  $\Lambda_\rho$  with a multilayer perceptron network that has either an activation function of  $-\alpha/2 \cdot \sigma(\cdot)$  or  $-\alpha/2 \cdot (1 + \tanh(\cdot))/2$  at its last layer. To satisfy constraint (ii), we consider the case when  $\mu_\theta$  is written in form of  $(f_\theta(x, a))^+$  for some function approximator  $f_\theta$ . This parameterization of  $\mu_\theta$  is justified by the closed-form solution policy of the Tsallis entropy-regularized MDP problem in (7). Specifically, (7) uses  $f_{\text{sp}}^*(x, a) = Q_{\text{sp}}^*(x, a)/\alpha - \mathcal{G}(Q_{\text{sp}}^*(x, \cdot)/\alpha)$ . Now suppose that  $\lambda_{\theta,\rho}$  is parameterized as follows:  $\lambda_{\theta,\rho}(a|x) = (-f_\theta(x, a))^+ \cdot F_\rho(x, a)$ , where  $F_\rho : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}^+$  is an auxiliary function approximator. Then by the property  $(x)^+ \cdot (-x)^+ = 0$ , constraint (ii) is immediately satisfied. A pseudo-code of our sparse PCL algorithm can be found in Algorithm 1 in the Appendix A.

**Unified Sparse PCL** Note that the closed-form optimal policy  $\mu_{\text{sp}}^*$  and value function  $V_{\text{sp}}^*$  are both functions of the optimal state-action value function  $Q_{\text{sp}}^*$ . As in soft PCL, based on this observation one can also parameterize both policy and value function in sparse PCL (see Eq. 7) with a single function approximator  $Q_\psi(x, a)$ . Although consistency does not imply optimality in sparse MDPs (as opposed to the case of soft MDPs), the justification of this parameterization comes from Corollary 4, where the unique optimal value function and optimal policy satisfy the consistency equation (14). From an actor-critic perspective, the significance of this is that both policy (*actor*) and value function (*critic*) can be updated simultaneously without affecting the convergence. Accordingly, the update rule for the model parameter  $\psi$  takes the form

$$\frac{\partial J(\xi)}{\partial \psi} = J(\xi; \psi, \rho) \left( \sum_{t=0}^{d-1} \gamma^t \nabla_\psi (\lambda_{\psi,\rho}(a_t|x_t) - \alpha \mu_\psi(a_t|x_t)) + \nabla_\psi V_\psi(x_0) - \gamma^d \nabla_\psi V_\psi(x_d) \right).$$

**The Saddle Point Reformulation** While sparse PCL presents an effective policy learning methodology, unless the immediate reward and transition function are deterministic, in general the gradient derived by minimizing the empirical mean-square consistency error  $\mathcal{J}_n(\theta, \rho, \phi)$  is biased w.r.t. the original consistency error  $\mathbb{E} [J(\xi; \theta, \rho, \phi)]^2$ . If the variance of  $\mathcal{J}_n$  is large, solving for a policy by minimizing this function could lead to an inaccurate solution. To resolve this issue, inspired by Dai et al. (2017); Liu et al. (2015), we propose to reformulate the PCL objective into an equivalent saddle-point prob-

lem by exploiting the Fenchel conjugate of the square function, i.e.,  $x^2 = \max_\nu (2\nu x - \nu^2)$  and by applying the interchangeability principle (Proposition 6.37 in Shapiro et al. 2009). Specifically, using the above properties one can easily show that the problem of minimizing  $\mathbb{E} [J(\xi; \theta, \rho, \phi)]^2$  is equivalent to the saddle-point problem  $\min_{\theta,\rho,\phi} \max_{\nu \in \mathcal{N}} \mathbb{E}_\xi [2\nu(x_0, a_{0:d-1}) \cdot J(\xi; \theta, \rho, \phi) - \nu(x_0, a_{0:d-1})^2]$ , where  $\mathcal{N}$  is the set of functions defined on the space  $\mathcal{X} \times \mathcal{A}^d \rightarrow \mathbb{R}$ . This means that the corresponding stochastic gradient is *unbiased*. Furthermore, using completion of squares and defining  $\tilde{\nu}(x_0, a_{0:d-1}) = \nu(x_0, a_{0:d-1}) + V_\phi(x_0)$ , the objective function of the saddle-point problem can be rewritten as  $\min_{\theta,\rho,\phi} \mathbb{E}_\xi [J(\xi; \theta, \rho, \phi)^2] - \min_{\tilde{\nu}} \mathbb{E}_\xi [(\tilde{\nu}(x_0, a_{0:d-1}) - \tilde{J}(\xi; \theta, \rho, \phi))^2]$ , where  $\tilde{J}(\xi; \theta, \rho, \phi) = \gamma^d V_\phi(x_d) + \sum_{t=0}^{d-1} \gamma^j (r(x_t, a_t) + \frac{\alpha}{2} - \alpha \mu_\theta(a_t|x_t) + \lambda_{\theta,\rho}(a_t|x_t) - \Lambda_\rho(x_t))$  is a  $d$ -step auxiliary cumulative return. In this formulation, the first term is the same as the original objective function of sparse PCL, while the second term cancels extra variance coming from  $\tilde{J}(\xi; \theta, \rho, \phi)$ . Leveraging this property, we can propose an alternative saddle point algorithm for sparse PCL, which we call sparse saddle point PCL (Algorithm 2 in Appendix B). Notice that the above problem is *not* a standard convex-concave saddle point problem, i.e., the objective is concave in  $\nu$  for fixed  $(\theta, \rho, \phi)$ , but it is not necessarily always convex in  $(\theta, \rho, \phi)$  for fixed  $\nu$ . Yet, utilizing recent results from saddle point optimization, one can still derive convergence analysis for this algorithm (see Appendix B for more details).

## 7. Experimental Results

We demonstrate the effectiveness of the sparse PCL algorithm by comparing its performance with that of the soft PCL algorithm on a number of RL environments available in the OpenAI Gym environment (Brockman et al., 2016).

### 7.1. Discrete Control

Here we compare the performance of these two algorithms on the following standard algorithmic tasks: **1)** Copy, **2)** DuplicatedInput, **3)** RepeatCopy, **4)** Reverse, and **5)** ReversedAddition (see Appendix C for more details). Each task can be viewed as a grid environment, where each cell stores a single character from a finite vocabulary  $\mathcal{V}$ . An agent moves on the grid of the environment and writes to output. At each time step the agent observes the character of the single cell in which it is located. After observing the character, the agent must take an action of the form  $(m, w, c)$ , where  $m$  determines the agent’s move to an adjacent cell, (in 1D environments,  $m \in \{\text{left}, \text{right}\}$ ; in 2D environments,  $m \in \{\text{left}, \text{right}, \text{up}, \text{down}\}$ ),  $w \in \{0, 1\}$  determines whether the agent writes to output or not, and  $c \in \mathcal{V}$  determines the character that the agent writes if  $w = 1$  (otherwise  $c = \emptyset$ ). Based on this problem setting,

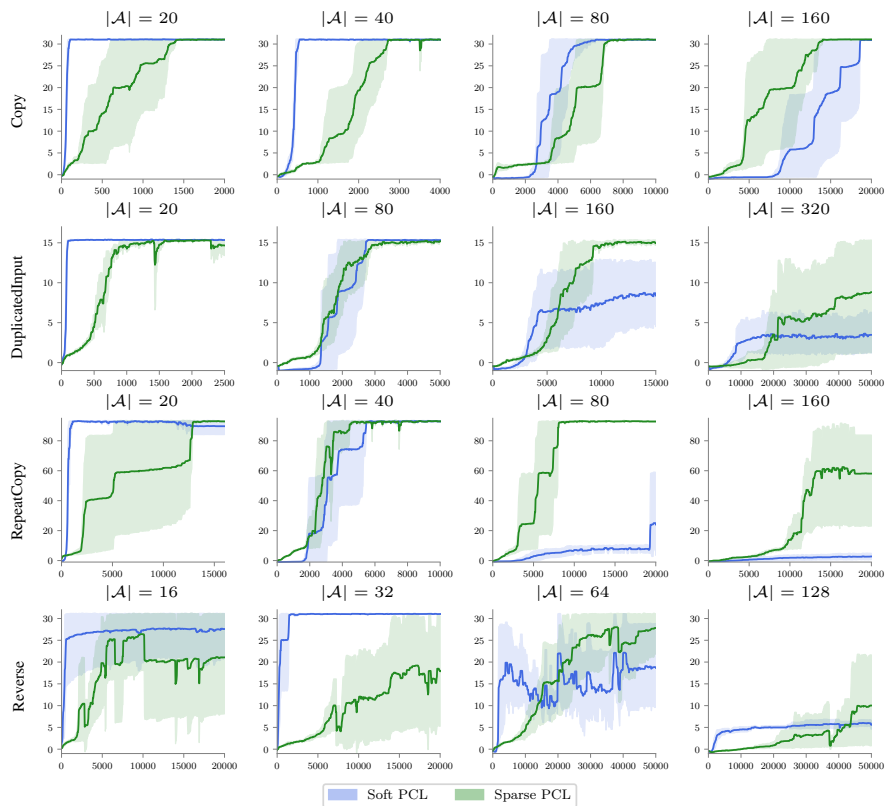


Figure 1. Results of the average reward from sparse PCL and standard soft PCL during training. Each row corresponds to a specific algorithmic task. For each particular task, the action space is increased from left to right across the rows, corresponding to an increase in difficulty. We observe that soft PCL returns a better solution when the action space is small, but its performance degrades quickly as the size of the action space grows. On the other hand, sparse PCL is not only able to learn good policies in tasks with small actions, but unlike soft PCL, also successfully learns high-reward policies in the higher-dimension variants (see Appendix C for additional results).

the action space  $\mathcal{A}$  has size  $|\mathcal{A}| = \Theta(|\mathcal{V}|)$ . Accordingly, the difficulty of these tasks grows with the size of the vocabulary. To illustrate the effectiveness of Tsallis entropy-regularized MDPs in problems with large action space, we evaluate these two PCL algorithms on 4 different choices of  $|\mathcal{V}|$ .

In each task, the agent has a different goal. In Copy, the environment is a 1D sequence of characters and the agent aims to copy the sequence to output. In DuplicatedInput, the environment is a 1D sequence of duplicated characters and the agent needs to write the de-duplicated sequence to output. In RepeatCopy, the environment is a 1D sequence of characters in which the agent must copy in forward order, reverse order, and finally forward order again. In Reverse, the environment is a 1D sequence of characters in which the agent must copy to output in reverse order. Finally, in ReversedAddition, the environment is a  $2n$  grid of digits representing two numbers in base- $|\mathcal{V}|$  that the agent needs to sum. In each task the agent receives a reward of 1 for each correctly output character. The episode is terminated either when the task is completed or when the agent

outputs an incorrect character.

We follow a similar experimental procedure as in Nachum et al. (2017), where the functions  $V$ ,  $\mu$ ,  $\lambda$ , and  $\Lambda$  in the consistency equations are parameterized with a recurrent neural network with multiple heads. For each task and each PCL algorithm, we perform a hyper-parameter search to find the optimal regularization weight  $\alpha$ , and the corresponding training curves for average reward are shown in Figure 1. To increase the statistical significance of these experiments, we also train these policies on 5 different Monte Carlo trials (Notice that these environments are inherently deterministic, and thus, no additional Monte Carlo evaluation is needed.). Details of the experimental setup and extra numerical results are included in Appendix C.

For each task we evaluated sparse PCL compared to the original soft PCL on a suite of variants which successively increase the vocabulary size. For low vocabulary sizes soft PCL achieves better results. This suggests that Shannon entropy encourages better exploration in small action spaces. Indeed, in such regimes, a greater proportion of the total ac-

tions are useful to explore and exploration is not as costly. Therefore, the decreased exploration of the Tsallis entropy may outweigh its asymptotic benefits. The sub-optimality bounds presented in this paper support this behavior: when  $|\mathcal{A}|$  is small,  $\alpha_{\text{soft PCL}} \log(|\mathcal{A}|) \leq \frac{3}{2} \alpha_{\text{sparse PCL}}$ .

As we increase the vocabulary size and as a result the action space, the picture changes. We see that the advantage of soft PCL over sparse PCL decreases until eventually the order is reversed and sparse PCL begins to show a significant improvement over the standard soft PCL. This supports our original hypothesis. In large action spaces, the tendency of soft PCL to assign a non-zero probability to many sub-optimal actions over-emphasizes exploration and is detrimental to the final reward performance. On the other hand, sparse PCL is able to handle exploration in large action spaces properly. These empirical results provide evidence for this unique advantage of sparse PCL.

## 7.2. Continuous Control

We further evaluate the two PCL algorithms on HalfCheetah, a continuous control problem in the OpenAI gym. The environment consists of a 6-dimensional action space, where each dimension corresponds to a torque of  $[-1, 1]$ . Here we discretize each continuous action with either one of the following grids:  $\{-1, 0, 1\}$  and  $\{-1, -0.5, 0, 0.5, 1\}$ . Even though the resolution of these discretization grids is coarse, the corresponding action spaces are quite large, with sizes of  $3^6 = 729$  and  $5^6 = 15625$ , respectively.

We present the results of sparse PCL and soft PCL on these discretized problems in Figure 2. Similar to the observations in the algorithmic tasks, here the policy learned by sparse PCL performs much better than that of soft PCL. Specifically sparse PCL achieves higher average reward and is able to learn much faster. To better visualize the learning progress of these two PCL algorithms in these problems, at each training step we also compare the average probability of the most-likely actions across all time-steps from the on-policy trajectory. Specifically in each iteration we collect a single on-policy trajectory of 1000 steps. Therefore, this metric is an average over 1000 samples of (greedy) action probabilities. Clearly, sparse PCL quickly converges to a near-deterministic policy, while the policy generated by soft PCL still allocates significant probability masses to non-optimal actions (as the average probability of most-likely actions barely ever exceeds 0.75). In environments like HalfCheetah, where the trajectory has a long horizon (1000 steps), the soft-max policy will in general suffer because it chooses a large number of sub-optimal actions in each episode for exploration.

Comparing with the performance of other continuous RL algorithms, such as deterministic policy gradient (DPG)

(Silver et al., 2014), we found that the policy generated by sparse PCL is sub-optimal. This is mainly due to the coarse discretization of the action space. Our main purpose here is to demonstrate the fast and improved convergence to deterministic policies in sparse PCL, compared to soft PCL. Further evaluation of sparse PCL will be left to future work.

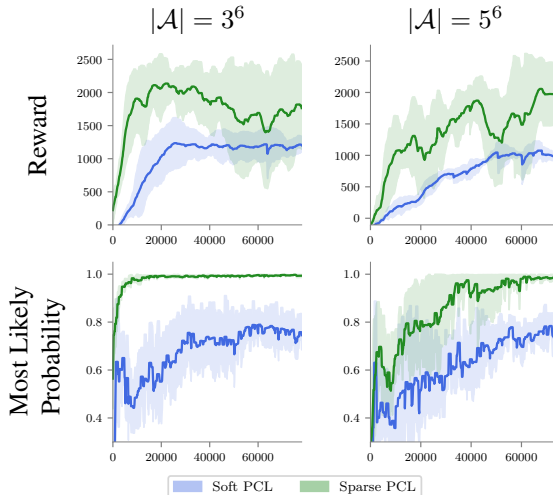


Figure 2. Results of sparse PCL and soft PCL in HalfCheetah with discretized actions. The top figure shows the average reward over 5 random runs during training, with best hyper-parameters. On the bottom we plot the average probability of the most-likely actions during training. The bottom figure illustrates the fast convergence of sparse PCL to a near-deterministic policy.

## 8. Conclusions

In this work we studied the sparse entropy-regularized problem in RL, whose optimal policy has non-zero probability for only a small number of actions. Similar to the work by Nachum et al. (2017), we derived a relationship between (near-)optimality and consistency for this problem. Furthermore, by leveraging the properties of the consistency equation, we proposed a class of sparse path consistency learning (sparse PCL) algorithms that are applicable to both on-policy and off-policy data and can learn from multi-step trajectories. We found that the theoretical advantages of sparse PCL correspond to empirical advantages as well. For tasks with a large number of actions, we find significant improvement in final performance and amount of time needed to reach that performance by using sparse PCL compared to the original soft PCL.

Future work includes **1)** extending the sparse PCL algorithm to the more general class of Tsallis entropy, **2)** investigating the possibility of combining sparse PCL and path following algorithms such as TRPO (Schulman et al., 2015), and **3)** comparing the performance of sparse PCL with other deterministic policy gradient algorithms, such as DPG (Silver et al., 2014) in the continuous domain.



## References

- Asadi, K. and Littman, M. An alternative softmax operator for reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning*, pp. 243–252, 2017.
- Azar, M., Gómez, V., and Kappen, H. Dynamic policy programming. *Journal of Machine Learning Research*, 13:3207–3245, 2012.
- Bellman, R. *Dynamic Programming*. Princeton University Press, 1957.
- Belousov, B. and Peters, J. F-divergence constrained policy improvement. *arXiv preprint arXiv:1801.00056*, 2017.
- Bertsekas, D. *Nonlinear programming*. Athena scientific Belmont, 1999.
- Bertsekas, D. *Constrained optimization and Lagrange multiplier methods*. Academic press, 2014.
- Bertsekas, D. and Tsitsiklis, J. *Neuro-Dynamic Programming*. Athena Scientific, 1996.
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. OpenAI Gym. *arXiv:1606.01540*, 2016.
- Dai, B., Shaw, A., Li, L., Xiao, L., He, N., Chen, J., and Song, L. Smoothed dual embedding control. *arXiv preprint arXiv:1712.10285*, 2017.
- Farahmand, A. M., Ghavamzadeh, M., Szepesvári, Cs., and Mannor, S. Regularized policy iteration. In *Proceedings of Advances in Neural Information Processing Systems 21*, pp. 441–448. MIT Press, 2008.
- Farahmand, A. M., Ghavamzadeh, M., Szepesvári, Cs., and Mannor, S. Regularized fitted Q-iteration for planning in continuous-space Markovian decision problems. In *Proceedings of the American Control Conference*, pp. 725–730, 2009.
- Fox, R., Pakman, A., and Tishby, N. G-learning: Taming the noise in reinforcement learning via soft update. In *Proceedings of the 32nd International Conference on Uncertainty in Artificial Intelligence*, pp. 202–211, 2016.
- Ghadimi, S. and Lan, G. Stochastic first and zeroth order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23(4):2341–2368, 2013.
- Ghavamzadeh, M., Lazaric, A., Munos, R., and Hoffman, M. Finite-sample analysis of lasso-TD. In *Proceedings of the Twenty-Eighth International Conference on Machine Learning*, pp. 1177–1184, 2011.
- Johns, J., Painter-Wakefield, C., and Parr, R. Linear complementarity for regularized policy evaluation and improvement. In *Proceedings of Advances in Neural Information Processing Systems 23*, pp. 1009–1017. MIT Press, 2010.
- Kappen, H. Path integrals and symmetry breaking for optimal control theory. *Journal of Statistical Mechanics*, 11, 2005.
- Kolter, Z. and Ng, A. Regularization and feature selection in least-squares temporal difference learning. In *Proceedings of the Twenty-Sixth International Conference on Machine Learning*, pp. 521–528, 2009.
- Konda, V. and Tsitsiklis, J. Actor-critic algorithms. In *Advances in neural information processing systems*, pp. 1008–1014, 2000.
- Lee, K., Choi, S., and Oh, S. Sparse Markov decision processes with causal sparse Tsallis entropy regularization for reinforcement learning. *IEEE Robotics and Automation Letters*, 2018.
- Liu, B., Liu, J., Ghavamzadeh, M., Mahadevan, S., and Petrik, M. Finite-sample analysis of proximal gradient td algorithms. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence*, pp. 504–513, 2015.
- Milgrom, P. and Segal, I. Envelope theorems for arbitrary choice sets. *Econometrica*, 70(2):583–601, 2002.
- Mnih, V., Badia, A., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., and Kavukcuoglu, K. Asynchronous methods for deep reinforcement learning. In *Proceedings of the 33rd International Conference on Machine Learning*, pp. 1928–1937, 2016.
- Nachum, O., Norouzi, M., Xu, K., and Schuurmans, D. Bridging the gap between value and policy based reinforcement learning. In *NIPS*, pp. 2772–2782, 2017.
- Nachum, O., Norouzi, M., Xu, K., and Schuurmans, D. Trust-PCL: An off-policy trust region method for continuous control. In *Proceedings of the 5th International Conference on Learning Representations*, 2018.
- Neu, G., Jonsson, A., and Gómez, V. A unified view of entropy-regularized Markov decision processes. *arXiv:1705.07798*, 2017.
- O’Donoghue, B., Munos, R., Kavukcuoglu, K., and Mnih, V. PGQ: Combining policy gradient and Q-learning. In *Proceedings of the 5th International Conference on Learning Representations*, 2017.

- Peters, J., Müling, K., and Altun, Y. Relative entropy policy search. In *Proceedings of the 24th Conference on Artificial Intelligence*, pp. 1607–1612, 2010.
- Puterman, M. *Markov Decision Processes*. Wiley Interscience, 1994.
- Schulman, J., Levine, S., Moritz, P., Jordan, M., and Abbeel, P. Trust region policy optimization. In *Proceedings of the 32nd International Conference on Machine Learning*, pp. 1889–1897, 2015.
- Shapiro, A., Dentcheva, D., and Ruszczyński, A. *Lectures on stochastic programming: modeling and theory*. SIAM, 2009.
- Silver, D., Lever, G., Heess, N., Degris, T., Wierstra, D., and Riedmiller, M. Deterministic policy gradient algorithms. In *ICML*, 2014.
- Sutton, R. and Barto, A. *An Introduction to Reinforcement Learning*. MIT Press, 1998.
- Sutton, R., McAllester, D., Singh, S., and Mansour, Y. Policy gradient methods for reinforcement learning with function approximation. In *Proceedings of Advances in Neural Information Processing Systems 12*, pp. 1057–1063, 2000.
- Todorov, E. Linearly-solvable Markov decision problems. In *Proceedings of the 19th Advances in Neural Information Processing*, pp. 1369–1376, 2006.
- Todorov, E. Policy gradients in linearly-solvable MDPs. In *Proceedings of the 23rd Advances in Neural Information Processing*, pp. 2298–2306, 2010.
- Tsallis, C. Possible generalization of Boltzmann-Gibbs statistics. *Journal of Statistical Physics*, 52(1):479–487, 1988.
- Ziebart, B. *Modeling Purposeful Adaptive Behavior with the Principle of Maximum Causal Entropy*. PhD thesis, Carnegie Mellon University, 2010.