

EVAL: EigenVector-based Average-reward Learning

Jacob Adamczyk^{1,2}, Volodymyr Makarenko³, Stas Tiomkin⁴, Rahul V. Kulkarni^{1,2}

¹Department of Physics, University of Massachusetts Boston

²The NSF Institute for Artificial Intelligence and Fundamental Interactions

³Department of Computer Engineering, San José State University

⁴Department of Computer Science, Whitacre College of Engineering, Texas Tech University

jacob.adamczyk001@umb.edu, volodymyr.makarenko@sjsu.edu, stas.tiomkin@ttu.edu, rahul.kulkarni@umb.edu

Abstract

In reinforcement learning, two objective functions have been developed extensively in the literature: discounted and averaged rewards. The generalization to an entropy-regularized setting has led to improved robustness and exploration for both of these objectives. Recently, the entropy-regularized average-reward problem was addressed using tools from large deviation theory in the tabular setting. This method has the advantage of linearity, providing access to both the optimal policy and average reward-rate through properties of a single matrix. In this paper, we extend that framework to more general settings by developing approaches based on function approximation by neural networks. This formulation reveals new theoretical insights into the relationship between different objectives used in RL. Additionally, we combine our algorithm with a posterior policy iteration scheme, showing how our approach can also solve the average-reward RL problem without entropy-regularization. Using classic control benchmarks, we experimentally find that our method compares favorably with other algorithms in terms of stability and rate of convergence.

Introduction

To solve the central problem of reinforcement learning, an agent continuously and autonomously interacts with its surroundings to maximize a long-term reward signal. The standard method of solving reinforcement learning (RL) tasks involves a discounted objective function: the agent seeks to maximize an infinite discounted sum of rewards, as expected under a chosen control policy. This geometrically discounted sum ensures a convergent objective function, making it a convenient representation of the problem to be solved. The discounted RL literature is vast in both theoretical (Sutton and Barto 2018; Kakade 2003; Bertsekas 2012) and algorithmic (Schulman et al. 2016; Andrychowicz et al. 2021) studies and has demonstrated great success in real-world problems of interest (Mnih et al. 2015; Schulman et al. 2015, 2017; Hessel et al. 2018; Haarnoja et al. 2018b). However, in many RL problems, this use of discounting is simply a useful proxy for the solving the true objective function: maximization of total episode reward. As a result, the actual choice of the discount factor, γ , is unphysical, not

grounded in any corresponding physical timescale. As such, it is treated as a hyperparameter which is often tuned over for best performance (or set to some fixed value, e.g. $\gamma = 0.99$, for simplicity). There is a precedent in past work for rejecting the discounted framework (Naik et al. 2019), for example the Heaven and Hell MDP (Schwartz 1993), suggesting that in many tasks the use of discounting is not only unphysical, but can lead to disastrous outcomes. Thus in long-term (“continuing”) decision-making settings, another objective is needed.

An alternative approach to ensuring convergence of rewards across infinitely-long trajectories is to instead use the **average-reward** objective function (Mahadevan 1996). Despite offering a principled alternative to the long-term optimization problem, the average reward framework has not been historically popular in the RL literature, perhaps due to the lack of successful algorithms in this setting. However, recent work has developed new algorithms and theory (Zhang and Ross 2021; Ma et al. 2021; Saxena et al. 2023b; Arriojas et al. 2023b; Wan, Naik, and Sutton 2021; Naik et al. 2024) specifically for the average-reward objective. This prior work has focused on the tabular setting or on using policy-gradient techniques to develop new algorithms. Although these methods have proven useful in their respective domains, there remain gaps in the field that can be addressed with new approaches.

One such gap in the field of average-reward algorithms is the lack of entropy-regularized objectives, which have become a cornerstone of state-of-the-art algorithms in discounted RL (Haarnoja et al. 2017a, 2018b; Schulman et al. 2017, 2015). Formally, entropy-regularization involves including an “information” cost for deviating from a pre-specified (e.g. prior, guide, or behavioral) policy. This entropy cost is weighted with an inverse temperature parameter, β . This entropy-based regularization, in combination with treating the rewards as negative energies, reveals a close connection to statistical mechanics (Arriojas et al. 2023b; Rose, Mair, and Garrahan 2021). Furthermore, including an entropy regularization term in the RL objective leads to more robust solutions (non-greedy optimal policies) and has theoretically been shown to be more adaptable to changes in reward and transition dynamics (Eysenbach and Levine 2022). Preventing the optimal policy from collapsing to a deterministic function can ensure additional exploration occurs dur-

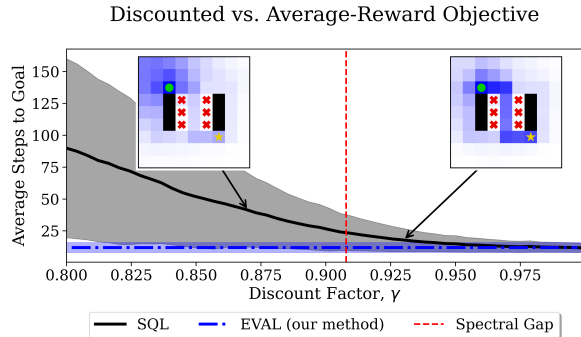


Figure 1: Performance of discounted soft Q-Learning (SQL) as a function of discount factor compared with solution using the proposed average-reward method (EVAL). Note that the average-reward solution (blue line) recovers the discounted solution as $\gamma \rightarrow 1$. For the discounted objective, computational cost grows as $(1 - \gamma)^{-1}$ and choosing a low discount factor to reduce computational cost can result in lower rewards. The boundary of low reward, low complexity and high reward, high complexity is demarcated by the discount factor derived from the spectral gap of the associated tilted matrix (cf. “Preliminaries”). Insets: state-occupation distributions following the SQL optimal policies at $\gamma = 0.87, 0.93$. Green dot denotes initial position of agent, and star denotes the goal. The agent can move in any of the cardinal directions. Since we use entropy-regularization, the optimal policy is stochastic, thus yielding a variance in the return (plotted with a shaded interval for each method). Inverse temperature $\beta = 15$.

ing training and deployment, often leading to faster learning (Ahmed et al. 2019).

In this paper, we introduce a novel algorithm for average-reward RL with entropy regularization: EigenVector-based Average-reward Learning (EVAL). Figure 1 illustrates a simple experiment highlighting one benefit of the EVAL approach to entropy-regularized RL compared to the discounted approach (Haarnoja et al. 2017b). When the RL problem is posed in the discounted framework, a discount factor γ is a required input parameter. However, there is often no principled approach for choosing the value of γ corresponding to the specific problem being addressed. Thus, the experimenter must treat γ as a hyperparameter. This reduces the choice of γ to a trade-off between large values to capture long-term rewards and small values to capture computational efficiency which typically scales polynomially with the horizon, $(1 - \gamma)^{-1}$ (Kakade 2003). The horizon introduces a natural timescale to the problem, but this timescale may not be well-aligned with the timescale corresponding to the optimal dynamics: the mixing time of the induced Markov chain. For the discounted approach to accurately estimate the optimal policy, the discounting timescale (horizon) must be larger than the mixing time. However, estimating the mixing time for the optimal dynamics can be challenging in the general case, even when the transition dynamics are known. Therefore, an arbitrary “sufficiently

large” choice of γ is often made without knowledge of the relevant problem-dependent timescale. This can be problematic from a computational standpoint as evidenced by recent work (Jiang et al. 2015; Schulman et al. 2017; Andrychowicz et al. 2021).

The approach outlined in this work for average-reward RL (EVAL) has the added benefit that it can lead to an estimation of the mixing timescale for the optimal dynamics, which can then be used to inform the choice of a discount factor, if the discounted objective is of interest. In Figure 1, the discount factor set by the spectral gap of the tilted matrix (see Eq. (8) for definition) is indicated by the vertical line. We empirically find that this “spectral gap discount factor” naturally separates discount factors between “small” and “large” values, seen by the distinct change in the optimal state distributions (inset). The discount factor set by the spectral gap indicates a point of diminishing return for further increasing γ in soft Q-Learning (solid black line). For comparison, we plot the return given by our average-reward algorithm (dashed blue line). Importantly, we see that the average-reward solution recovers the discounted solution in the $\gamma \rightarrow 1$ limit, as expected (Blackwell 1962; Mahadevan 1996).

Despite the desirable features of both the average-reward and entropy-regularized objectives, the combination of these formulations (Neu, Jonsson, and Gómez 2017) is not as well-studied, and no function approximator algorithms exist for this setting. Here, we extend the ideas introduced in (Rawlik 2013; Arriojas 2022) to develop a new off-policy learning-based approach to find the optimal policy of average-reward MDPs. Our main contributions are as follows:

Main Contributions

- We provide a novel off-policy solution to average-reward RL in both entropy-regularized and un-regularized settings with general function approximators.
- We demonstrate experimentally the advantage of our approach against standard baselines in classic control environments.

Notably, our implementation requires minimal changes relative to the common DQN setup, making it accessible for researchers and allowing for multiple future extensions.

Preliminaries

Reinforcement learning treats decision-making problems under the framework of Markov Decision Processes (MDPs). In this section, we provide a brief account of the problem to be solved (the entropy-regularized average-reward) objective under this framework. As a point of notation, we will denote the set of distributions over a generic space \mathcal{X} as $\Delta(\mathcal{X})$. To begin, we describe the defining quantities of an MDP: the state space \mathcal{S} denotes all possible configurations of the agent within its environment; the initial state distribution $\mu \in \Delta(\mathcal{S})$ describes the initialization of the agent at each episode; the action space \mathcal{A} denotes the set of allowed controls the agent may exert to affect its environment; the transition function (also “dynamics”) is a function

$p : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ describing the probability of transitioning into a future state given a current state and action; and the reward function $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is a real-valued scalar provided to the agent at each time-step, encoding the desired behavior of the agent. In the following equations, we focus on the discrete state-action case for ease of notation. The applicability of the framework for continuous spaces is demonstrated empirically below.

Now, we state some of the usual assumptions for average-reward MDPs:

Assumption 1. The Markov chain induced by the dynamics p and any stationary policy π is irreducible and aperiodic.

Assumption 2. The reward function is upper bounded.¹

As in the discounted case, one seeks a control policy π which maximizes the expected return. In the discounted formulation, this objective is defined as a discounted infinite sum; but in the average-reward formulation, we instead consider the limiting value of the average trajectory reward, for increasingly long trajectories:

$$\sum_{t=0}^{\infty} \gamma^t r_t \rightarrow \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{t=0}^{N-1} r_t \quad (1)$$

This *reward rate*, ρ^π , becomes the new objective:

$$\rho^\pi = \lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}_{\tau \sim \pi, p, \mu} \left[\sum_{t=0}^{N-1} r(s_t, a_t) \right] = \mathbb{E}_\nu [r(s, a)], \quad (2)$$

with the expectation above taken with respect to the probability of a trajectory as generated by the dynamics p , fixed policy π , and initial state distribution μ . In the rightmost expression, ν is used to denote the stationary state-action distribution induced by the transition dynamics and choice of control policy. Because of Assumption 1, this distribution is well-defined as the eigenvector of the stochastic transition matrix P^π :

$$\sum_{s, a} p(s' | s, a) \pi(a' | s') \nu(s, a) = \nu(s', a'), \quad (3)$$

or in a more compact notation, $P^\pi \nu = \nu$. Related eigenvector equations will be our primary concern in approaching the entropy-regularized average-reward objective.

With the scalar reward-rate (“bias”) defined, we now turn to the corresponding “differential” value function, which plays the role of the standard Q function in discounted RL. Specifically, the (s, a) -dependent contribution to a trajectory’s value is denoted as Q_ρ^π .

$$Q_\rho^\pi(s, a) = \mathbb{E}_{\tau \sim p, \pi} \left[\sum_{t=0}^{\infty} r(s_t, a_t) - \rho^\pi \middle| s_0 = s, a_0 = a \right] \quad (4)$$

We now consider a more general version of this MDP which includes an entropy regularization term. Note that the

¹In the following we assume the reward is non-positive, which is equivalent to $r(s, a)$ having a finite upper-bound: the reward function can always be shifted to be entirely negative without affecting the optimal policy, merely shifting the reward rate accordingly.

original objective can be recovered in the zero temperature limit ($\beta \rightarrow \infty$). For convenience we will refer to entropy-regularized average-reward MDPs as ERAR MDPs. The ERAR MDP possesses the same components as an average-reward MDP as described, in addition to a pre-specified prior policy² $\pi_0 : \mathcal{S} \rightarrow \Delta(\mathcal{A})$ and “inverse temperature”, $\beta \in \mathbb{R}_{>0}$. The modified objective function for an ERAR MDP now includes a relative entropy based regularization term, such that the agent now aims to optimize the expected *entropy-regularized reward-rate*, denoted θ^π below:

$$\theta^\pi = \lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}_{\tau \sim p, \pi, \mu} \left[\sum_{t=0}^{N-1} r(s_t, a_t) - \frac{1}{\beta} \log \frac{\pi(a_t | s_t)}{\pi_0(a_t | s_t)} \right]. \quad (5)$$

The optimal policy for an ERAR MDP is defined analogously as the solution to

$$\pi^*(a | s) = \operatorname{argmax}_\pi \theta^\pi. \quad (6)$$

From the above, we see that the agent must balance the returns given by directly optimizing the reward with the cost of deviating from the prior policy π_0 . This also emphasizes the dependence of π_0 on the resulting optimal policy. In most work on entropy-regularized RL, the uniform distribution $\pi_0(a | s) \propto 1$ is used to express the prior belief that all actions are equally likely (preferable). However, this “maximum entropy” (MaxEnt) objective has a lack of flexibility, not allowing non-trivial priors or “guide” policies to be exploited, despite its apparent theoretical and experimental advantages (Adamczyk et al. 2023; Yan et al. 2024; Ahmed et al. 2019; Grau-Moya, Leibfried, and Vrancx 2018; Wu, Tucker, and Nachum 2019). In the following, we show that the more general non-uniform prior π_0 plays a crucial role in solving the un-regularized problem.

Our Assumption 1 guarantees the expression in Equation (5) is indeed independent of the initial state-action. If on the other hand the induced Markov chain were reducible, there may be some recurrent classes whose long-term reward rate depends on the initial state and action.

In the following, for convenience, we will simply write $\theta = \theta^{\pi^*}$ for the optimal entropy-regularized reward-rate. Now, in comparing to Equation (2), we see that the present objective, the “ERAR rate”, includes an entropic contribution: the relative entropy between the control and prior policies: $\text{KL}(\pi | \pi_0)$, where KL denotes the Kullback-Liebler divergence.

Corresponding to Eq. (4), the differential entropy-regularized action-value function is then given by³

$$Q_\theta^\pi(s, a) = \mathbb{E}_{\tau \sim p, \pi} \left[\sum_{t=0}^{\infty} r(s_t, a_t) - \frac{1}{\beta} \log \frac{\pi(a_t | s_t)}{\pi_0(a_t | s_t)} - \theta^\pi \right] \quad (7)$$

In this section, θ and ρ subscripts were used to distinguish the two (un-regularized and regularized) differential value

²We assume that π_0 has full support across \mathcal{A} , ensuring the Kullback-Liebler divergence between any policy π and π_0 remains finite.

³we have suppressed the conditioning on the initial state and action for brevity, but the conditioning is identical to that of Eq. (4).

functions. In the following sections, we omit the subscript as we focus solely on the entropy-regularized objective. Similarly, we use the shorthand $Q(s, a) = Q_\theta^*(s, a)$.

Prior Work

Classical approaches to the average-reward problem typically involve the following logic: Solve the discounted MDP for discount factors $\gamma_1 < \gamma_2 \dots < 1$. As the value of the discount factor approaches 1, the (appropriately re-scaled) optimal value function will approach that of the average-reward MDP (Blackwell 1962). Later work directly approaches the average-reward objective by attacking the associated Bellman equation head-on (Mahadevan 1996). Indeed, (Schwartz 1993) introduced *R-learning* (without proof of convergence), emphasizing the benefits of the average-reward algorithm: faster reward propagation, better value disambiguation, simpler learning dynamics near initialization, and potentially speedups. Schwartz (1993) also mentions how the discounted framework can be subsumed in the more general average-reward problem (Sec. 6.5 therein).

More recently, a variety of papers have considered the average-reward problem in a modern light, beginning to understand the theoretical properties of this objective and associated algorithms: (Even-Dar, Kakade, and Mansour 2009; Abbasi-Yadkori et al. 2019; Abounadi, Bertsekas, and Borkar 2001; Neu, Jonsson, and Gómez 2017; Wan, Naik, and Sutton 2021; Zhang et al. 2021). We note this set of prior works addresses the average-reward objective in the case of tabular settings, with discrete, finite state and action spaces. To address the continuous setting, recent work has developed (Zhang and Ross 2021; Ma et al. 2021; Saxena et al. 2023a) a variety of policy-based methods. In contrast, we will focus on the value-based setting here, learning the Q function directly in an off-policy manner. It is worth noting that these studies have found average-reward algorithms to be superior to discounted methods in continuous control Mujoco (Todorov, Erez, and Tassa 2012) tasks. A more complete description of the history of average-reward algorithms can be found in a recent survey: (Dewanto et al. 2020).

Alongside the average-reward literature, the discounted objective has seen the introduction of many techniques to generally improve sample complexity, by tackling planning (Hafner et al. 2023), estimation bias (Van Hasselt, Guez, and Silver 2016), and exploration (Park et al. 2023). Prior work has shown that the addition of an entropy-regularization term can improve robustness (Eysenbach and Levine 2022), exploration (Eysenbach et al. 2018), composability (Haarnoja et al. 2018a) and sample complexity (Haarnoja et al. 2018b). The theory of entropy regularization also has a long history, discussed in a series of work such as (Ziebart 2010; Rawlik 2013; Todorov 2006, 2009; Haarnoja et al. 2018b; Geist, Scherrer, and Pietquin 2019). This innovation yields controllably stochastic optimal policies, a flexibility which has led Soft Actor-Critic (Haarnoja et al. 2018c) and its variants to become state-of-the-art solution methods for addressing the discounted objective. To date, there are no direct combinations of the average-reward and entropy-regularized objectives for deep

reinforcement learning. However, recent work by (Rawlik 2013; Rose, Mair, and Garrahan 2021; Arriojas et al. 2023b) has established a framework for combining these formulations. In this work, we leverage the results derived in combination with function approximation techniques to establish value-based algorithms for the average-reward objective. In the following section we give an overview of the solution which builds on results obtained in (Arriojas et al. 2023b) before introducing our algorithms.

Solution Method

The presented solution method is based on insights and results from recent work (Mitter and Newton 2000; Rawlik, Toussaint, and Vijayakumar 2012; Rose, Mair, and Garrahan 2021; Arriojas et al. 2023b; Levine 2018) establishing a connection between the optimization problems in RL and statistical mechanics. Specifically, there is a mapping at the trajectory level between free energy minimization in statistical mechanics and (entropy-regularized) average reward maximization⁴. This mapping is based on control-as-inference approaches to RL as outlined below.

The control-as-inference perspective of RL reviewed in (Levine 2018) introduces a binary optimality variable to solve the MaxEnt RL problem for the case of deterministic dynamics. Specifically, Bayesian inference of the posterior dynamics resulting from conditioning on this optimality variable gives the solution to the optimal control problem. Furthermore, this approach also reveals that the optimal trajectory distribution follows a Boltzmann (softmax) distribution, similar to trajectory distributions arising from generalizations of the canonical ensemble in statistical mechanics (Rose, Mair, and Garrahan 2021; Chetrite and Touchette 2015b,a).

This connection to statistical mechanics indicates that the optimal trajectory distribution can be generated by a “tilted” transition matrix, which encodes information about the RL solution, as shown in (Arriojas et al. 2023b). Specifically, the trajectory distribution corresponding to the optimal policy for the ERAR objective can be generated from the tilted transition matrix:

$$\tilde{P}_{(s',a'),(s,a)} = p(s'|s, a)\pi_0(a'|s')e^{\beta r(s,a)}. \quad (8)$$

For a fixed value of β , the tilted matrix generates the dynamics for trajectories having an average return that is “tilted” or biased away from their typical value. Thus, by tuning β , the matrix \tilde{P} gives access to those trajectories with significantly larger expected returns than those obtained by running the system’s unbiased dynamics: $p(s'|s, a)\pi_0(a'|s')$.

Denoting the dominant eigenvalue and corresponding eigenvector of \tilde{P} as $e^{\beta\theta}$ and $u(s, a)$, the solution to the ERAR MDP is given by:

Lemma 1 (Arriojas et al. (2023b)). *The entropy-regularized average reward rate is given by θ , and the optimal differen-*

⁴Since energies can be viewed as negative rewards, these optimization problems are formally equivalent.

tial value function and optimal policy are given by:

$$Q(s, a) = \frac{1}{\beta} \log u(s, a), \quad (9)$$

$$\pi(a|s) = \frac{\pi_0(a|s)u(s, a)}{\sum_a \pi_0(a|s)u(s, a)}. \quad (10)$$

The solution method used requires the transition dynamics to be deterministic, though we discuss the stochastic generalization at the end of this section. It should also be noted that this method obtains the differential value and corresponding policy for the *optimal* entropy-regularized reward-rate, as opposed to a policy-evaluation technique, which may gather information on arbitrary suboptimal policies.

To gain further insights into the utility of the tilted matrix, consider the class of allowed trajectories (of length N) that start at a given state-action pair (s, a) and terminate at (s', a') . If we condition on optimality throughout the trajectory, then the posterior probability of observing this class of trajectories can be shown (Arriojas et al. 2023b) to be proportional to $[\tilde{P}^N]_{(s', a'), (s, a)}$. In the long-time limit (Eq. (5)), the probability for this class of trajectories is well-approximated by the dominant contribution in the spectral decomposition (thus corresponding to the largest eigenvalue) of \tilde{P} :

$$[\tilde{P}^N]_{(s', a'), (s, a)} \approx e^{N\beta\theta} u(s, a)v(s', a'), \quad (11)$$

where the left and right dominant eigenvectors u, v correspond to the dominant eigenvalue (Perron root) $\exp(\beta\theta)$. Now, to connect to the solution of the entropy-regularized RL problem, we recall that the soft value function is given by the log-probability of a trajectory being optimal (Levine 2018), given the initial state-action (s, a) . To obtain the corresponding probability, we thus need to sum over all possible final state-action pairs (s', a') in Eq. (11). Taking the logarithm then gives the soft value function as shown in Eq. (21) of (Arriojas et al. 2023b). This leads to the identification of θ as the ERAR rate and the derived results shown in Lemma 1.

In practice it may not be possible to construct this tilted matrix beyond the tabular setting without model-based techniques. Nevertheless, only a single component of its spectrum is required for an analytic solution to the ERAR MDP. As such, we will consider only the relevant component of the spectrum to devise a model-free learning algorithm, circumventing the need for learning the entire matrix.

The preceding discussion shows that the solution to the ERAR MDP is given by solving an eigenvalue equation for the tilted dynamics:

$$e^{\beta\theta} u(s, a) = \sum_{s', a'} u(s', a') \tilde{P}_{(s', a'), (s, a)}. \quad (12)$$

Because of the assumptions imposed, the Perron-Frobenius theorem applies to \tilde{P} , guaranteeing a *unique* positive left eigenvector u and well-defined ERAR rate θ . We can rewrite the above equation in the following form, reminiscent of a temporal-difference backup equation:

$$u(s, a) = e^{\beta(r(s, a) - \theta)} \mathbb{E}_{s' \sim p, a' \sim \pi_0} u(s', a'), \quad (13)$$

for an eigenvector $u > 0$ and (the unique real) eigenvalue with $\theta < 0$. The next largest eigenvalue (in magnitude), denoted $e^{\beta\xi}$, determines the spectral gap of \tilde{P} : $(\beta\xi - \beta\theta)^{-1}$. As discussed in the Introduction, the spectral gap determines an important timescale for the optimal dynamics. This mixing time controls the rate of convergence to the “(quasi) steady-state distribution” (v above) (Mélard and Villemonais 2012; Arriojas et al. 2023b) and hence controls the effective time horizon of the agent in its environment.

Given that we do not rely on policy-gradient techniques, and since the expectation in Equation (13) is over the pre-specified prior policy π_0 , we have an off-policy algorithm whose trajectory data can be collected by any rollout policy. Note that in log-space, this equation resembles the TD equation for soft Q-learning, without a discount factor, and with an additional correction for the reward-rate, θ .

As discussed by (Wan, Naik, and Sutton 2021), all average reward algorithms (except their Centered Differential TD-learning) are only able to learn an “uncentered” value function: i.e. a differential value function that may be off by some global constant. In contrast, since our solution method is based on a dominant eigenvalue and corresponding normalized eigenvector (guaranteed to exist through the Perron-Frobenius theorem, and normalized as a probability against the corresponding right eigenvector), our algorithm necessarily finds a “centered” differential value function for ERAR MDPs.

For the case of stochastic transition dynamics, the policy derived from the tilted matrix is optimal only if the agent can also change the transition dynamics correspondingly. However, controlling the dynamics is often infeasible and the transition dynamics is generally assumed to be fixed by the problem statement. Through Bayesian inference and matching the two objective functions, (Arriojas et al. 2023a) shows how to address this issue by iteratively biasing the reward function and transition dynamics of the original MDP so that the solution for a controlled problem coincides with the uncontrolled transition dynamics. To avoid these complications, in the following we focus on the case of deterministic transition functions. We discuss possible extensions to stochastic dynamics for future work in the penultimate section.

Proposed Algorithms

In this section, we present pseudocode of our proposed algorithms. We would like to highlight that the core of the algorithms here are built on DQN (Mnih et al. 2015). This means that existing codebases can easily adapt DQN-style methods with significant advantage (Fig. 2) with limited additional complexity. However, there are several important distinctions specific to the ERAR objective. In the next section we present Algorithm 1 which highlights these differences in **red**. In the following section, we present an algorithm to solve the unregularized average-reward objective. In the Appendix, the pseudocode for Algorithm 3 highlights the differences from Algorithm 1 in **blue**.

Algorithm 1: EVAL

```
1: IN: sample budget, environment,  $\beta$ , hyperparameters
2: Initialize:
3: Online network weights:  $\psi_i \sim \text{init. distribution}$ 
4: Target network weights:  $\psi_i = \psi_i$ 
5: Entropy-regularized reward-rate:  $\theta = 0$ 
6: Replay buffer:  $\mathcal{D} = \{\}$ 
7: while  $t < \text{sample budget}$  do
8:   Collect experience:
9:   Sample action  $a \sim \pi_0$ 
10:  Take step in environment  $s' \sim p(\cdot|s, a)$ 
11:  Save to replay buffer:  $\mathcal{D} \leftarrow \{s, a, r, s'\}$ 
12:  if train this step then
13:    for each gradient step do
14:      Sample minibatch  $\mathcal{B} \subset \mathcal{D}$ 
15:      Calculate loss via Eq. (14)
16:      Do gradient descent on  $u$  network(s)
17:      Calculate  $\theta_{\text{new}}$  via Equation (16)
18:    end for
19:  end if
20:  if update  $\theta$  then
21:    Update  $\theta$  estimate:  $\theta \leftarrow \theta \cdot (1 - \tau_\theta) + \theta_{\text{new}} \cdot \tau_\theta$ 
22:  end if
23:  if update target parameters then
24:    Update target parameters (Polyak averaging with
      parameters  $\tau_\psi$ ).
25:  end if
26: end while
27: OUT: Optimal policy for ERAR-MDP
```

Solution to ERAR-MDP

Our first algorithm implements three key components present in many value-based deep RL methods: (1) an estimate of the value function (left eigenvector) parameterized by two deep neural nets (inspired by (Van Hasselt, Guez, and Silver 2016)), (2) stochastic gradient descent on a temporal-difference error with Adam (Kingma and Ba 2015) and (3) a replay buffer of stored experience. In principle, the replay buffer can be collected by any behavior policy such as an ϵ -greedy policy (Mnih et al. 2015), but we use the learnt policy (Eq. (10)) for more effective exploration as in (Haarnoja et al. 2018b). Since we use an average-reward objective, we must also maintain a running estimate of the entropy-regularized reward-rate, θ .

Our model-free algorithm uses the update equations prescribed by (Arriojas et al. 2023b) (cf. Appendix) to learn $u(s, a)$ and θ through stochastic approximation.

For updating the u network, we use a mean squared error loss between the online network and corresponding estimate calculated through a target network:

$$\mathcal{J}(\psi) = \frac{1}{2} \mathbb{E}_{s, a \sim \mathcal{D}} (u_\psi(s, a) - \hat{u}_{\bar{\psi}}(s, a))^2. \quad (14)$$

We denote the trained online network as u_ψ and its temporal difference (TD) target as $\hat{u}_{\bar{\psi}}$ (we use \hat{u} to emphasize this is not a neural net itself, but a stand-in for the target value calculated by Eq. (15)). The TD target is calculated based

on the lagging network’s weights, denoted $\bar{\psi}$ (cf. Appendix for further implementation details). To find the TD target equation, we read off Equation (13) as:

$$\hat{u}_{\bar{\psi}}(s, a) = e^{\beta(r(s, a) - \theta)} \mathbb{E}_{s' \sim p, a' \sim \pi_0} u_{\bar{\psi}}(s', a'). \quad (15)$$

The value of θ must be updated online as well. To this end, we re-interpret Eq. (13) as an equation for the entropy-regularized reward-rate, θ . Since such an equation is valid for any s, a , we can use the entire batch of data \mathcal{B} (sampled uniformly from the replay buffer) to obtain a more accurate estimate for θ . This (s, a) -dependent calculation of θ is based on the current estimate of the left eigenvector, u_ψ (online network as opposed to target network). To preserve the linear structure of the eigenvector equation (Eq. (12)), we propose to track θ through the eigenvalue itself (exponential of θ):

$$e^{\beta\theta} = \frac{1}{|\mathcal{B}|} \sum_{\{s, a, r, s'\} \in \mathcal{B}} \frac{e^{\beta r} \mathbb{E}_{a' \sim \pi_0} u_\psi(s', a')}{u_\psi(s, a)}. \quad (16)$$

The value of θ is updated after averaging its value over the “gradient steps” loop in Algorithm 1. To improve stability of the next iteration of Equation (15), we find it helpful to allow θ to slowly mix with previous estimates, and hence use a (constant) step size to update its value (Line 21 in Algorithm 1).

Inspired by (Van Hasselt, Guez, and Silver 2016; Fujimoto, Hoof, and Meger 2018), we will train two online networks in parallel. We find this to help considerably improve the evaluation reward, as shown in Figure 4. Interestingly, we have found that the popular choice of taking a pessimistic estimate (i.e. min) is not optimal in our experiments. Instead we treat the “aggregation function” as a hyperparameter, and tune it over the choices of min, max, mean. Across all environments, we have found max to be the optimal choice for aggregation. We use two online networks and two corresponding target networks for updates, and upon calculating an aggregated estimate of $u_{\bar{\psi}}$ (shown in Eq. (15)) we use the max over the two target networks, calculated (s, a) -wise: $u_{\bar{\psi}}(s, a) = \max \{u_{\bar{\psi}}^{(1)}(s, a), u_{\bar{\psi}}^{(2)}(s, a)\}$. We train both of the online networks independently to minimize the same loss in Equation (14), measured against the aggregated target value. This is the same technique employed in (Fujimoto, Hoof, and Meger 2018) for discounted un-regularized RL and in (Haarnoja et al. 2018c) for discounted regularized RL. We use the same aggregation method (max) on the online networks $u_\psi^{(i)}$ to calculate θ across the current batch of data via Equation (16).

Posterior Policy Iteration

Although we plot the greedy policy’s reward as an evaluation metric, EVAL inherently aims to maximize the entropy-regularized reward-rate (a combination of rewards and policy relative entropy) shown in Equation (5). Thus, the ERAR MDP’s corresponding optimal policy is necessarily stochastic as discussed. The greedy policy used in evaluation, $\hat{\pi}(a|s) = \text{argmax}_a \pi_0(a|s)u(s, a)$ may therefore not be the

Algorithm 2: Posterior Policy Iteration (PPI)

Initialize: Prior policy $\pi_0, \beta > 0$, solve budget.
while $N < \text{solve budget}$ **do**
 $\pi_0 \leftarrow \text{Solve}(\pi_0, \beta)$
end while
Output: Deterministic optimal policy $\pi_{\beta=\infty}^* = \pi_0$

correct policy for maximizing the un-regularized average reward rate.

Instead, it may be of interest to directly obtain the true greedy solution corresponding to the un-regularized average-reward RL objective ($\beta \rightarrow \infty$) shown in Equation (2). A simple way to find such a solution is to slowly increase the value of β throughout training, as done in (Haarnoja et al. 2018b). However, this can lead to numerical instabilities because of the exponential term $\exp \beta (\tau(s, a) - \theta)$ in Equation 13. One workaround was proposed by an updated version of SAC (Haarnoja et al. 2018c), where the authors implement a learning-based method of temperature annealing based on entropy constraints in the dual space, facilitated by Lagrange multipliers.

A different method of recovering the un-regularized solution ($\beta \rightarrow \infty$) is based on an observation by (Rawlik 2013), which we outline below. Consider for notational convenience the function $\text{Solve} : (\mathbb{R}^{|S||A|}, \Delta(\mathcal{A}), \mathbb{R}_+) \rightarrow \Delta(\mathcal{A})$, which solves an ERAR MDP for a given choice of reward function, prior policy $\pi_0 \in \Delta(\mathcal{A})$, and inverse temperature β by outputting the corresponding optimal policy, $\pi^* \in \Delta(\mathcal{A})$ (which maximizes the objective in Equation (5)). Then, Theorem 4 in (Rawlik, Toussaint, and Vijayakumar 2012) indicates that the iteration shown in Algorithm 2 will converge to the greedy optimal policy: the solution of the un-regularized MDP. Thus, we implement this ‘‘Posterior Policy Iteration’’ (PPI) technique using a function approximator (MLP) to parameterize the prior policy. Additionally, we note that this method has the benefits of not requiring a set temperature, an ‘‘annealing schedule’’ $\beta_t \rightarrow \infty$ or requiring a minimum entropy with additional computational complexity for a learning scheme (Haarnoja et al. 2018c). These benefits come at the expense of some additional overhead (the memory and training of another network) and two new hyperparameters: the prior network’s update frequency and update weight (Polyak parameter, τ_ϕ, ω_ϕ in Algorithm 3).

To update the prior as prescribed in Algorithm 2, we periodically replace the policy π_0 with the current estimate of the optimal policy $\pi^*(a|s) \propto \pi_0(a|s)u(s, a)$ at a fixed frequency throughout training. Given sufficient time for convergence at each iteration, this process is guaranteed to converge to the solution for $\beta \rightarrow \infty$, thereby solving the average-reward MDP without entropy-regularization. In practice, we do not perform hard updates of the prior network’s weights, and instead employ Polyak averaging between the parameters of an online prior and a target prior. The online prior is trained to minimize the Kullback-Leibler (KL) divergence to the (current estimate of the) optimal pol-

icy via Equation (17):

$$\mathcal{L}_\phi = \frac{1}{|\mathcal{B}|} \sum_{\{s, a, r, s'\} \in \mathcal{B}} \text{KL} \left(\pi^{\bar{\phi}\psi}(\cdot|s) \mid \pi_0^\phi(\cdot|s) \right) \quad (17)$$

where \mathcal{B} is a mini-batch randomly sampled from the replay buffer and π is the current estimate for the optimal policy calculated from a combination of the online u network and target π_0 network according to Equation (10),

$$\pi^{\bar{\phi}\psi}(a|s) = \frac{\pi_0^\phi(\cdot|s)u^\psi(s, \cdot)}{\sum_a \pi_0^\phi(a|s)u^\psi(s, a)}. \quad (18)$$

Notice that since the current estimate of the policy is dependent upon the estimate of π_0 , we use the online version of both to select actions. Further implementation details and pseudocode of this extension are provided in Algorithm 3 below. Experimentally, we find that even for short timescales of iteration (compared to the timescale for a full solution), the PPI method is still able to recover the $\beta \rightarrow \infty$ solution, in agreement with Section 4.1.2 of (Rawlik 2013).

Experiments

We focus on the classic control environments from OpenAI’s Gymnasium (Brockman et al. 2016) with varying levels of complexity: CartPole-v1, Acrobot-v1, MountainCar-v0. First we will provide benchmark experiments, comparing EVAL against DQN (Mnih et al. 2015) (as it shares many implementation details) and SQL (Haarnoja et al. 2017a) (for its entropy-regularized objective). Although we recognize that these algorithms are not designed to optimize the same objective function, we do not have other benchmarks with which to compare in the discrete-action setting. On the other hand, we emphasize that the metric of interest in discounted RL is often the average evaluation reward. Against this metric, EVAL and EVAL-PPI have stronger performance in terms of sample complexity compared to DQN and SQL.

We compare EVAL with fixed (tuned) β and with PPI (denoted EVAL and EVAL-PPI, respectively) to DQN (Raffin et al. 2021) and SQL (Haarnoja et al. 2017a) on classic control environments in Figure 2. Hyperparameter tuning has been performed for each algorithm on each environment, with the associated values shown in Appendix .

For proper comparison with other un-regularized methods, the evaluation phase for EVAL uses the greedy policy derived from the learned stochastic policy. Every one thousand environment steps, we pause the training to perform a greedy evaluation, averaged over ten episodes. The resulting reward from each algorithm is averaged over twenty random initializations. All code for reproducing the experiments is available at <https://anonymous.4open.science/r/EVAL-D25E>.

Since the average-reward objective is aimed at solving continuing tasks, we devise a paradigmatic experiment (inspired by Fig. 2 in (Zhang and Ross 2021)) based on CartPole. showcasing the ability of EVAL in the long-time setting. Although both agents are trained in environments with




Figure 2: Classic control benchmark comparing soft Q-learning (SQL), deep Q network (DQN) and our two proposed methods (EVAL, EVAL+PPI). We find EVAL and EVAL+PPI to generally obtain higher reward with less variance than SQL or DQN.




Figure 3: As a demonstration of the usefulness of EVAL+PPI, we consider a modified version of CartPole which represents a continuing task. After training for 5000 steps (in the standard CartPole-v1 environment with a maximum episode length of 500), we compare the evaluation performance of SQL with EVAL+PPI. Specifically, we set the time-limit of the environment much higher: to 100,000 steps. We find that EVAL+PPI consistently reaches the maximum number of steps while SQL only rarely achieves similarly high reward. We find that EVAL+PPI can continue episodes for at least 10 billion steps (as of submission).

a default (500) timestep limit, the EVAL agent is able to generalize further. The results of this experiment, shown in Figure 3, indicate that EVAL+PPI is far superior to SQL at solving the (effectively) continuing task of balancing the pole forever.

Limitations and Future Work

The current work focuses on the case of deterministic transition dynamics. (Arriojas et al. 2023a) has developed a method for the more general case of stochastic transition dynamics, by iteratively learning biases for the dynamics and rewards. With a model-based algorithm, this seems to be a promising avenue for future exploration for the general case of stochastic transition dynamics. While we have focused on

the left eigenvector of \tilde{P} , the right eigenvector $v(s, a)$ can also be learned online. Learning v gives access to the optimal steady-state distribution (uv) a quantity of interest as mentioned in recent work (Saxena et al. 2023a). Although we have tested our method on classic control environments, other more challenging benchmarks such as the ALE (Bellemaire et al. 2013) can be used as well. Preliminary experiments show that EVAL is able to find the optimal policy in some Atari environments, but further hyperparameter tuning is needed to improve its stability and sample efficiency.

Our current method features discrete action spaces, but adapting the present algorithm to actor-based approaches could allow for a straightforward extension to continuous actions, analogous to the soft actor-critic approach (Haarnoja et al. 2018b). Since we use an off-policy method, it would be interesting to apply specific exploration policies for rollout collection. Since PPI (Rawlik, Toussaint, and Vijayakumar 2012) does not require a particular initial policy, this method can similarly benefit from an improved exploration policy at initialization. As a value-based technique, other ideas from the literature such as $TD(n)$, REDQ (Chen et al. 2021), PER (Schaul et al. 2015), or dueling architectures (Wang et al. 2016) may be included. An important contribution for future work is studying the sample complexity and convergence properties of the proposed algorithm. Further work may also provide a connection between PPI and temperature annealing procedures, and port this idea to the discounted framework as well.

Conclusion

By leveraging the connection of the ERAR objective to the solution of a particular eigenvalue problem, we have presented the first solution to deterministic ERAR MDPs in continuous state spaces by use of function approximation. Our experiments suggest that EVAL compares favorably in several respects to DQN and Soft Q-Learning. Our algorithm leverages the existing DQN framework allowing for a straightforward and easily extendable model applicable to the ERAR objective. We have also provided a solution to the un-regularized, greedy average-reward objective through an iterative procedure in the prior policy space. Our algorithms

show an advantage in the classic control suite, especially given that there are no other algorithms for this setting. We believe that the average-reward objective with entropy regularization is a fruitful direction for further research and real-world application, with this work addressing an important gap in the existing literature.

Acknowledgements

JA would like to acknowledge the use of the supercomputing facilities managed by the Research Computing Department at UMass Boston; the Unity high-performance computing cluster; and funding support from the Alliance Innovation Lab – Silicon Valley. RVK and JA would like to acknowledge funding support from the NSF through Award No. DMS-1854350 and PHY-2425180. ST was supported in part by NSF Award (2246221), PAZY grant (195-2020), and WCoE, Texas Tech U. This work is supported by the National Science Foundation under Cooperative Agreement PHY-2019786 (The NSF AI Institute for Artificial Intelligence and Fundamental Interactions, <http://iaifi.org/>).

References

- Abbasi-Yadkori, Y.; Bartlett, P.; Bhatia, K.; Lazic, N.; Szepesvári, C.; and Weisz, G. 2019. Politex: Regret bounds for policy iteration using expert prediction. In *International Conference on Machine Learning*, 3692–3702. PMLR.
- Abounadi, J.; Bertsekas, D.; and Borkar, V. S. 2001. Learning Algorithms for Markov Decision Processes with Average Cost. *SIAM Journal on Control and Optimization*, 40(3): 681–698.
- Adamczyk, J.; Arriojas, A.; Tiomkin, S.; and Kulkarni, R. V. 2023. Utilizing Prior Solutions for Reward Shaping and Composition in Entropy-Regularized Reinforcement Learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(6): 6658–6665.
- Adamczyk, J.; Makarenko, V.; Tiomkin, S.; and Kulkarni, R. V. 2024. Boosting Soft Q-Learning by Bounding. *Reinforcement Learning Journal*, 5: 2373–2399.
- Ahmed, Z.; Le Roux, N.; Norouzi, M.; and Schuurmans, D. 2019. Understanding the impact of entropy on policy optimization. In *International conference on machine learning*, 151–160. PMLR.
- Andrychowicz, M.; Raichuk, A.; Stańczyk, P.; Orsini, M.; Girgin, S.; Marinier, R.; Hussenot, L.; Geist, M.; Pietquin, O.; Michalski, M.; et al. 2021. What matters for on-policy deep actor-critic methods? a large-scale study. In *International conference on learning representations*.
- Arriojas, A.; Adamczyk, J.; Tiomkin, S.; and Kulkarni, R. V. 2023a. Bayesian inference approach for entropy regularized reinforcement learning with stochastic dynamics. In Evans, R. J.; and Shpitser, I., eds., *Proceedings of the Thirty-Ninth Conference on Uncertainty in Artificial Intelligence*, volume 216, 99–109. PMLR.
- Arriojas, A.; Adamczyk, J.; Tiomkin, S.; and Kulkarni, R. V. 2023b. Entropy regularized reinforcement learning using large deviation theory. *Phys. Rev. Res.*, 5: 023085.
- Arriojas, A. A. 2022. *Analytical Framework for Entropy Regularized Reinforcement Learning Using Probabilistic Inference*. Ph.D. thesis, University of Massachusetts Boston.
- Bellemare, M. G.; Naddaf, Y.; Veness, J.; and Bowling, M. 2013. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47: 253–279.
- Bertsekas, D. 2012. *Dynamic programming and optimal control: Volume I*, volume 4. Athena scientific.
- Blackwell, D. 1962. Discrete dynamic programming. *The Annals of Mathematical Statistics*, 719–726.
- Brockman, G.; Cheung, V.; Pettersson, L.; Schneider, J.; Schulman, J.; Tang, J.; and Zaremba, W. 2016. OpenAI gym. *arXiv preprint arXiv:1606.01540*.
- Ceron, J. S. O.; and Castro, P. S. 2021. Revisiting Rainbow: Promoting more insightful and inclusive deep reinforcement learning research. In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning*, volume 139, 1373–1383. PMLR.
- Chen, X.; Wang, C.; Zhou, Z.; and Ross, K. W. 2021. Randomized Ensembled Double Q-Learning: Learning Fast Without a Model. In *International Conference on Learning Representations*.
- Chetrite, R.; and Touchette, H. 2015a. Nonequilibrium Markov Processes Conditioned on Large Deviations. *Annales Henri Poincaré*, 16(9): 2005–2057.
- Chetrite, R.; and Touchette, H. 2015b. Variational and optimal control representations of conditioned and driven processes. *Journal of Statistical Mechanics: Theory and Experiment*, 2015(12): P12001.
- Dewanto, V.; Dunn, G.; Eshragh, A.; Gallagher, M.; and Roosta, F. 2020. Average-reward model-free reinforcement learning: a systematic review and literature mapping. *arXiv preprint arXiv:2010.08920*.
- Even-Dar, E.; Kakade, S. M.; and Mansour, Y. 2009. On-line Markov decision processes. *Mathematics of Operations Research*, 34(3): 726–736.
- Eysenbach, B.; Gupta, A.; Ibarz, J.; and Levine, S. 2018. Diversity is all you need: Learning skills without a reward function. *arXiv preprint arXiv:1802.06070*.
- Eysenbach, B.; and Levine, S. 2022. Maximum Entropy RL (Provably) Solves Some Robust RL Problems. In *International Conference on Learning Representations*.
- Fujimoto, S.; Hoof, H.; and Meger, D. 2018. Addressing function approximation error in actor-critic methods. In *International conference on machine learning*, 1587–1596. PMLR.
- Geist, M.; Scherrer, B.; and Pietquin, O. 2019. A theory of regularized Markov decision processes. In *International Conference on Machine Learning*, 2160–2169. PMLR.
- Grau-Moya, J.; Leibfried, F.; and Vrancx, P. 2018. Soft q-learning with mutual-information regularization. In *International conference on learning representations*.
- Haarnoja, T.; Pong, V.; Zhou, A.; Dalal, M.; Abbeel, P.; and Levine, S. 2018a. Composable deep reinforcement learning for robotic manipulation. In *2018 IEEE international*

- conference on robotics and automation (ICRA), 6244–6251. IEEE.
- Haarnoja, T.; Tang, H.; Abbeel, P.; and Levine, S. 2017a. Reinforcement Learning with Deep Energy-Based Policies. In Precup, D.; and Teh, Y. W., eds., *Proceedings of the 34th International Conference on Machine Learning*, volume 70, 1352–1361. PMLR.
- Haarnoja, T.; Tang, H.; Abbeel, P.; and Levine, S. 2017b. Reinforcement learning with deep energy-based policies. In *International conference on machine learning*, 1352–1361. PMLR.
- Haarnoja, T.; Zhou, A.; Abbeel, P.; and Levine, S. 2018b. Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor. In Dy, J.; and Krause, A., eds., *Proceedings of the 35th International Conference on Machine Learning*, volume 80, 1861–1870. PMLR.
- Haarnoja, T.; Zhou, A.; Hartikainen, K.; Tucker, G.; Ha, S.; Tan, J.; Kumar, V.; Zhu, H.; Gupta, A.; Abbeel, P.; et al. 2018c. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*.
- Hafner, D.; Pasukonis, J.; Ba, J.; and Lillicrap, T. 2023. Mastering diverse domains through world models. *arXiv preprint arXiv:2301.04104*.
- Hessel, M.; Modayil, J.; Van Hasselt, H.; Schaul, T.; Ostrovski, G.; Dabney, W.; Horgan, D.; Piot, B.; Azar, M.; and Silver, D. 2018. Rainbow: Combining improvements in deep reinforcement learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Jiang, N.; Kulesza, A.; Singh, S.; and Lewis, R. 2015. The dependence of effective planning horizon on model accuracy. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, 1181–1189.
- Kakade, S. M. 2003. *On the sample complexity of reinforcement learning*. Ph.D. thesis, University College London.
- Kingma, D.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations (ICLR)*. San Diego, CA, USA.
- Lee, K.; Laskin, M.; Srinivas, A.; and Abbeel, P. 2021. SUNRISE: A Simple Unified Framework for Ensemble Learning in Deep Reinforcement Learning. In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning*, volume 139, 6131–6141. PMLR.
- Levine, S. 2018. Reinforcement learning and control as probabilistic inference: Tutorial and review. *arXiv preprint arXiv:1805.00909*.
- Ma, X.; Tang, X.; Xia, L.; Yang, J.; and Zhao, Q. 2021. Average-Reward Reinforcement Learning with Trust Region Methods. In Zhou, Z.-H., ed., *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, 2797–2803. International Joint Conferences on Artificial Intelligence Organization. Main Track.
- Mahadevan, S. 1996. Average reward reinforcement learning: Foundations, algorithms, and empirical results. *Machine learning*, 22: 159–195.
- Méléard, S.; and Villemonais, D. 2012. Quasi-stationary distributions and population processes. *Probability Surveys*, 9: 340–410.
- Mitter, S. K.; and Newton, N. 2000. The duality between estimation and control. *Published in Festschrift for A. Benoussan*.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; et al. 2015. Human-level control through deep reinforcement learning. *nature*, 518(7540): 529–533.
- Naik, A.; Shariff, R.; Yasui, N.; Yao, H.; and Sutton, R. S. 2019. Discounted reinforcement learning is not an optimization problem. *arXiv preprint arXiv:1910.02140*.
- Naik, A.; Wan, Y.; Tomar, M.; and Sutton, R. S. 2024. Reward Centering. *arXiv preprint arXiv:2405.09999*.
- Neu, G.; Jonsson, A.; and Gómez, V. 2017. A unified view of entropy-regularized Markov decision processes. *arXiv preprint arXiv:1705.07798*.
- Park, S.; Lee, K.; Lee, Y.; and Abbeel, P. 2023. Controllability-Aware Unsupervised Skill Discovery. In Krause, A.; Brunskill, E.; Cho, K.; Engelhardt, B.; Sabato, S.; and Scarlett, J., eds., *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, 27225–27245. PMLR.
- Raffin, A.; Hill, A.; Gleave, A.; Kanervisto, A.; Ernestus, M.; and Dormann, N. 2021. Stable-Baselines3: Reliable Reinforcement Learning Implementations. *Journal of Machine Learning Research*, 22(268): 1–8.
- Rawlik, K.; Toussaint, M.; and Vijayakumar, S. 2012. On stochastic optimal control and reinforcement learning by approximate inference. *Proceedings of Robotics: Science and Systems VIII*.
- Rawlik, K. C. 2013. *On probabilistic inference approaches to stochastic optimal control*. Ph.D. thesis, The University of Edinburgh.
- Rose, D. C.; Mair, J. F.; and Garrahan, J. P. 2021. A reinforcement learning approach to rare trajectory sampling. *New Journal of Physics*, 23(1): 013013.
- Saxena, N.; Khastagir, S.; Kolathaya, S.; and Bhatnagar, S. 2023a. Off-policy average reward actor-critic with deterministic policy search. In *International Conference on Machine Learning*, 30130–30203. PMLR.
- Saxena, N.; Khastagir, S.; Shishir, N.; and Bhatnagar, S. 2023b. Off-policy average reward actor-critic with deterministic policy search. In *International Conference on Machine Learning*, 30130–30203. PMLR.
- Schaul, T.; Quan, J.; Antonoglou, I.; and Silver, D. 2015. Prioritized experience replay. *arXiv preprint arXiv:1511.05952*.
- Schulman, J.; Levine, S.; Abbeel, P.; Jordan, M.; and Moritz, P. 2015. Trust region policy optimization. In *International conference on machine learning*, 1889–1897. PMLR.

Schulman, J.; Moritz, P.; Levine, S.; Jordan, M.; and Abbeel, P. 2016. High-Dimensional Continuous Control Using Generalized Advantage Estimation. In *Proceedings of the International Conference on Learning Representations (ICLR)*.

Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.

Schwartz, A. 1993. A reinforcement learning method for maximizing undiscounted rewards. In *Proceedings of the tenth international conference on machine learning*, volume 298, 298–305.

Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement learning: An introduction*. MIT press.

Todorov, E. 2006. Linearly-solvable Markov decision problems. In Schölkopf, B.; Platt, J.; and Hoffman, T., eds., *Advances in Neural Information Processing Systems*, volume 19. MIT Press.

Todorov, E. 2009. Efficient computation of optimal actions. *Proceedings of the national academy of sciences*, 106(28): 11478–11483.

Todorov, E.; Erez, T.; and Tassa, Y. 2012. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ international conference on intelligent robots and systems*, 5026–5033. IEEE.

Van Hasselt, H.; Guez, A.; and Silver, D. 2016. Deep reinforcement learning with double Q-learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 30.

Wan, Y.; Naik, A.; and Sutton, R. S. 2021. Learning and planning in average-reward Markov decision processes. In *International Conference on Machine Learning*, 10653–10662. PMLR.

Wang, Z.; Schaul, T.; Hessel, M.; Hasselt, H.; Lanctot, M.; and Freitas, N. 2016. Dueling network architectures for deep reinforcement learning. In *International conference on machine learning*, 1995–2003. PMLR.

Wu, Y.; Tucker, G.; and Nachum, O. 2019. Behavior regularized offline reinforcement learning. *arXiv preprint arXiv:1911.11361*.

Yan, X.; Song, Y.; Feng, X.; Yang, M.; Zhang, H.; Ammar, H. B.; and Wang, J. 2024. Efficient Reinforcement Learning with Large Language Model Priors. *arXiv:2410.07927*.

Zhang, S.; Wan, Y.; Sutton, R. S.; and Whiteson, S. 2021. Average-reward off-policy policy evaluation with function approximation. In *international conference on machine learning*, 12578–12588. PMLR.

Zhang, Y.; and Ross, K. W. 2021. On-policy deep reinforcement learning for the average-reward criterion. In *International Conference on Machine Learning*, 12535–12545. PMLR.

Ziebart, B. D. 2010. *Modeling purposeful adaptive behavior with the principle of maximum causal entropy*. Carnegie Mellon University.

Theory

For convenience, we provide the update equations from (Arriojas et al. 2023b) (Eq. 26, 27) which will be learned with EVAL:

$$u(s, a) \rightarrow (1 - \alpha)u(s, a) + \alpha e^{\beta(r(s,a) - \theta)} u(s', a') \quad (19)$$

$$e^{\beta\theta} \rightarrow (1 - \alpha_\theta)e^{\beta\theta} + \alpha_\theta e^{\beta r(s,a)} \frac{u(s', a')}{u(s, a)} \quad (20)$$

which resembles our Eq. (15) and (16), respectively.

An important distinction between Eq. 26 in (Arriojas et al. 2023b) and Eq. (19) is that we have changed the sign of θ in their definition. This allows us to write the original objective in Eq. (5) in a way consistent with the average-reward literature (e.g. Eq. (2)).

Experimental Details

Implementation

Algorithm 3: EVAL with PPI

IN: sample budget, environment, β , hyperparameters

Initialize:

Online network weights: $\psi_i, \phi_i \sim W(\cdot)$

Target network weights: $\bar{\psi}_i = \psi_i, \bar{\phi}_i = \phi_i$

Reward-rate: $\theta = 0$

Replay buffer: $\mathcal{D} = \{\}$

while $t < \text{sample budget}$ **do**

 Collect experience:

 Sample action $a \sim \pi_0^\phi \propto u^\psi \cdot \pi_0^{\hat{\phi}}$

 Take step in environment $s' \sim p(\cdot | s, a)$

 Save to replay buffer: $\mathcal{D} \leftarrow \{s, a, r, s'\}$

for each gradient step **do**

 Sample minibatch $\mathcal{B} \subset \mathcal{D}$

 Calculate loss via Eq. (14) and Eq. (17)

 Do gradient descent using Adam on u networks and π_0 network

 Calculate θ_{new} via Equation (16)

 Update θ estimate: $\theta \leftarrow \theta \cdot (1 - \tau_\theta) + \theta_{\text{new}} \cdot \tau_\theta$

end for

if update target params **then**

 Update target parameters (Polyak averaging with parameters τ_ψ, τ_ϕ).

end if

end while

OUT: Optimal policy for AR-MDP

We implement our algorithm in PyTorch, and follow the style of Stable Baselines3 (Raffin et al. 2021). The function approximators for $u(s, a)$ (and similarly for π_0 , discussed below) are MLPs with state as input, and $|\mathcal{A}|$ output heads. Below, we show the pseudocode for EVAL with the PPI method (a combination of Algorithm 1 and 2 in the main text), allowing one to solve the un-regularized RL objective. We highlight the differences from Algorithm 1 (main text) in blue. As mentioned in the main text, we require a parameterization for the (online and target) prior policy as well as an update frequency and rolling average weight for the Polyak update.

We highlight the use of Eq. (16) in this algorithm, as it involves an expectation operation taken over the prior policy. We take the target prior for this calculation, as opposed to the online prior network for the calculation of the optimal policy, as seen in the loss function of Eq. (17).

Choice of Architecture

We note that since $u(s, a) > 0$ (by the Perron-Frobenius theorem), we require the u network outputs to be strictly positive and thus use a soft-plus activation function at the output layer. Although a ReLU activation is possible, we found it to give much worse performance across all tasks, despite additional hyperparameter tuning.

We similarly parameterize the prior net as an MLP (same hidden dimension as u -network), but with a softmax output. We have experimented with sharing weights between u and π_0 networks without finding significant performance gains, but this may be an interesting avenue for future exploration.

Soft Q-Learning

We implement our own soft Q-learning algorithm, in a similar style as EVAL (i.e. with a minimization over two networks, and use of a separate lagging target network). Similar to EVAL’s evaluation, we take the greedy policy $\pi(s) = \operatorname{argmax}_a Q(s, a)$ for evaluating the agent.

Hyperparameters

Table 1: Hyperparameter Ranges Used for Finetuning EVAL & PPI

Hyperparameter	Value Range
Learning Rate, η	$(10^{-4}, 10^{-1})$
Inverse Temperature, β	0.01, 0.05, 0.1, 0.5, 1, 2, 5, 10
Batch Size, b	16, 32, 64, 128, 256
Target Update Interval, ω	10, 100, 500, 1000, 2000, 5000
Gradient Steps, UTD	1,5,10,50

For each run, we choose the buffer size to be the same number of steps in the sample budget (except for DQN, which has a tuned value of buffer size B , given by (Raffin et al. 2021)). We also tune the hidden dimension “hid. dim.” of the MLP over values in $\{32, 64, 128\}$. In the following Table 2, ω_π refers to the prior policy’s target network update interval (in terms of environment steps).

For simplicity, we have found the algorithms to generally work well for some fixed hyperparameters, reducing the search space and potential sensitivity: The Polyak update ratio (for u and π_0 in PPI) is fixed to 1.0 (“hard updates”); The ERAR rate, θ is kept frozen at zero (except for irreducible tabular dynamics); the number of gradient steps per environment step is fixed to 5 throughout; The replay buffer size set to the maximum value (total number of environment steps during training); and the agent is trained after every environment step (cf. “training frequency” in (Raffin et al. 2021)). The finetuned hyperparameters for each environment are listed below. Each is the result of a sweep of roughly 200 runs (in random search), each:

Table 2: Finetuned Hyperparameter Values for EVAL(-PPI)

Environment	hid. dim.	η	b	β	ω	learn starts	ω_π
CartPole-v1	16	$1 \cdot 10^{-3}$	64	2	10	0	500
Acrobot-v1	64	$5 \cdot 10^{-4}$	64	0.01	10	0	500
MountainCar-v0	32	$6 \cdot 10^{-4}$	128	20	100	5000	2000

Table 3: The final two columns show the additionally tuned (with all others held fixed) hyperparameters specific to PPI.

Table 4: Finetuned Hyperparameter Values for DQN

Environment	η	b	B	γ	ω	$\varepsilon_{\text{final}}$	$\varepsilon_{\text{frac}}$	grad step	train freq	learn starts
CartPole-v1	$2.3 \cdot 10^{-3}$	64	100,000	0.99	10	0.04	0.16	128	256	1000
Acrobot-v1	$6.3 \cdot 10^{-4}$	128	50,000	0.99	250	0.1	0.12	-1	4	0
MountainCar-v0	$4 \cdot 10^{-3}$	128	10,000	0.98	600	0.07	0.2	8	16	1000

Table 5: $\varepsilon_{\text{frac}}$ denotes the exploration fraction over which to decay $\varepsilon = 1.0$ to $\varepsilon = \varepsilon_{\text{final}}$. A training frequency of -1 indicates that training of the Q networks occurs only after the end of each rollout episode. The provided optimal values for $\tau = 1.0$ and hidden dimension of 256 throughout all environments.

Table 6: Finetuned Hyperparameter Values for SQL

Environment	hid. dim	η	b	β	γ	τ	ω	grad steps	learn starts
CartPole-v1	64	$2 \cdot 10^{-2}$	64	0.1	0.98	0.95	100	9	1,000
Acrobot-v1	32	$6.6 \cdot 10^{-3}$	128	2.6	0.999	0.92	100	9	2,000
MountainCar-v0	64	$2 \cdot 10^{-3}$	128	0.7	0.99	0.97	100	2	9,000

Table 7: We use the finetuned hyperparameters given in (Adamczyk et al. 2024) for soft Q -learning in the classic control benchmark.

We use implementations of DQN from stable-baselines3, with the finetuned hyperparameters for each environment given by (Raffin et al. 2021).

Additional Results

Number of networks

In light of several recent works maintaining multiple estimates or an entire ensemble of the agent’s Q function (Lee et al. 2021; Chen et al. 2021), we consider the effect of learning more than two (online and target) networks on training performance. The results shown in Figure 4 indicate that having greater than one or two networks can drastically improve performance in some environments, with diminishing returns as seen in prior work. Our experiments indicate that a minimum of two networks is required for learning the optimal policy, so we use two networks throughout for simplicity, despite this generally being a tunable parameter. For convenience, we use two independent target networks as well, to separately perform rolling averages of the target parameters before aggregating.




Figure 4: For EVAL (without PPI) we compare the performance for training multiple networks in parallel. All networks are aggregated with the max function as discussed in the “Proposed Algorithms” section.