
Exploration Conscious Reinforcement Learning Revisited

Lior Shani^{*1} Yonathan Efroni^{*1} Shie Mannor¹

Abstract

The Exploration-Exploitation tradeoff arises in Reinforcement Learning when one cannot tell if a policy is optimal. Then, there is a constant need to explore new actions instead of exploiting past experience. In practice, it is common to resolve the tradeoff by using a fixed exploration mechanism, such as ϵ -greedy exploration or by adding Gaussian noise, while still trying to learn an optimal policy. In this work, we take a different approach and study exploration-conscious criteria, that result in optimal policies with respect to the exploration mechanism. Solving these criteria, as we establish, amounts to solving a surrogate Markov Decision Process. We continue and analyze properties of exploration-conscious optimal policies and characterize two general approaches to solve such criteria. Building on the approaches, we apply simple changes in existing tabular and deep Reinforcement Learning algorithms and empirically demonstrate superior performance relatively to their non-exploration-conscious counterparts, both for discrete and continuous action spaces.

1. Introduction

The main goal of Reinforcement Learning (RL) (Sutton et al., 1998) is to find an optimal policy for a given decision problem. A major difficulty arises due to the Exploration-Exploitation tradeoff, which characterizes the omnipresent tension between exploring new actions and exploiting the so-far acquired knowledge. Considerable line of work has been devoted for dealing with this tradeoff. Algorithms that explicitly balance between exploration and exploitation were developed for tabular RL (Kearns & Singh, 2002; Brافman & Tennenholtz, 2002; Jaksch et al., 2010; Osband et al., 2013). However, generalizing these results to approximate

RL, i.e. when using function approximation, remains an open problem. On the practical side, recent works combined more advanced exploration schemes in approximate RL (e.g. Bellemare et al. (2016); Fortunato et al. (2017)), inspired by the theory of tabular RL. Nonetheless, even in the presence of more advanced mechanisms, ϵ -greedy exploration is still applied (Bellemare et al., 2017; Dabney et al., 2018; Osband et al., 2016). More generally, the traditional and simpler ϵ -greedy scheme (Sutton et al., 1998; Asadi & Littman, 2016) in discrete RL, and Gaussian action noise in continuous RL, are still very useful and popular in practice (Mnih et al., 2015; 2016; Silver et al., 2014; Schulman et al., 2017; Horgan et al., 2018), especially due to their simplicity.

These types of exploration schemes share common properties. First, they all fix some exploration parameter beforehand, e.g. ϵ , the ‘inverse temperature’ β , or the action variance σ for the ϵ -greedy, soft-max and Gaussian exploration schemes, respectively. By doing so, the balance between exploration and exploitation is set. Second, they all explore using a random policy, and exploit using current estimate of the *optimal policy*. In this work, we follow a different approach, when using these fixed exploration schemes: exploiting by using an estimate of the optimal policy w.r.t. the *exploration mechanism*.

Exploration-Consciousness is the main reason for the improved performance of on-policy methods like Sarsa and Expected-Sarsa (Van Seijen et al., 2009) over Q-learning during training (Sutton et al., 1998)[Example 6.6: Cliff Walking]. Imagine a simple Cliff-Walking problem: The goal of the agent is to reach the end without falling of the cliff, where the optimal policy is to go alongside the cliff. While using a fixed-exploration scheme, playing a near optimal policy which goes alongside the cliff will lead to a significant sub-optimal performance. This, in turn, will hurt the acquisition of new experience needed to learn the optimal policy. However, learning to act optimally w.r.t. the exploration scheme can mitigate this difficulty; the agent learns to reach the goal while keeping a safe enough distance from the cliff.

In the past, tabular q-learning-like exploration-conscious algorithms were suggested (John, 1994; Littman et al., 1997; Van Seijen et al., 2009). Here we take a different approach, and focus on exploration conscious *policies*. The main contributions of this work are as follows:

^{*}Equal contribution ¹Department of Electrical Engineering, Technion, Haifa, Israel. Correspondence to: Lior Shani <shanlior@gmail.com>, Yonathan Efroni <jonathan.efroni@gmail.com>.

- We define exploration-consciousness optimization criteria, for discrete and continuous actions spaces. The criteria are interpreted as finding an optimal policy within a restricted set of policies. Both, we show, can be reduced to solving a surrogate MDP. The surrogate MDP approach, to the best of our knowledge, is a new one, and serves us repeatedly in this work.
- We formalize a bias-error sensitivity tradeoff. The solutions are biased w.r.t. the optimal policy, yet, are less sensitive to approximation errors.
- We establish two fundamental approaches to practically solve Exploration-Conscious optimization problems. Based on these, we formulate algorithms in discrete and continuous action spaces, and empirically test the algorithms on the Atari and MuJoCo domains.

2. Preliminaries

Our framework is the infinite-horizon discounted Markov Decision Process (MDP). An MDP is defined as the 5-tuple $(\mathcal{S}, \mathcal{A}, P, R, \gamma)$ (Puterman, 1994), where \mathcal{S} is a finite state space, \mathcal{A} is a compact space, $P \equiv P(s'|s, a)$ is a transition kernel, $R \equiv r(s, a) \in [0, R_{\max}]$ is a bounded reward function, and $\gamma \in [0, 1)$. Let $\pi : \mathcal{S} \rightarrow \mathcal{P}(\mathcal{A})$ be a stationary policy, where $\mathcal{P}(\mathcal{A})$ is a probability distribution on \mathcal{A} , and denote Π as the set of deterministic policies, $\pi \in \Pi : \mathcal{S} \rightarrow \mathcal{A}$. Let $v^\pi \in \mathbb{R}^{|\mathcal{S}|}$ be the value of a policy π , defined in state s as $v^\pi(s) \equiv \mathbb{E}_s^\pi[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t)]$, where $a_t \sim \pi(s_t)$, and \mathbb{E}_s^π denotes expectation w.r.t. the distribution induced by π and conditioned on the event $\{s_0 = s\}$. It is known that $v^\pi = \sum_{t=0}^{\infty} \gamma^t (P^\pi)^t r^\pi = (I - \gamma P^\pi)^{-1} r^\pi$, with the component-wise values $[P^\pi]_{s,s'} \triangleq \mathbb{E}_{a \sim \pi}[P(s' | s, a)]$ and $[r^\pi]_s \triangleq \mathbb{E}_{a \sim \pi}[r(s, a)]$. Furthermore, the q -function of π is given by $q^\pi(s, a) = r(s, a) + \gamma \sum_{s'} P(s' | s, a) v^\pi(s')$, and represents the value of taking an action a from state s and then using the policy π .

Usually, the goal is to find π^* yielding the optimal value, $\pi^* \in \arg \max_{\pi \in \Pi} \mathbb{E}^\pi[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t)]$, and the optimal value is $v^* = v^{\pi^*}$. It is known that optimal deterministic policy always exists (Puterman, 1994). To achieve this goal the following classical operators are defined (with equalities holding component-wise). $\forall v, \pi$:

$$T^\pi v = r^\pi + \gamma P^\pi v, \quad T v = \max_{\pi} T^\pi v, \quad (1)$$

$$\mathcal{G}(v) = \{\pi : T^\pi v = T v\}, \quad (2)$$

where T^π is a linear operator, T is the optimal Bellman operator and both T^π and T are γ -contraction mappings w.r.t. the max norm. It is known that the unique fixed points of T^π and T are v^π and v^* , respectively. $\mathcal{G}(v)$ is the standard set of 1-step greedy policies w.r.t. v . Furthermore, given v^* , the set $\mathcal{G}(v^*)$ coincides with that of stationary optimal

policies. It is also useful to define the q -optimal Bellman operator, which is a γ -contraction, with fixed point q^* .

$$T^q q(s, a) = r(s, a) + \gamma \sum_{s'} P(s' | s, a) \max_{a'} q(s', a'), \quad (3)$$

In this work, the use of *mixture policies* is abundant. We denote the $\alpha \in [0, 1]$ -convex mixture of policies π_1, π_2 by $\pi^\alpha(\pi_1, \pi_2) \triangleq (1 - \alpha)\pi_1 + \alpha\pi_2$. Importantly, $\pi^\alpha(\pi_1, \pi_2)$ can be interpreted as a stochastic policy s.t with w.p $(1 - \alpha)$ the agent acts with π_1 and w.p α acts with π_2 .

3. The α -optimal criterion

In this section, we define the notion of α -optimal policy w.r.t. a policy, π_0 . We then claim that finding an α -optimal policy can be done by solving a *surrogate* MDP. We continue by defining the surrogate MDP, and analyze some basic properties of the α -optimal policy.

Let $\alpha \in [0, 1]$. We define π_{α, π_0}^* to be the α -optimal policy w.r.t. π_0 , and is contained in the following set,

$$\pi_{\alpha, \pi_0}^* \in \arg \max_{\pi' \in \Pi} \mathbb{E}^{\pi^\alpha(\pi', \pi_0)} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \right], \quad (4)$$

or, $\pi_{\alpha, \pi_0}^* \in \arg \max_{\pi'} v^{\pi^\alpha(\pi', \pi_0)}$, where $a_t \sim \pi^\alpha(\pi', \pi_0)$ and $\pi^\alpha(\pi', \pi_0)$ is the α -convex mixture of π' and π_0 , and thus a probability distribution. For brevity, we omit the subscript π_0 , and denote the α -optimal policy by π_α^* throughout the rest of the paper. The α -optimal value (w.r.t. π_0) is $v^{\pi^\alpha(\pi_\alpha^*, \pi_0)}$, the value of the policy $\pi_\alpha^*(\pi_\alpha^*, \pi_0)$. In the following, we will see the problem is equivalent to solving a surrogate MDP, for which an optimal deterministic policy is known to exist. Thus, there is no loss optimizing over the set of deterministic policies Π .

Optimization problem (4) can be viewed as optimizing over a restricted set of policies: all policies that are a convex combination of π_0 with a fixed α . Naturally, we can consider in (4) a state-dependent $\alpha(s)$ as well, and some of the results in this work will consider this scenario. In other words, π_α^* is the best policy an agent can act with, if it plays w.p $(1 - \alpha)$ according to π_α^* , and w.p α according to π_0 , where π_0 can be any policy. The relation to the ϵ -greedy exploration setup becomes clear when π_0 is a uniform distribution on the actions, and set $\alpha = \epsilon$ instead of α . Then, π_α^* is optimal w.r.t. the ϵ -greedy exploration scheme; the policy would have the largest accumulated reward, relatively to all other policies, when acting in an ϵ -greedy fashion w.r.t. it.

We choose to name the policy as the α - and not ϵ -optimal to prevent confusion with other frameworks. The ϵ -optimal policy is a notation used in the context of PAC-MDP type of analysis (Strehl et al., 2009), and has a different meaning than the objective in this work (4).

3.1. The α -optimal Bellman operator, α -optimal policy and policy improvement

In the previous section, we defined the α -optimal policy and the α -optimal value, π_α^* and $v^{\pi_\alpha^*}(\pi_\alpha^*, \pi_0)$, respectively. We start this section by observing that problem (4) can be viewed as solving a *surrogate MDP*, denoted by \mathcal{M}_α . We define the Bellman operators of the surrogate MDP, and use them to prove an important improvement property.

Define the surrogate MDP as $\mathcal{M}_\alpha = (\mathcal{S}, \mathcal{A}, P_\alpha, R_\alpha, \gamma)$.

$$\begin{aligned} \forall a \in \mathcal{A}, r_\alpha(s, a) &= (1 - \alpha)r(s, a) + \alpha r^{\pi_0}(s), \\ P_\alpha^\pi(s' | s, a) &= (1 - \alpha)P(s' | s, a) + \alpha P^{\pi_0}(s' | s), \end{aligned} \quad (5)$$

are its reward and dynamics, and rest of its ingredients are similar to \mathcal{M} . We denote the value of a policy π on \mathcal{M}_α by v_α^π , and the optimal value on \mathcal{M}_α by v_α^* . The following simple lemma relates the value of a policy π , measured on \mathcal{M} and \mathcal{M}_α (see proof in Appendix D).

Lemma 1. *For any policy π , $v_\alpha^\pi = v^{\pi_\alpha^*}(\pi, \pi_0)$. Thus, an optimal policy on \mathcal{M}_α is the α -optimal policy π_α^* (4).*

The fixed-policy and optimal Bellman operators of \mathcal{M}_α are denoted by T_α^π and T_α , respectively. Again, for brevity we omit π_0 from the definitions. Notice that T_α^π and T_α are γ -contractions as being Bellman operators of a γ -discounted MDP. The following Lemma relates T_α^π and T_α to the Bellman operators of the original MDP, \mathcal{M} . Furthermore, it stresses a non-trivial relation between the α -optimal policy π_α^* and the α -optimal value, $v^{\pi_\alpha^*}(\pi_\alpha^*, \pi_0)$.

Proposition 2. *The following claims hold for any policy π :*

1. $T_\alpha^\pi = (1 - \alpha)T^\pi + \alpha T^{\pi_0}$, with fixed point $v_\alpha^\pi = v^{\pi_\alpha^*}(\pi, \pi_0)$.
2. $T_\alpha = (1 - \alpha)T + \alpha T^{\pi_0}$, with fixed point $v_\alpha^* = v^{\pi_\alpha^*}(\pi_\alpha^*, \pi_0)$.
3. An α -optimal policy is an optimal policy of \mathcal{M}_α and is greedy w.r.t. v_α^* , $\pi_\alpha^* \in \mathcal{G}(v_\alpha^*) = \{\pi' : T^{\pi'} v_\alpha^* = T v_\alpha^*\}$.

In previous works, e.g. (Asadi & Littman, 2016), the operator $(1 - \epsilon)T + \epsilon T^{\pi_0}$ was referred to as the ϵ -greedy operator. Lemma 2 shows this operator is T_α (with $\alpha = \epsilon$), the optimal Bellman operator of the defined surrogate MDP \mathcal{M}_α . This lemma leads to the following important property.

Proposition 3. *Let $\alpha \in [0, 1)$, $\beta \in [0, \alpha]$, π_0 be a policy, and π_α^* be the α -optimal policy w.r.t π_0 . Then, $v^{\pi_0} \leq v^{\pi_\alpha^*}(\pi_\alpha^*, \pi_0) \leq v^{\pi_\beta}(\pi_\alpha^*, \pi_0)$, with equality iff $v^{\pi_0} = v^*$.*

The first relation $v^{\pi_0} \leq v^{\pi_\alpha^*}(\pi_\alpha^*, \pi_0)$, $\pi_\alpha^*(\pi_\alpha^*, \pi_0)$ is better than π_0 , is trivial and holds by definition (4). The non-trivial statement is the second one. It asserts that given π_α^* , it is worthwhile to use the mixture policy $\pi_\beta(\pi_\alpha^*, \pi_0)$ with $\beta < \alpha$; use π_0 with smaller probability. Specifically, better performance, compared to $\pi_\alpha^*(\pi_\alpha^*, \pi_0)$, is assured when using the deterministic policy π_α^* , by setting $\beta = 0$.

In section 6, we demonstrate the empirical consequences of the improvement lemma, which, to our knowledge, has not yet been stated. Furthermore, the improvement lemma is unique to the defined optimization criterion (4). We will show that alternative definitions of exploration conscious criteria does not necessarily have this property. Moreover, one can use Proposition 3 to generalize the notion of the 1-step greedy policy (2), as was done in Efroni et al. (2018) with multiple-step greedy improvement. We leave studying this generalization and its Policy Iteration scheme for future work, and focus on solving (4) a single time.

3.2. Performance bounds in the presence of approximations

We now consider an approximate setting and quantify a bias - error sensitivity tradeoff in $\pi_\alpha^*(\hat{\pi}_\alpha^*, \pi_0)$, where $\hat{\pi}_\alpha^*$ is an approximated α -optimal policy. We formalize an intuitive argument; as α increases the bias relatively to the optimal policy increases. Yet, the sensitivity to errors decreases, since the agent uses π_0 w.p. α regardless of errors.

Definition 1. *Let v^* be the optimal value of an MDP, \mathcal{M} . We define $L(s) \triangleq v^*(s) - T^{\pi_0} v^*(s) \geq 0$, to be the Lipschitz constant w.r.t. π_0 of the MDP at state s . We further define the upper bound on the Lipschitz constant $L \triangleq \max_s L(s)$.*

Definition 1 defines the ‘Lipschitz’ property of the optimal value, v^* . Intuitively, $L(s)$ quantifies a degree of ‘smoothness’ of the *optimal value*. A small value of $L(s)$ indicates that if one acts according to π_0 once and then continue playing the optimal policy from state s , a great loss will not occur. Large values of $L(s)$ indicate that using π_0 from state s leads to an irreparable outcome (e.g, falling off a cliff). The following theorem formalizes a bias-error sensitivity tradeoff. As α increases, the bias increases, while the sensitivity to errors decreases (see proof in Appendix H).

Theorem 4. *Let $\alpha \in [0, 1]$. Assume \hat{v}_α^* is an approximate α -optimal value s.t $\|v_\alpha^* - \hat{v}_\alpha^*\| = \delta$ for some $\delta \geq 0$. Let $\hat{\pi}_\alpha^*$ be the greedy policy w.r.t. \hat{v}_α^* , $\hat{\pi}_\alpha^* \in \mathcal{G}(\hat{v}_\alpha^*)$. Then, the performance relatively to the optimal policy is bounded by,*

$$\|v^* - v^{\pi_\alpha^*}(\hat{\pi}_\alpha^*, \pi_0)\| \leq \underbrace{\frac{\alpha L}{1 - \gamma}}_{\text{Bias}} + \underbrace{\frac{2(1 - \alpha)\gamma\delta}{1 - \gamma}}_{\text{Sensitivity}}.$$

When the bias of the α -optimal value relatively to the optimal one is small, solving (4) does not lead to a great loss relatively to the optimal performance. The bias can be bounded by the ‘Lipschitz’ property L of the MDP. For a state dependent $\alpha(s)$, the bias bound changes to be dependent on $\max_s \alpha(s)L(s)$. This highlights the importance of prior knowledge when using (4). Choosing π_0 (possibly state-wise) s.t. $\max_s \alpha(s)L(s)$ is small, allows to use a bigger α , while the bias is small. The sensitivity term upper bounds the performance of $\pi_\alpha^*(\hat{\pi}_\alpha^*, \pi_0)$ relatively to the α -optimal value, and is less sensitive to errors as α increase.

The bias term is derived by using the structure of \mathcal{M}_α , and is not a direct application of the Simulation Lemma (Kearns & Singh, 2002; Strehl et al., 2009); applying it would lead to a bias of $\frac{\alpha R_{\max}}{(1-\gamma)^2}$. For the sensitivity term, we generalize (Bertsekas & Tsitsiklis, 1995)[Proposition 6.1] (see Appendix G). There, a $(1 - \alpha)$ factor does not exist.

4. Exploration-Conscious Continuous Control

The α -greedy approach from Section 3 relies on an exploration mechanism which is fixed beforehand: π_0 and α are fixed, and an optimal policy w.r.t. them is being calculated (4). However, in continuous control RL algorithms, such as DDPG and PPO (Lillicrap et al., 2015; Schulman et al., 2017), different approach is used. Usually, a policy is being learned, and the exploration noise is injected by perturbing the policy, e.g., by adding to it a Gaussian noise.

We start this section by defining an exploration-conscious optimality criterion that captures such perturbation for the simple case of Gaussian noise. Then, results from Section 3 are adapted to the newly defined criterion, while highlighting commonalities and differences relatively to (4). As in Section 3, we define an appropriate surrogate MDP and we show it can be solved by the usual machinery of Bellman operators. Unlike Section 3, we show that improvement when decreasing the stochasticity does not generally hold. Finally, we prove a similar bias-error sensitivity result: As σ grows, the bias increases, but the sensitivity term decreases.

Instead of restricting the set of policies to the one defined in (4), we restrict our set of policies to be the set of Gaussian policies with a fixed σ^2 variance. Formally, we wish to find the optimal deterministic policy $\mu_\sigma^* : \mathcal{S} \rightarrow \mathcal{A}$ in this set,

$$\mu_\sigma^* \in \arg \max_{\mu \in \Pi} \mathbb{E}^{\pi_{\mu, \sigma}} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \right], \quad (6)$$

where $\pi_{\mu, \sigma}(\cdot | s) = \mathcal{N}(\mu(s), \sigma^2)$, is a Gaussian policy with mean $\mu(s)$ and a fixed variance σ^2 . We name μ_σ^* and π_σ^* as the mean and σ -optimal policy, respectively. As in (4), we show in the following that solving (6) is equivalent for solving a surrogate MDP. Thus, optimal policy can always be found in the deterministic class of policies Π ; mixture of Gaussians would not lead to a better performance in (6).

Similarly to (5), we define a surrogate MDP \mathcal{M}_σ w.r.t. to the Gaussian noise and relate it to values of Gaussian policies on the original MDP \mathcal{M} . Then, we characterize its Bellman operators and thus establish it can be solved using Dynamic Programming. Define the surrogate MDP as $\mathcal{M}_\sigma = (\mathcal{S}, \mathcal{A}, P_\sigma, R_\sigma, \gamma)$. For every $a \in \mathcal{A}$,

$$\begin{aligned} r_\sigma(s, a) &= \int_{\mathcal{A}} \mathcal{N}(a'; a, \sigma) r(s, a') da', \\ P_\sigma(s' | s, a) &= \int_{\mathcal{A}} \mathcal{N}(a'; a, \sigma) P(s' | s, a') da', \end{aligned} \quad (7)$$

are its reward and dynamics, and denote a value of a policy on \mathcal{M}_σ by v_σ^μ . The following results correspond to Lemma 1 and Proposition 2 for the class of Gaussian policies.

Lemma 5. *For any policy π , $v_{\mu, \sigma}^\pi = v_\sigma^\mu$. Thus, an optimal policy on \mathcal{M}_σ is the mean optimal policy μ_σ^* (6).*

Proposition 6. *Let π be a mixture of Gaussian policies. Then, the following holds:*

1. $T_\sigma^\mu = \mathbb{E}^{\pi \sim \pi_{\mu, \sigma}} T^\pi$, with fixed point $v_\sigma^\mu = v^{\pi_{\mu, \sigma}}$.
2. $T_\sigma = \max_{\mu \in \mathcal{A}} \mathbb{E}^{\pi \sim \pi_{\mu, \sigma}} T^\pi$, with fixed point $v_\sigma^* = v^{\pi_{\mu_\sigma^*, \sigma}}$.
3. The mean σ -optimal policy μ_σ^* is an optimal policy of \mathcal{M}_σ and, $\mu_\sigma^* \in \{\mu : T^{\pi_{\mu, \sigma}} v_\sigma^* = \max_{\mu} T^{\pi_{\mu, \sigma}} v_\sigma^*\}$.

Surprisingly, given a σ -optimal policy mean μ_σ^* , an improvement is not assured when lowering the stochasticity by decreasing σ in $\pi_{\mu_\sigma^*, \sigma}$. This comes in contrast to Proposition 3 and highlights its uniqueness (proof in Appendix J).

Proposition 7. *Let $0 \leq \sigma' < \sigma$ and let μ_σ^* be the mean σ -optimal policy. There exists an MDP s.t $v^{\pi_{\mu_\sigma^*, \sigma'}} \not\leq v^{\pi_{\mu_\sigma^*, \sigma}}$.*

Definition 2. *Let \mathcal{M} be a continuous action space MDP. Assume that exists $L_r, L_p \geq 0$, s.t. $\forall s \in \mathcal{S}, \forall a_1, a_2 \in \mathcal{A}, |r(s, a_1) - r(s, a_2)| \leq L_r \|a_1 - a_2\|_1$ and $\|p(\cdot | s, a_1) - p(\cdot | s, a_2)\|_{TV} \leq L_p \|a_1 - a_2\|_1$. The Lipschitz constant of \mathcal{M} is $\mathcal{L} \triangleq (1 - \gamma)L_r + \gamma L_p R_{\max}$.*

The following theorem quantifies a bias-error sensitivity tradeoff in σ , similarly to Theorem 4 (see Appendix K).

Theorem 8. *Let \mathcal{M} be an MDP with Lipschitz constant \mathcal{L} and let $\sigma \in \mathbb{R}_+^{|\mathcal{A}|}$. Let v_σ^* be the σ -optimal value of \mathcal{M}_σ . Let \hat{v}_σ^* be an approximation of v_σ^* s.t. $\|v_\sigma^* - \hat{v}_\sigma^*\| = \delta$ for $\delta \geq 0$. Let $\mu_\sigma^*, \hat{\mu}_\sigma^* \in \mathbb{R}^{\mathcal{A}}$ be the greedy mean policy w.r.t. v_σ^* and \hat{v}_σ^* respectively. Let $\|\cdot\|_{\sigma^{-2}}$ is the σ^{-2} -weighted euclidean norm. Then,*

$$\|v^* - v^{\pi_{\mu_\sigma^*, \sigma}}\| \leq \underbrace{\frac{\mathcal{L} \|\sigma\|_1}{2(1-\gamma)^2}}_{\text{Bias}} + \underbrace{\frac{\gamma \delta \min\{\frac{1}{2} \|\mu_\sigma^* - \hat{\mu}_\sigma^*\|_{\sigma^{-2}}, 2\}}{1-\gamma}}_{\text{Sensitivity}}.$$

5. Algorithms

In this section, we offer two fundamental approaches to solve exploration conscious criteria using sample-based algorithms: the *Expected* and *Surrogate* approaches. For both, we formulate converging, q-learning-like, algorithms. Next, by adapting DDPG, we show the two approaches can be used in exploration-conscious continuous control as well.

Consider any fixed exploration scheme. Generally, these schemes operate in two stages: (i) Choose a greedy action, a_{chosen} . (ii) Based on a_{chosen} and some randomness generator, choose an action to be applied on the environment, a_{env} . E.g., for ϵ -greedy exploration, w.p. $1 - \alpha$ the agent acts with

a_{chosen} , otherwise, with a random uniform policy. While in RL the common update rules use a_{env} , the saved experience is $(s, a_{\text{env}}, r, s')$, in the following we motivate the use of a_{chosen} , and view the data as $(s, a_{\text{chosen}}, a_{\text{env}}, r, s')$.

The two approaches characterized in the following are based on two, inequivalent, ways to define the q -function. For the *Expected* approach the q -function is defined as usual: $q^\pi(s, a)$ represents the value obtained when **taking an action** $a = a_{\text{env}}$ and then acting with π , meaning a is the action chosen in step (ii). Alternatively, for the *Surrogate* approach, the q -function is defined on the ‘Surrogate’ MDP, i.e., the exploration is viewed as stochasticity of the environment. Then, $q_\alpha^\pi(s, a)$ is the value obtained when a is the action of step (i), i.e., **choosing action** $a = a_{\text{chosen}}$.

5.1. Exploration Conscious Q-Learning

We focus on solving the α -optimal policy (4), and formulate q -learning-like algorithms using the two aforementioned approaches. The *Expected* α -optimal q -function is,

$$q^{\pi^\alpha(\pi_\alpha^*, \pi_0)}(s, a) \triangleq r(s, a) + \gamma \sum_{s'} P(s' | s, a) v_\alpha^*(s') \quad (8)$$

Indeed, $q^{\pi^\alpha(\pi_\alpha^*, \pi_0)}$ is the usually defined q -function of the policy $\pi^\alpha(\pi_\alpha^*, \pi_0)$ on an MDP \mathcal{M} . Here, the action a represents the actual performed action, a_{env} . By relating $q^{\pi^\alpha(\pi_\alpha^*, \pi_0)}$ to v_α^* it can be easily verified that $q^{\pi^\alpha(\pi_\alpha^*, \pi_0)}$ satisfies the fixed point equation (see Appendix L),

$$\begin{aligned} q^{\pi^\alpha(\pi_\alpha^*, \pi_0)}(s, a) = & \\ & r(s, a) + \gamma(1 - \alpha) \sum_{s'} P(s' | s, a) \max_{a'} q^{\pi^\alpha(\pi_\alpha^*, \pi_0)}(s', a') \\ & + \gamma\alpha \sum_{s', a'} P(s' | s, a) \pi_0(a' | s') q^{\pi^\alpha(\pi_\alpha^*, \pi_0)}(s', a'). \end{aligned} \quad (9)$$

Alternatively, consider the optimal q -function of the surrogate MDP \mathcal{M}_α (5). It satisfies the fixed-point equation

$$q_\alpha^*(s, a) \triangleq r_\alpha(s, a) + \gamma \sum_{s'} P_\alpha(s' | s, a) \max_{a'} q_\alpha^*(s', a').$$

The following lemma formalizes the relation between the two q -functions, and shows they are related by a function of the state, and not of the action.

Lemma 9. $q_\alpha^*(s, a) = (1 - \alpha)q^{\pi^\alpha(\pi_\alpha^*, \pi_0)}(s, a) + f(s)$.

The α -optimal policy π_α^* is also an optimal policy of \mathcal{M}_α (Lemma 1). Thus, it is greedy w.r.t. q_α^* , the optimal q of \mathcal{M}_α . By Proposition 2.3 it is also greedy w.r.t. $q^{\pi^\alpha(\pi_\alpha^*, \pi_0)}$, i.e.,

$$\pi_\alpha^*(s) \in \arg \max_{a'} q_\alpha^*(s, a') = \arg \max_{a'} q^{\pi^\alpha(\pi_\alpha^*, \pi_0)}(s, a').$$

Lemma 9 describes this fact by different means; the two q -functions are related by a function of the state and, thus, the

greedy action w.r.t. each is equal. Furthermore, it stresses the fact that the two q -function are not equal.

Before describing the algorithms, we define the following notation for any $q(s, a)$,

$$v(s) = \max_{a'} q(s, a'), \quad v^\pi(s) = \sum_{a'} \pi(a' | s) q(s, a').$$

We now describe the Expected α -Q-learning algorithm (see Algorithm 1), also given in (John, 1994; Littman et al., 1997), and re-interpret it in light of the previous discussion.

The fixed point equation (9), leads us to define the operator T_α^{Eq} for which $q^{\pi^\alpha(\pi_\alpha^*, \pi_0)} = T_\alpha^{Eq} q^{\pi^\alpha(\pi_\alpha^*, \pi_0)}$. Expected α -Q-learning (Alg. 1) is a Stochastic Approximation (SA) alg. based on the operator T_α^{Eq} . Given a sample of the form $(s, a_{\text{chosen}}, a_{\text{env}}, r, s')$, it updates $q(s, a_{\text{env}})$ by

$$(1 - \eta)q(s, a_{\text{env}}) + \eta(r_t + \gamma((1 - \alpha)v(s_{t+1}) + \alpha v^{\pi_0}(s_{t+1}))) \quad (10)$$

Algorithm 1 Expected α -Q-Learning

Initialize: $\alpha \in [0, 1]$, π_0 , q , learning rate η_t .

for $t = 0, 1, \dots$ **do**

$a_{\text{chosen}} \leftarrow \arg \max_a q_t(s_t, a)$

$X_t \sim \text{Bernoulli}(1 - \alpha)$

$a_{\text{env}} = \begin{cases} a_{\text{chosen}}, & \text{if } X_t = 1 \\ a \sim \pi_0(\cdot | s), & \text{if } X_t = 0 \end{cases}$

$r_t, s_{t+1} \leftarrow \text{ACT}(a_{\text{env}})$

$y_t \leftarrow r_t + \gamma(1 - \alpha)v_t(s_{t+1}) + \gamma\alpha v_t^{\pi_0}(s_{t+1})$

$q(s_t, a_{\text{env}}) \leftarrow (1 - \eta_t)q(s_t, a_{\text{env}}) + \eta_t y_t$

end for

return: $\pi \in \arg \max_a q(\cdot, a)$

Its convergence proof is standard and follows by showing T_α^{Eq} is a γ -contraction and using (Bertsekas & Tsitsiklis, 1995)[Proposition 4.4] (see proof in Appendix L.1).

We now turn to describe an alternative algorithm, which operates on the surrogate MDP, \mathcal{M}_α , and converges to q_α^* . Naively, given a sample $(s, a_{\text{chosen}}, r, s')$, regular q -learning on \mathcal{M}_α can be used by updating $q(s, a_{\text{chosen}})$ as,

$$(1 - \eta_t)q(s, a_{\text{chosen}}) + \eta_t(r_t + \gamma v(s_{t+1})), \quad (11)$$

Yet, this approach does not utilize a meaningful knowledge; when the exploration policy π_0 is played, i.e., when $X_t = 0$, the sample (r_t, s_{t+1}) can be used to update all the action entries from the current state. These entries are also affected by the policy π_0 . In fact, we cannot prove the convergence of the naive update based on current techniques; if the greedy action is repeatedly chosen, ‘infinitely often’ visit in all (s, a) pairs cannot be guaranteed.

This reasoning leads us to formulate Surrogate α -Q-learning (see Algorithm 2). The Surrogate α -Q-learning updates two q -functions, q and q_α . The first, q , has the same update

Algorithm 2 Surrogate α -Q-Learning

Initialize: $\alpha \in [0, 1]$, π_0 , q_α , q , learning rate η_t .

for $t = 0, 1, \dots$ **do**

$a_{\text{chosen}} \leftarrow \arg \max_a q(s_t, a)$

$X_t \sim \text{Bernoulli}(1 - \alpha)$

$a_{\text{env}} = \begin{cases} a_{\text{chosen}}, & \text{if } X_t = 1 \\ a \sim \pi_0(\cdot | s), & \text{if } X_t = 0 \end{cases}$

$r_t, s_{t+1} \leftarrow \text{ACT}(a_{\text{env}})$

for $\bar{a} \in \mathcal{A}$ **do**

$y_t^{\bar{a}} = \begin{cases} r_t + \gamma v_\alpha(s_{t+1}), & \bar{a} = a_{\text{chosen}} \\ X_t q(s_t, \bar{a}) + (1 - X_t)(r_t + \gamma v_\alpha(s_{t+1})), & \text{o.w.} \end{cases}$

$q_\alpha(s_t, \bar{a}) \leftarrow (1 - \eta) q_\alpha(s_t, \bar{a}) + \eta y_t^{\bar{a}}$

end for

$y_t \leftarrow r_t + \gamma(1 - \alpha)v(s_{t+1}) + \gamma\alpha v^{\pi_0}(s_{t+1})$

$q(s_t, a_{\text{env}}) \leftarrow (1 - \eta_t)q(s_t, a_{\text{env}}) + \eta_t y_t$

end for

return $\pi \in \arg \max_a q_\alpha(\cdot, a)$

as in Expected α -Q-learning, and thus converges (w.p 1) to $q^{\pi^\alpha(\pi_\alpha^*, \pi_0)}$. The second, q_α , updates the chosen greedy action using equation (11), when the exploration policy is not played ($X_t = 1$). By bootstrapping on q , the algorithm updates all other actions when the exploration policy π_0 is played ($X_t = 0$). Using (Singh et al., 2000)[Lemma 1], the convergence of Surrogate α -Q-learning to $(q^{\pi^\alpha(\pi_\alpha^*, \pi_0)}, q_\alpha^*)$ is established (see proof in Appendix L.2). Interestingly, and unlike other q -learning algorithms (e.g. Expected α -Q-learning, Q-learning, etc.), Surrogate α -Q-learning updates the entire action set given a single sample. For completeness, we state the convergence result for both algorithms.

Theorem 10. *Consider the processes described in Alg. 1, 2. Assume $\{\eta_t\}_{t=0}^\infty$ satisfies $\forall s \in \mathcal{S}, \forall a \in \mathcal{A}, \sum_{t=0}^\infty \eta_t = \infty$, and $\sum_{t=0}^\infty \eta_t^2 < \infty$, where $\eta_t \equiv \eta_t(s_t = s, a_{\text{env}, t} = a)$. Then, for both 1, 2 the sequence $\{q_n\}_{n=0}^\infty$ converges w.p. 1 to $q^{\pi^\alpha(\pi_\alpha^*, \pi_0)}$, and for 2, $\{q_{\alpha, n}\}_{n=0}^\infty$ converges w.p. 1 to q_α^* .*

5.2. Continuous Control

Building on the two approaches for solving Exploration Conscious criteria, we suggest two techniques to find an optimal Gaussian policy (6) using gradient based Deep RL (DRL) algorithms, and specifically, DDPG (Lillicrap et al., 2015). Nonetheless, the techniques are generalizable to other actor-critic, DRL algorithms (Schulman et al., 2017).

Assume we wish to find an optimal Gaussian policy by parameterizing its mean $\mu(\phi)$. Nachum et al. (2018)[Eq. 13] showed the gradient of the value w.r.t. ϕ is similar to Silver et al. (2014),

$$\nabla_\phi v^{\pi_{\mu, \sigma}} = \int_{\mathcal{S}} \partial_a q_\sigma^{\pi_{\mu, \sigma}}(s, a) \nabla_\phi \mu^\theta(s) d\rho^{\pi_{\mu, \sigma}}(s), \quad (12)$$

where $q_\sigma^\mu(s, a) = r_\sigma(s, a) + \gamma \int_{\mathcal{S}} p_\sigma(s' | s, a) v^{\pi_{\mu, \sigma}}(s') ds'$, is the q -function of the surrogate MDP. In light of previ-

ous section, we interpret q_σ^μ as the q -function of the surrogate MDP's \mathcal{M}_σ (7). Furthermore, we have the following relation between the surrogate and expected q -functions, $q_\sigma^\mu(s, a) = \int_{a' \in \mathcal{A}} \mathcal{N}(a' | a, \sigma) q^{\pi_{\mu, \sigma}}(s, a') da'$, from which it is easy to verify that (see Appendix L.3),

$$\nabla_a q_\sigma^{\pi_{\mu, \sigma}}(s, b) = \int_{\mathcal{A}} \mathcal{N}(b | a, \sigma) \nabla_b q^{\pi_{\mu, \sigma}}(s, b) db. \quad (13)$$

Thus, we can update the actor in two inequivalent ways, by using gradients on the surrogate MDP's q -function (12), or by using gradients of the expected q -function (13).

The updates of the critic, q_σ^μ or $q^{\pi_{\mu, \sigma}}$, can be done using the same notion that led to the two forms of updates in (11)-(10). When using Gaussian noise, one performs the two stages defined in Section 5, where a_{chosen} is the output of the actor $\mu(s)$, and $a_{\text{env}} \sim \mathcal{N}(a_{\text{chosen}}, \sigma)$. Then, the sample $(s, a_{\text{chosen}}, a_{\text{env}}, r, s')$ is obtained by interacting with the environment. Based on the the fixed policy TD-error defined in (11), we define the following loss function, for learning q_σ^μ , q -function of the fixed policy μ over \mathcal{M}_σ ,

$$(q_\sigma^\theta(s, a_{\text{chosen}}) - r - \gamma q_\sigma^{\theta-}(s', \mu^{\phi-}(s')))^2.$$

On the other hand, we can define a loss function derived from the fixed-policy TD-error defined in (10), for learning $q^{\pi_{\mu, \sigma}}$, the q -function of the Gaussian policy with mean and variance μ, σ^2 over \mathcal{M} ,

$$(q^\theta(s, a_{\text{env}}) - r - \gamma \int_{\mathcal{A}} \mathcal{N}(b | \mu^{\phi-}(s'), s') q^{\theta-}(s', b) db)^2.$$

6. Experiments

In this section, we test the theory and algorithms¹ suggested in this work. In all experiments we used $\gamma = 0.99$. The tested DRL algorithms in this section (See Appendix B) are simple variations of DDQN (Van Hasselt et al., 2016) and DDPG (Lillicrap et al., 2015), without any parameter tuning, and based on Section 5. For example, for the surrogate approach in both DDQN and DDPG we merely save $(s, a_{\text{chosen}}, r, s')$ instead of $(s, a_{\text{env}}, r, s')$ in the replay buffer (see Section 5 for definitions of $a_{\text{env}}, a_{\text{chosen}}$).

We observe a significant improved empirical performance, both in **training** and **evaluation** for both the surrogate and expected approaches relatively to the baseline performance. The improved training performance is predictable; the learned policy is optimal w.r.t. the noise which is being played. In large portion of the results, the exploration-conscious criteria leads to better performance in evaluation.

6.1. Exploration Consciousness with Prior Knowledge

We use an adaptation of the Cliff-Walking maze (Sutton et al., 1998) we term T-Cliff-Walking (see Appendix C).

¹Implementation of the proposed algorithms can be found in <https://github.com/shanlior/ExplorationConsciousRL>.




Figure 1. T-Cliff-Walking for the expected (E) and surrogate (S) approaches. (Left) $\alpha=0.3$. (Right) $\alpha(s)$ from prior knowledge.

The agent starts at the bottom-left side of a maze, and needs to get to the bottom-right side goal state with value $+1$. If the agent falls off the cliff, the episode terminates with reward -1 . When the agent visits any of the first three steps on top of the cliff, it gets a reward of $0.01 \cdot (1 - \gamma)$.

We tested Expected α -Q-learning, Surrogate α -Q-learning, and compared their performance to Q-learning in the presence of ϵ -greedy exploration. Figure 1 stresses the typical behaviour of the α -optimality criterion. It is easier to approximate $\pi^\alpha(\pi_\alpha^*, \pi_0)$ than the optimal policy. Further, by being exploration-consciousness, the value of the approximated policy improves faster using the α -optimal algorithms; it learns faster which regions to avoid. As Proposition 4 suggests, the value of the learned policy is biased w.r.t v^* . Next, as suggested by Proposition 3, acting greedily w.r.t. the approximated value attains better performance. Such improvement is not guaranteed while the value had not yet converged to v_α^* . However, the empirical results suggest that if the agent performs well over the mixture policy, it is worth using the greedy policy.

We show that it is possible to incorporate prior knowledge to decrease the bias caused by being Exploration-Conscious. The T-Cliff-Walking example demands high exploration, $\alpha = \epsilon = 0.3$, because of the bottleneck state between the two sides of the maze. The α -optimal policy in such case is to stay at the left part of the maze. We used the prior knowledge that $L(s)$ close to the barrier is high. The knowledge was injected through the choice of α , i.e., we chose a state-wise exploration scheme with $\alpha(s) = \epsilon(s) = 0.1$ in the passage and the two states around it, and $\alpha(s) = 0.3$ elsewhere, for all three algorithms. The results in Figure 1 suggests that using prior knowledge to set $\alpha(s)$, can increase the performance by reducing the bias. In contrast, such prior knowledge does not help the baseline q-learning.

6.2. Exploration Consciousness in Atari

We tested the α -optimal criterion in the more complex function approximation setting (see Appendix Alg. 3, 4). We used five Atari 2600 games (5) from the ALE (Bellemare

Table 1. Train and Test rewards for the Atari 2600 environment, with 90% confidence interval

| | Game | DDQN | Expected α -DDQN | Surrogate α -DDQN |
|-------|------------|---------------------|-------------------------|--------------------------|
| Train | Breakout | 350 \pm 4 | 356 \pm 6 | 357 \pm 4 |
| | FishingDer | -45 \pm 9 | -35 \pm 27 | -8 \pm 8 |
| | Frostbite | 1191 \pm 171 | 794 \pm 158 | 1908 \pm 162 |
| | Qbert | 13221 \pm 565 | 13431 \pm 178 | 14240 \pm 225 |
| | Riverraid | 8602 \pm 205 | 8811 \pm 645 | 11476 \pm 79 |
| Test | Breakout | 402 \pm 14 | 390 \pm 5 | 392 \pm 5 |
| | FishingDer | -37 \pm 15 | -19 \pm 34 | -3 \pm 19 |
| | Frostbite | 1720 \pm 191 | 1638 \pm 292 | 2686 \pm 278 |
| | Qbert | 15627 \pm 497 | 15780 \pm 206 | 16082 \pm 338 |
| | Riverraid | 9049 \pm 443 | 9491 \pm 802 | 12846 \pm 241 |

et al., 2013). We chose games that resemble the Cliff Walking scenario, where the wrong choice of action can lead to a sudden termination of the episode. Thus, being unaware of the exploration strategy can lead to poor training results. We used the same deep neural network as in DQN (Mnih et al., 2015), using the openAI Baselines implementation (Dhariwal et al., 2017), without *any parameter tuning*, except for the update equations. We chose to use the Double-DQN variant of DQN (Van Hasselt et al., 2016) for simplicity and generality. Nonetheless, changing the optimality criterion is orthogonal to any of the suggested add-ons to DQN (Hessel et al., 2017). We used $\alpha = \epsilon = 0.01$ in the train phase, and $\epsilon = 0.001$ in the evaluation phase. For the *surrogate* version, we used a naive implementation based on equation (11).

Table 1 shows that our method improves upon using the optimal criterion. That is, while bias exists, the algorithm still converges to a better policy. This result holds both on the exploratory training regime and the evaluation regime. Again, acting greedy w.r.t. the approximation of the α -optimal policy proved beneficial: The evaluation phase results surpasses the train phase results as shown in the table, and the training figures in Appendix (2). The evaluation is usually done with an $\epsilon = 0.001 > 0$. Proposition 3 put formal grounds for using smaller ϵ in the evaluation phase than in the training phase; improvement is assured. Being accurate is extremely important in most Atari games, so Exploration-Consciousness can also hurt the performance. Still, one can use prior knowledge to overcome this obstacle.

6.3. Exploration Consciousness in MuJoCo

We tested the Expected σ -DDPG (5) and Surrogate σ -DDPG (6) on continuous control tasks from the MuJoCo environment (Todorov et al., 2012). We used the OpenAI implementation of DDPG as the baseline, where we only changed the update equations to match our proposed algorithms. We used the default hyper-parameters, and independent Gaussian noise with $\sigma = 0.2$, for all tasks and

Table 2. Train and Test rewards for the MuJoCo environment.

| | Game | DDPG | Expected σ -DDPG | Surrogate σ -DDPG |
|-------|-------------|-----------------------|-------------------------|--------------------------|
| Train | Ant | 809 \pm 47 | 1013 \pm 49 | 993 \pm 110 |
| | HalfCheetah | 2255 \pm 804 | 2634 \pm 828 | 3848 \pm 248 |
| | Hopper | 1864 \pm 139 | 1866 \pm 132 | 2566 \pm 155 |
| | Humanoid | 1281 \pm 142 | 1416 \pm 155 | 1703 \pm 272 |
| | InPendulum | 694 \pm 109 | 882 \pm 33 | 998 \pm 3 |
| | Walker | 1722 \pm 170 | 2144 \pm 145 | 2587 \pm 214 |
| Test | Ant | 1611 \pm 120 | 1924 \pm 126 | 1754 \pm 184 |
| | HalfCheetah | 2729 \pm 936 | 3147 \pm 986 | 4579 \pm 298 |
| | Hopper | 3099 \pm 113 | 3071 \pm 50 | 3037 \pm 78 |
| | Humanoid | 1688 \pm 223 | 1994 \pm 389 | 2154 \pm 408 |
| | InPendulum | 999 \pm 2 | 1000 \pm 0 | 1000 \pm 0 |
| | Walker | 3031 \pm 298 | 3315 \pm 147 | 3501 \pm 240 |

algorithms. The results in Table 2 were averaged over 10 different seeds. The performance of the σ -optimal variants superseded the baseline DDPG, for most of the training and test results. Interestingly, although improvement is not guaranteed (Proposition 7), the σ -optimal policy improved when using μ^ϕ deterministically, i.e., in the test phase. This suggests that improvement can be expected on certain scenarios, although that generally it is not guaranteed. We also found that the training process was faster using the σ -optimal algorithms, as can be seen in the learning curves in Appendix 3. Interestingly, again, the surrogate approach proved superior.

7. Relation to existing work

Lately, several works have tackled the exploration problem for deep RL. In some, like Bootstrapped-DQN (see appendix [D.1] in (Osband et al., 2016)), the authors still employ an ϵ -greedy mechanism on top of their methods. Moreover, methods like Distributional-DQN (Bellemare et al., 2017; Dabney et al., 2018) and the state-of-the-art Ape-X DQN (Horgan et al., 2018), still uses ϵ -greedy and Gaussian noise, for discrete and continuous actions, respectively. Hence, all the above works are applicable for the α -optimal criterion by using the simple techniques described in Section 5.

Existing on-policy methods produce variants of Exploration-Consciousness. In TRPO and A3C (Schulman et al., 2015; Mnih et al., 2016), the exploration is *implicitly injected* into the agent policy through entropy regularization, and the agent improves upon the value of the explorative policy. Simple derivation shows the α -greedy and the Gaussian approaches are both equivalent to regularizing the entropy to be higher than a certain value by setting α or σ appropriately.

Expected α -Q-learning highlights a relation to algorithms analysed in (John, 1994; Littman et al., 1997) and to Expected-Sarsa (ES) (Van Seijen et al., 2009). The focus of (John, 1994; Littman et al., 1997) is exploration-conscious

q-based methods. In ES, when setting the ‘estimation policy’ (Van Seijen et al., 2009) to be $\pi = (1 - \alpha_t)\pi_G + \alpha_t\pi_0$, we get similar updating equations as in lines 1-1, and similarly to (John, 1994; Littman et al., 1997). However, in ES α_t decays to zero, and the *optimal policy* is obtained in the infinite time limit. In (Nachum et al., 2018), the authors offer a gradient based mechanism for updating the mean and variance of the actor. Here, we offer and analyze the approach of setting α_t and σ_t to a constant value. This would be of interest especially when a ‘good’ mechanism for decaying α_t and σ_t lacks; the decay mechanism is usually chosen by trial-and-error, and is not clear how it should be set.

Lastly, (4) and (6) can be understood as defining a ‘surrogate problem’, rather than finding an optimal policy. In this sense, it offers an alternative approach to biasing the problem by lowering the discount-factor, i.e., solve a surrogate MDP with $\bar{\gamma} < \gamma$ (Petrik & Scherrer, 2009; Jiang et al., 2015). Interestingly, the introduced bias when solving (4) is proportional to a *local property* of v^* , $L(s)$, that can be estimated using prior-knowledge on the MDP, where solving an MDP with $\bar{\gamma}$ introduces a bias proportional to a *non-local* term, which is harder to estimate. More importantly, the performance of an α -optimal policy π_α^* is assured to improve when tested on the original MDP \mathcal{M} (Proposition 3), while the performance of an optimal policy in an MDP with $\bar{\gamma}$ might decline when tested on \mathcal{M} with γ -discounting.

8. Summary

In this paper, we revisited the notion of an agent being conscious to an exploration process. To our view, this notion did not receive the proper attention, though it is implicitly and repeatedly used.

We started by formally defining *optimal policy* w.r.t. an exploration mechanism (4), (6). This expanded the view on exploration-conscious q-learning (John, 1994; Littman et al., 1997) to a more general one, and lead us to derive new algorithms, as well as re-interpreting existing ones (Van Seijen et al., 2009). We formulated the surrogate MDP notion, which helped us to establish that exploration-conscious criteria can be solved by Dynamic Programming, or, more generally, by an MDP solver. From the practical side, based on the theory, we tested DRL algorithms – by simply modifying existing ones, with no further hyper-parameter tuning – and empirically showed their superiority.

Although a bias - error sensitivity tradeoff was formulated, we did not prove (4), (6) are easier to solve than an MDP. We believe proving whether the claim is true is of interest. Furthermore, analyzing more exploration-conscious criteria, e.g., exploration-conscious w.r.t. Ornstein-Uhlenbeck noise, is of interest, as well as defining a unified framework for exploration-conscious criteria.

Acknowledgments

We would like to thank Chen Tessler, Nadav Merlis and Tom Zahavy for helpful discussions.

References

- Asadi, K. and Littman, M. L. An alternative softmax operator for reinforcement learning. *arXiv preprint arXiv:1612.05628*, 2016.
- Bellemare, M., Srinivasan, S., Ostrovski, G., Schaul, T., Saxton, D., and Munos, R. Unifying count-based exploration and intrinsic motivation. In *Advances in Neural Information Processing Systems*, pp. 1471–1479, 2016.
- Bellemare, M. G., Naddaf, Y., Veness, J., and Bowling, M. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47:253–279, 2013.
- Bellemare, M. G., Dabney, W., and Munos, R. A distributional perspective on reinforcement learning. *arXiv preprint arXiv:1707.06887*, 2017.
- Bertsekas, D. P. and Tsitsiklis, J. N. Neuro-dynamic programming: an overview. In *Decision and Control, 1995., Proceedings of the 34th IEEE Conference on*, volume 1, pp. 560–564. IEEE, 1995.
- Brafman, R. I. and Tennenholtz, M. R-max-a general polynomial time algorithm for near-optimal reinforcement learning. *Journal of Machine Learning Research*, 3(Oct): 213–231, 2002.
- Dabney, W., Rowland, M., Bellemare, M. G., and Munos, R. Distributional reinforcement learning with quantile regression. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Dhariwal, P., Hesse, C., Klimov, O., Nichol, A., Plappert, M., Radford, A., Schulman, J., Sidor, S., and Wu, Y. Openai baselines. <https://github.com/openai/baselines>, 2017.
- Efroni, Y., Dalal, G., Scherrer, B., and Mannor, S. Beyond the one-step greedy approach in reinforcement learning. In *Proceedings of the 35th International Conference on Machine Learning*, pp. 1386–1395, 2018.
- Fortunato, M., Azar, M. G., Piot, B., Menick, J., Osband, I., Graves, A., Mnih, V., Munos, R., Hassabis, D., Pietquin, O., et al. Noisy networks for exploration. *arXiv preprint arXiv:1706.10295*, 2017.
- Hessel, M., Modayil, J., Van Hasselt, H., Schaul, T., Ostrovski, G., Dabney, W., Horgan, D., Piot, B., Azar, M., and Silver, D. Rainbow: Combining improvements in deep reinforcement learning. *arXiv preprint arXiv:1710.02298*, 2017.
- Horgan, D., Quan, J., Budden, D., Barth-Maron, G., Hessel, M., Van Hasselt, H., and Silver, D. Distributed prioritized experience replay. *arXiv preprint arXiv:1803.00933*, 2018.
- Jaksch, T., Ortner, R., and Auer, P. Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research*, 11(Apr):1563–1600, 2010.
- Jiang, N., Kulesza, A., Singh, S., and Lewis, R. The dependence of effective planning horizon on model accuracy. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pp. 1181–1189. International Foundation for Autonomous Agents and Multiagent Systems, 2015.
- John, G. H. When the best move isn’t optimal: Q-learning with exploration. Citeseer, 1994.
- Kearns, M. and Singh, S. Near-optimal reinforcement learning in polynomial time. *Machine learning*, 49(2-3):209–232, 2002.
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- Littman, M. L. et al. Generalized markov decision processes: Dynamic-programming and reinforcement-learning algorithms. 1997.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529, 2015.
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., and Kavukcuoglu, K. Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning*, pp. 1928–1937, 2016.
- Nachum, O., Norouzi, M., Tucker, G., and Schuurmans, D. Smoothed action value functions for learning gaussian policies. *arXiv preprint arXiv:1803.02348*, 2018.
- Osband, I., Russo, D., and Van Roy, B. (more) efficient reinforcement learning via posterior sampling. In *Advances in Neural Information Processing Systems*, pp. 3003–3011, 2013.
- Osband, I., Blundell, C., Pritzel, A., and Van Roy, B. Deep exploration via bootstrapped dqn. In *Advances in neural information processing systems*, pp. 4026–4034, 2016.

- Petrik, M. and Scherrer, B. Biasing approximate dynamic programming with a lower discount factor. In *Advances in neural information processing systems*, pp. 1265–1272, 2009.
- Puterman, M. L. Markov decision processes. j. *Wiley and Sons*, 1994.
- Schulman, J., Levine, S., Abbeel, P., Jordan, M., and Moritz, P. Trust region policy optimization. In *International Conference on Machine Learning*, pp. 1889–1897, 2015.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Silver, D., Lever, G., Heess, N., Degris, T., Wierstra, D., and Riedmiller, M. Deterministic policy gradient algorithms. In *ICML*, 2014.
- Singh, S., Jaakkola, T., Littman, M. L., and Szepesvári, C. Convergence results for single-step on-policy reinforcement-learning algorithms. *Machine learning*, 38(3):287–308, 2000.
- Strehl, A. L., Li, L., and Littman, M. L. Reinforcement learning in finite mdps: Pac analysis. *Journal of Machine Learning Research*, 10(Nov):2413–2444, 2009.
- Sutton, R. S., Barto, A. G., et al. *Reinforcement learning: An introduction*. MIT press, 1998.
- Todorov, E., Erez, T., and Tassa, Y. Mujoco: A physics engine for model-based control. In *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, pp. 5026–5033. IEEE, 2012.
- Van Hasselt, H., Guez, A., and Silver, D. Deep reinforcement learning with double q-learning. In *AAAI*, volume 2, pp. 5. Phoenix, AZ, 2016.
- Van Seijen, H., Van Hasselt, H., Whiteson, S., and Wiering, M. A theoretical and empirical analysis of expected sarsa. In *Adaptive Dynamic Programming and Reinforcement Learning, 2009. ADPRL'09. IEEE Symposium on*, pp. 177–184. IEEE, 2009.