
A Theory of Regularized Markov Decision Processes

Matthieu Geist¹ Bruno Scherrer² Olivier Pietquin¹

Abstract

Many recent successful (deep) reinforcement learning algorithms make use of regularization, generally based on entropy or Kullback-Leibler divergence. We propose a general theory of regularized Markov Decision Processes that generalizes these approaches in two directions: we consider a larger class of regularizers, and we consider the general modified policy iteration approach, encompassing both policy iteration and value iteration. The core building blocks of this theory are a notion of regularized Bellman operator and the Legendre-Fenchel transform, a classical tool of convex optimization. This approach allows for error propagation analyses of general algorithmic schemes of which (possibly variants of) classical algorithms such as Trust Region Policy Optimization, Soft Q-learning, Stochastic Actor Critic or Dynamic Policy Programming are special cases. This also draws connections to proximal convex optimization, especially to Mirror Descent.

1. Introduction

Many reinforcement learning algorithms make use of some kind of entropy regularization, with various motivations, such as improved exploration and robustness. Trust Region Policy Optimization (TRPO) (Schulman et al., 2015) is a policy iteration scheme where the greedy step is penalized with a Kullback-Leibler (KL) penalty between two consecutive policies. Dynamic Policy Programming (DPP) (Azar et al., 2012) is a reparametrization of a value iteration scheme regularized by a KL penalty between consecutive policies. Soft Q-learning, eg. (Fox et al., 2016; Schulman et al., 2017; Haarnoja et al., 2017), uses a Shannon entropy regularization in a value iteration scheme, while Soft Actor Critic (SAC) (Haarnoja et al., 2018a) uses it in a policy iteration scheme. Value iteration has also been combined with a

¹Google Research, Brain Team. ²Université de Lorraine, CNRS, Inria, IECL, F-54000 Nancy, France. Correspondence to: Matthieu Geist <mfggeist@google.com>.

Tsallis entropy (Lee et al., 2018), with the motivation of having a sparse regularized greedy policy. Other approaches are based on a notion of temporal consistency equation, somehow extending the notion of Bellman residual to the regularized case (Nachum et al., 2017; Dai et al., 2018; Nachum et al., 2018), or on policy gradient (Williams, 1992; Mnih et al., 2016).

This non-exhaustive set of algorithms share the idea of using regularization, but they are derived from sometimes different principles, consider each time a specific regularization, and have ad-hoc analysis, if any. Here, we propose a general theory of regularized Markov Decision Processes (MDPs). To do so, a key observation is that (approximate) dynamic programming, or (A)DP, can be derived solely from the core definition of the Bellman evaluation operator. The framework we propose is built upon a regularized Bellman operator, and on an associated Legendre-Fenchel transform. We study the theoretical properties of these regularized MDPs and of the related regularized ADP schemes. This generalizes many existing theoretical results and provides new ones. Notably, it allows for an error propagation analysis for many of the aforementioned algorithms. This framework also draws connections to convex optimization, especially to Mirror Descent (MD).

A unified view of entropy-regularized MDPs has already been proposed by Neu et al. (2017). They focus on regularized DP through linear programming for the average reward case. Our contribution is complementary to this work (different MDP setting, we do not regularize the same quantity, we do not consider the same DP approach). Our use of the Legendre-Fenchel transform is inspired by Mensch & Blondel (2018), who consider smoothed finite horizon DP in directed acyclic graphs. Our contribution is also complementary to this work, that does not allow recovering aforementioned algorithms nor analyzing them. After a brief background, we introduce regularized MDPs and various related algorithmic schemes based on approximate modified policy iteration (Scherrer et al., 2015), as well as their analysis. All proofs are provided in the appendix.

2. Background

In this section, we provide the necessary background for building the proposed regularized MDPs. We write Δ_X the

set of probability distributions over a finite set X and Y^X the set of applications from X to the set Y . All vectors are column vectors, except distributions, for left multiplication. We write $\langle \cdot, \cdot \rangle$ the dot product and $\| \cdot \|_p$ the ℓ_p -norm.

2.1. Unregularized MDPs

An MDP is a tuple $\{\mathcal{S}, \mathcal{A}, P, r, \gamma\}$ with \mathcal{S} the finite¹ state space, \mathcal{A} the finite action space, $P \in \Delta_{\mathcal{S} \times \mathcal{A}}^{\mathcal{S}}$ the Markovian transition kernel ($P(s'|s, a)$ denotes the probability of transitioning to s' when action a is applied in state s), $r \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$ the reward function and $\gamma \in (0, 1)$ the discount factor.

A policy $\pi \in \Delta_{\mathcal{A}}^{\mathcal{S}}$ associates to each state a distribution over actions. The associated Bellman operator is defined as, for any function $v \in \mathbb{R}^{\mathcal{S}}$,

$$\forall s \in \mathcal{S}, [T_{\pi}v](s) = \mathbb{E}_{a \sim \pi(\cdot|s)} [r(s, a) + \gamma \mathbb{E}_{s'|s, a} [v(s')]].$$

This operator is a γ -contraction in supremum norm and its unique fixed-point is the value function v_{π} . With $r_{\pi}(s) = \mathbb{E}_{a \sim \pi(\cdot|s)} [r(s, a)]$ and $P_{\pi}(s'|s) = \mathbb{E}_{a \sim \pi(\cdot|s)} [P(s'|s, a)]$, the operator can be written as $T_{\pi}v = r_{\pi} + \gamma P_{\pi}v$. For any function $v \in \mathbb{R}^{\mathcal{S}}$, we associate the function $q \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$,

$$q(s, a) = r(s, a) + \gamma \mathbb{E}_{s'|s, a} [v(s')].$$

Thus, the Bellman operator can also be written as $[T_{\pi}v](s) = \langle \pi(\cdot|s), q(s, \cdot) \rangle = \langle \pi_s, q_s \rangle$. With a slight abuse of notation, we will write $T_{\pi}v = \langle \pi, q \rangle = (\langle \pi_s, q_s \rangle)_{s \in \mathcal{S}}$.

From this evaluation operator, one can define the Bellman optimality operator as, for any $v \in \mathbb{R}^{\mathcal{S}}$,

$$T_*v = \max_{\pi} T_{\pi}v.$$

This operator is also a γ -contraction in supremum norm, and its fixed point is the optimal value function v_* . From the same operator, one can also define the notion of a policy being greedy respectively to a function $v \in \mathbb{R}^{\mathcal{S}}$:

$$\pi' \in \mathcal{G}(v) \Leftrightarrow T_*v = T_{\pi'}v \Leftrightarrow \pi' \in \underset{\pi}{\operatorname{argmax}} T_{\pi}v.$$

Given this, we could derive value iteration, policy iteration, modified policy iteration, and so on. Basically, we can do all these things from the core definition of the Bellman evaluation operator. We'll do so from a notion of regularized Bellman evaluation operator.

2.2. Legendre-Fenchel transform

Let $\Omega : \Delta_{\mathcal{A}} \rightarrow \mathbb{R}$ be a strongly convex function. The Legendre-Fenchel transform (or convex conjugate) of Ω is $\Omega^* : \mathbb{R}^{\mathcal{A}} \rightarrow \mathbb{R}$, defined as

$$\forall q_s \in \mathbb{R}^{\mathcal{A}}, \Omega^*(q_s) = \max_{\pi_s \in \Delta_{\mathcal{A}}} \langle \pi_s, q_s \rangle - \Omega(\pi_s).$$

¹We assume a finite space for simplicity of exposition, our results extend to more general cases.

We'll make use of the following properties (Hiriart-Urruty & Lemaréchal, 2012; Mensch & Blondel, 2018).

Proposition 1. *Let Ω be strongly convex, we have the following properties.*

i *Unique maximizing argument:* $\nabla \Omega^*$ is Lipschitz and satisfies $\nabla \Omega^*(q_s) = \operatorname{argmax}_{\pi_s \in \Delta_{\mathcal{A}}} \langle \pi_s, q_s \rangle - \Omega(\pi_s)$.

ii *Boundedness:* if there are constants L_{Ω} and U_{Ω} such that for all $\pi_s \in \Delta_{\mathcal{A}}$, we have $L_{\Omega} \leq \Omega(\pi_s) \leq U_{\Omega}$, then $\max_{a \in \mathcal{A}} q_s(a) - U_{\Omega} \leq \Omega^*(q_s) \leq \max_{a \in \mathcal{A}} q_s(a) - L_{\Omega}$.

iii *Distributivity:* for any $c \in \mathbb{R}$ (and $\mathbf{1}$ the vector of ones), we have $\Omega^*(q_s + c\mathbf{1}) = \Omega^*(q_s) + c$.

iv *Monotonicity:* $q_{s,1} \leq q_{s,2} \Rightarrow \Omega^*(q_{s,1}) \leq \Omega^*(q_{s,2})$.

A classical example is the negative entropy $\Omega(\pi_s) = \sum_a \pi_s(a) \ln \pi_s(a)$. Its convex conjugate is the smoothed maximum $\Omega^*(q_s) = \ln \sum_a \exp q_s(a)$ and the unique maximizing argument is the usual softmax $\nabla \Omega^*(q_s) = \frac{\exp q_s(a)}{\sum_b \exp q_s(b)}$. For a positive regularizer, one can consider $\Omega(\pi_s) = \sum_a \pi_s(a) \ln \pi_s(a) + \ln |\mathcal{A}|$, that is the KL divergence between π_s and a uniform distribution. Its convex conjugate is $\Omega^*(q_s) = \ln \sum_a \frac{1}{|\mathcal{A}|} \exp q_s(a)$, that is the Mellowmax operator (Asadi & Littman, 2017). The maximizing argument is still the softmax. Another less usual example is the negative Tsallis entropy (Lee et al., 2018), $\Omega(\pi_s) = \frac{1}{2} (\|\pi_s\|_2^2 - 1)$. The analytic convex conjugate is more involved, but it leads to the sparsemax as the maximizing argument (Martins & Astudillo, 2016).

3. Regularized MDPs

The core idea of our contribution is to regularize the Bellman evaluation operator. Recall that $[T_{\pi}v](s) = \langle \pi_s, q_s \rangle$. A natural idea is to replace it by $[T_{\pi, \Omega}v](s) = \langle \pi_s, q_s \rangle - \Omega(\pi_s)$. To get the related optimality operator, one has to perform state-wise maximization over $\pi_s \in \Delta_{\mathcal{A}}$, which gives the Legendre-Fenchel transform of $[T_{\pi, \Omega}v](s)$. This defines a smoothed maximum (Nesterov, 2005). The related maximizing argument defines the notion of greedy policy.

3.1. Regularized Bellman operators

We now define formally these regularized Bellman operators. With a slight abuse of notation, we write $\Omega(\pi) = (\Omega(\pi_s))_{s \in \mathcal{S}}$ (and similarly for Ω^* and $\nabla \Omega^*$).

Definition 1 (Regularized Bellman operators). *Let $\Omega : \Delta_{\mathcal{A}} \rightarrow \mathbb{R}$ be a strongly convex function. For any $v \in \mathbb{R}^{\mathcal{S}}$ define $q \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$ as $q(s, a) = r(s, a) + \gamma \mathbb{E}_{s'|s, a} [v(s')]$. The regularized Bellman evaluation operator is defined as*

$$T_{\pi, \Omega} : v \in \mathbb{R}^{\mathcal{S}} \rightarrow T_{\pi, \Omega}v = T_{\pi}v - \Omega(\pi) \in \mathbb{R}^{\mathcal{S}},$$

that is, state-wise, $[T_{\pi,\Omega}v](s) = \langle \pi_s, q_s \rangle - \Omega(\pi_s)$. The regularized Bellman optimality operator is defined as

$$T_{*,\Omega} : v \in \mathbb{R}^{\mathcal{S}} \rightarrow T_{*,\Omega}v = \max_{\pi \in \Delta_{\mathcal{A}}^{\mathcal{S}}} T_{\pi,\Omega}v = \Omega^*(q) \in \mathbb{R}^{\mathcal{S}},$$

that is, state-wise, $[T_{*,\Omega}v](s) = \Omega^*(q_s)$. For any function $v \in \mathbb{R}^{\mathcal{S}}$, the associated unique greedy policy is defined as

$$\pi' = \mathcal{G}_{\Omega}(v) = \nabla \Omega^*(q) \Leftrightarrow T_{\pi',\Omega}v = T_{*,\Omega}v,$$

that is, state-wise, $\pi'_s = \nabla \Omega^*(q_s)$.

To be really useful, these operators should satisfy the same properties as the classical ones. It is indeed the case (we recall that all proofs are provided in the appendix).

Proposition 2. *The operator $T_{\pi,\Omega}$ is affine and we have the following properties.*

i *Monotonicity:* let $v_1, v_2 \in \mathbb{R}^{\mathcal{S}}$ such that $v_1 \geq v_2$. Then,

$$T_{\pi,\Omega}v_1 \geq T_{\pi,\Omega}v_2 \text{ and } T_{*,\Omega}v_1 \geq T_{*,\Omega}v_2.$$

ii *Distributivity:* for any $c \in \mathbb{R}$, we have that

$$\begin{aligned} T_{\pi,\Omega}(v + c\mathbf{1}) &= T_{\pi,\Omega}v + \gamma c\mathbf{1} \\ \text{and } T_{*,\Omega}(v + c\mathbf{1}) &= T_{*,\Omega}v + \gamma c\mathbf{1}. \end{aligned}$$

iii *Contraction:* both operators are γ -contractions in supremum norm. For any $v_1, v_2 \in \mathbb{R}^{\mathcal{S}}$,

$$\begin{aligned} \|T_{\pi,\Omega}v_1 - T_{\pi,\Omega}v_2\|_{\infty} &\leq \gamma \|v_1 - v_2\|_{\infty} \\ \text{and } \|T_{*,\Omega}v_1 - T_{*,\Omega}v_2\|_{\infty} &\leq \gamma \|v_1 - v_2\|_{\infty}. \end{aligned}$$

3.2. Regularized value functions

The regularized operators being contractions, we can define regularized value functions as their unique fixed-points. Notice that from the following definitions, we could also easily derive regularized Bellman operators on q -functions.

Definition 2 (Regularized value function of policy π). *Noted $v_{\pi,\Omega}$, it is defined as the unique fixed point of the operator $T_{\pi,\Omega}$: $v_{\pi,\Omega} = T_{\pi,\Omega}v_{\pi,\Omega}$. We also define the associated state-action value function $q_{\pi,\Omega}$ as*

$$\begin{aligned} q_{\pi,\Omega}(s, a) &= r(s, a) + \gamma \mathbb{E}_{s'|s,a}[v_{\pi,\Omega}(s')] \\ \text{with } v_{\pi,\Omega}(s) &= \mathbb{E}_{a \sim \pi(\cdot|s)}[q_{\pi,\Omega}(s, a)] - \Omega(\pi(\cdot|s)). \end{aligned}$$

Thus, the regularized value function is simply the unregularized value of π for the reward $r_{\pi} - \Omega(\pi)$, that is $v_{\pi,\Omega} = (I - \gamma P_{\pi})^{-1}(r_{\pi} - \Omega(\pi))$.

Definition 3 (Regularized optimal value function). *Noted $v_{*,\Omega}$, it is the unique fixed point of the operator $T_{*,\Omega}$: $v_{*,\Omega} = T_{*,\Omega}v_{*,\Omega}$. We also define the associated state-action value function $q_{*,\Omega}(s, a)$ as*

$$\begin{aligned} q_{*,\Omega}(s, a) &= r(s, a) + \gamma \mathbb{E}_{s'|s,a}[v_{*,\Omega}(s')] \\ \text{with } v_{*,\Omega}(s) &= \Omega^*(q_{*,\Omega}(s, \cdot)). \end{aligned}$$

The function $v_{*,\Omega}$ is indeed the optimal value function, thanks to the following result.

Theorem 1 (Optimal regularized policy). *The policy $\pi_{*,\Omega} = \mathcal{G}_{\Omega}(v_{*,\Omega})$ is the unique optimal regularized policy, in the sense that for all $\pi \in \Delta_{\mathcal{A}}^{\mathcal{S}}$, $v_{\pi_{*,\Omega}} = v_{*,\Omega} \geq v_{\pi,\Omega}$.*

When regularizing the MDP, we change the problem at hand. The following result relates value functions in (un)regularized MDPs.

Proposition 3. *Assume that $L_{\Omega} \leq \Omega \leq U_{\Omega}$. Let π be any policy. We have that $v_{\pi} - \frac{U_{\Omega}}{1-\gamma}\mathbf{1} \leq v_{\pi,\Omega} \leq v_{\pi} - \frac{L_{\Omega}}{1-\gamma}\mathbf{1}$ and $v_{*} - \frac{U_{\Omega}}{1-\gamma}\mathbf{1} \leq v_{*,\Omega} \leq v_{*} - \frac{L_{\Omega}}{1-\gamma}\mathbf{1}$.*

Regularization changes the optimal policy, the next result shows how it performs in the original MDP.

Theorem 2. *Assume that $L_{\Omega} \leq \Omega \leq U_{\Omega}$. We have that*

$$v_{*} - \frac{U_{\Omega} - L_{\Omega}}{1-\gamma} \leq v_{\pi_{*,\Omega}} \leq v_{*}.$$

3.3. Related Works

Some of these results already appeared in the literature, in different forms and with specific regularizers. For example, the contraction of $T_{*,\Omega}$ (Prop. 2) was shown in various forms, e.g. (Fox et al., 2016; Asadi & Littman, 2017; Dai et al., 2018), as well as the relation between (un)regularized optimal value functions (Th. 2), e.g. (Lee et al., 2018; Dai et al., 2018). The link to Legendre-Fenchel has also been considered before, e.g. (Dai et al., 2018; Mensch & Blondel, 2018; Richemond & Maginnis, 2017).

The core contribution of Sec. 3 is the regularized Bellman operator, inspired by Nesterov (2005) and Mensch & Blondel (2018). It allows building in a principled and general way regularized MDPs, and generalizing existing results easily. More importantly, it is the core building block of regularized (A)DP, studied in the next sections. The framework and analysis we propose next rely heavily on this formalism.

4. Regularized Modified Policy Iteration

Having defined the notion of regularized MDPs, we still need algorithms that solve them. As the regularized Bellman operators have the same properties as the classical ones, we can apply classical dynamic programming. Here, we consider directly the modified policy iteration approach (Puterman & Shin, 1978), that we regularize (reg-MPI for short):

$$\begin{cases} \pi_{k+1} = \mathcal{G}_{\Omega}(v_k) \\ v_{k+1} = (T_{\pi_{k+1},\Omega})^m v_k \end{cases} \quad (1)$$

Given an initial v_0 , reg-MPI iteratively performs a regularized greedy step to get π_{k+1} and a partial regularized evaluation step to get v_{k+1} .

With $m = 1$, we retrieve a regularized value iteration algorithm, that can be simplified as $v_{k+1} = T_{*,\Omega} v_k$ (as π_{k+1} is greedy resp. to v_k , we have $T_{\pi_{k+1},\Omega} v_k = T_{*,\Omega} v_k$). With $m = \infty$, we obtain a regularized policy iteration algorithm, that can be simplified as $\pi_{k+1} = \mathcal{G}_\Omega(v_{\pi_k,\Omega})$ (indeed, with a slight abuse of notation, $(T_{\pi_k,\Omega})^\infty v_{k-1} = v_{\pi_k,\Omega}$).

Before studying the convergence and rate of convergence of this general algorithmic scheme (with approximation), we discuss its links to state of the art algorithms (and more generally how it can be practically instantiated).

4.1. Related algorithms

Most existing schemes consider the negative entropy as the regularizer. Usually, it is also more convenient to work with q-functions. First, we consider the case $m = 1$. In the exact case, the regularized value iteration scheme can be written

$$q_{k+1}(s, a) = r(s, a) + \gamma \mathbb{E}_{s'|s, a} [\Omega^*(q_k(s', \cdot))].$$

In the entropic case, $\Omega^*(q_k(s, \cdot)) = \ln \sum_a \exp q_k(s, a)$. In an approximate setting, the q-function can be parameterized by parameters θ (for example, the weights of a neural network), write $\bar{\theta}$ the target parameters (computed during the previous iteration) and $\hat{\mathbb{E}}$ the empirical expectation over sampled transitions (s_i, a_i, r_i, s'_i) , an iteration amounts to minimize the expected loss

$$J(\theta) = \hat{\mathbb{E}} \left[(\hat{q}_i - q_\theta(s_i, a_i))^2 \right] \quad (2)$$

with $\hat{q}_i = r_i + \gamma \Omega^*(q_{\bar{\theta}}(s'_i, \cdot))$.

Getting a practical algorithm may require more work, for example for estimating $\Omega^*(q_{\bar{\theta}}(s'_i, \cdot))$ in the case of continuous actions (Haarnoja et al., 2017), but this is the core principle of soft Q-learning (Fox et al., 2016; Schulman et al., 2017). This idea has also been applied using the Tsallis entropy as the regularizer (Lee et al., 2018).

Alternatively, assume that q_k has been estimated. One could compute the regularized greedy policy analytically, $\pi_{k+1}(\cdot|s) = \nabla \Omega^*(q_k(s, \cdot))$. Instead of computing this for any state-action couple, one can generalize this from observed transitions to any state-action couple through a parameterized policy π_w , by minimizing the KL divergence between both distributions:

$$J(w) = \hat{\mathbb{E}} [\text{KL}(\pi_w(\cdot|s_i) || \nabla \Omega^*(q_k(s_i, \cdot)))]. \quad (3)$$

This is done in SAC (Haarnoja et al., 2018a), with an entropic regularizer (and thus $\nabla \Omega^*(q_k(s, \cdot)) = \frac{\exp q_k(s, \cdot)}{\sum_a \exp q_k(s, a)}$). This is also done in Maximum A Posteriori Policy Optimization (MPO) (Abdolmaleki et al., 2018b) with a KL regularizer (a case we discuss Sec. 5), or by Abdolmaleki et al. (2018a) with more general ‘‘conservative’’ greedy policies.

Back to SAC, q_k is estimated using a TD-like approach, by minimizing² for the current policy π :

$$J(\theta) = \hat{\mathbb{E}} [(\hat{q}_i - q_\theta(s_i, a_i))^2] \quad (4)$$

with $\hat{q}_i = r_i + \gamma (\mathbb{E}_{a \sim \pi(\cdot|s'_i)} [q_{\bar{\theta}}(s'_i, a)] - \Omega(\pi(\cdot, s'_i)))$.

For SAC, we have $\Omega(\pi(\cdot, s)) = \mathbb{E}_{a \sim \pi(\cdot|s)} [\ln \pi(a|s)]$ specifically (negative entropy). This approximate evaluation step corresponds to $m = 1$, and SAC is therefore more a VI scheme than a PI scheme, as presented by Haarnoja et al. (2018a) (the difference with soft Q-learning lying in how the greedy step is performed, implicitly or explicitly). It could be extended to the case $m > 1$ in two ways. One possibility is to minimize m times the expected loss (4), updating the target parameter vector $\bar{\theta}$ between each optimization, but keeping the policy π fixed. Another possibility is to replace the 1-step rollout of Eq. (4) by an m -step rollout (similar to classical m -step rollouts, up to the additional regularizations correcting the rewards). Both are equivalent in the exact case, but not in the general case.

Depending on the regularizer, Ω^* or $\nabla \Omega^*$ might not be known analytically. In this case, one can still solve the greedy step directly. Recall that the regularized greedy policy satisfies $\pi_{k+1} = \max_\pi T_{\pi,\Omega} v_k$. In an approximate setting, this amounts to maximize³

$$J(w) = \hat{\mathbb{E}} [\mathbb{E}_{a \sim \pi_w(\cdot|s_i)} [q_k(s_i, a)] - \Omega(\pi_w(\cdot|s_i))]. \quad (5)$$

This improvement step is used by Riedmiller et al. (2018) with an entropy, as well as by TRPO (up to the fact that the objective is constrained rather than regularized), with a KL regularizer (see Sec. 5).

To sum up, for any regularizer Ω , with $m = 1$ one can concatenate greedy and evaluation steps as in Eq. (2), with $m \geq 1$ one can estimate the greedy policy using either Eqs. (3) or (5), and estimate the q-function using Eq. 4, either performed m times repeatedly or combined with m -step rollouts, possibly combined with off-policy correction such as importance sampling or Retrace (Munos et al., 2016).

4.2. Analysis

We analyze the propagation of errors of the scheme depicted in Eq. (1), and as a consequence, its convergence and rate of convergence. To do so, we consider possible errors in both the (regularized) greedy and evaluation steps,

$$\begin{cases} \pi_{k+1} = \mathcal{G}_\Omega^{\epsilon_{k+1}}(v_k) \\ v_{k+1} = (T_{\pi_{k+1},\Omega})^m v_k + \epsilon_{k+1} \end{cases}, \quad (6)$$

²Actually, a separate network is used to estimate the value function, but it is not critical here.

³One could add a state-dependant baseline to q_k , eg. v_k , this does not change the maximizer but can reduce the variance.

with $\pi_{k+1} = \mathcal{G}_\Omega^{\epsilon'_{k+1}}(v_k)$ meaning that for any policy π , we have $T_{\pi, \Omega} v_k \leq T_{\pi_{k+1}, \Omega} v_k + \epsilon'_{k+1}$. The following analysis is basically the same as the one of Approximate Modified Policy Iteration (AMPI) (Scherrer et al., 2015), thanks to the results of Sec. 3 (especially Prop. 2).

The distance we bound is the loss $l_{k, \Omega} = v_{*, \Omega} - v_{\pi_k, \Omega}$. The bound will involve the terms $d_0 = v_{*, \Omega} - v_0$ and $b_0 = v_0 - T_{\pi_1, \Omega} v_0$. It requires also defining the following.

Definition 4 (Γ -matrix (Scherrer et al., 2015)). For $n \in \mathbb{N}^*$, \mathbb{P}_n is the set of transition kernels defined as **1**) for any set of n policies $\{\pi_1, \dots, \pi_n\}$, $\prod_{i=1}^n (\gamma P_{\pi_i}) \in \mathbb{P}_n$ and **2**) for any $\alpha \in (0, 1)$ and $(P_1, P_2) \in \mathbb{P}_n \times \mathbb{P}_n$, $\alpha P_1 + (1 - \alpha) P_2 \in \mathbb{P}_n$. Any element of \mathbb{P}_n is denoted Γ^n .

We first state a point-wise bound on the loss. This is the same bound as for AMPI, generalized to regularized MDPs.

Theorem 3. After k iterations of scheme (6), we have

$$l_{k, \Omega} \leq 2 \sum_{i=1}^{k-1} \sum_{j=i}^{\infty} \Gamma^j |\epsilon_{k-i}| + \sum_{i=0}^{k-1} \sum_{j=i}^{\infty} \Gamma^j |\epsilon'_{k-i}| + h(k)$$

with $h(k) = 2 \sum_{j=k}^{\infty} \Gamma^j |d_0|$ or $h(k) = 2 \sum_{j=k}^{\infty} \Gamma^j |b_0|$.

Next, we provide a bound on the weighted ℓ_p -norm of the loss, defined for a distribution ρ as $\|l_k\|_{p, \rho}^p = \rho |l_k|^p$. Again, this is the AMPI bound generalized to regularized MDPs.

Corollary 1. Let ρ and μ be distributions. Let p, q and q' such that $\frac{1}{q} + \frac{1}{q'} = 1$. Define the concentrability coefficients $C_q^i = \frac{1-\gamma}{\gamma^i} \sum_{j=i}^{\infty} \gamma^j \max_{\pi_1, \dots, \pi_j} \left\| \frac{\rho P_{\pi_1} P_{\pi_2} \dots P_{\pi_j}}{\mu} \right\|_{q, \mu}$. After k iterations of scheme (6), the loss satisfies

$$\begin{aligned} \|l_{k, \Omega}\|_{p, \rho} &\leq 2 \sum_{i=1}^{k-1} \frac{\gamma^i}{1-\gamma} (C_q^i)^{\frac{1}{p}} \|\epsilon_{k-i}\|_{pq', \mu} \\ &\quad + \sum_{i=0}^{k-1} \frac{\gamma^i}{1-\gamma} (C_q^i)^{\frac{1}{p}} \|\epsilon'_{k-i}\|_{pq', \mu} + g(k) \end{aligned}$$

with $g(k) = \frac{2\gamma^k}{1-\gamma} (C_q^i)^{\frac{1}{p}} \min(\|d_0\|_{pq', \mu}, \|b_0\|_{pq', \mu})$.

As this is the same bound (up to the fact that it deals with regularized MDPs) as the one of AMPI, we refer to Scherrer et al. (2015) for a broad discussion about it. It is similar to other error propagation analyses in reinforcement learning, and generalizes those that could be obtained for regularized value or policy iteration. The factor m does not appear in the bound. This is also discussed by Scherrer et al. (2015), but basically this depends on where the error is injected. We could derive a regularized version of Classification-based Modified Policy Iteration (CBMPI, see Scherrer et al. (2015) again) and make it appear.

So, we get the same bound for reg-MPI that for unregularized AMPI, no better nor worse. This is a good thing,

as it justifies considering regularized MDPs, but it does not explain the good empirical results of related algorithms.

With regularization, policies will be more stochastic than in classical approximate DP (that tends to produce deterministic policies). Such stochastic policies can induce lower concentrability coefficients. We also hypothesize that regularizing the greedy step helps controlling the related approximation error, that is the $\|\epsilon'_{k-i}\|_{pq', \mu}$ terms. Digging this question would require instantiating more the algorithmic scheme and performing a finite sample analysis of the resulting optimization problems. We left this for future work, and rather pursue the general study of solving regularized MDPs, with varying regularizers now.

5. Mirror Descent Modified Policy Iteration

Solving a regularized MDP provides a solution that differs from the one of the unregularized MDP (see Thm. 2). The problem we address here is estimating the original optimal policy while solving regularized greedy steps. Instead of considering a fixed regularizer $\Omega(\pi)$, the key idea is to penalize a divergence between the policy π and the policy obtained at the previous iteration of an MPI scheme. We consider more specifically the Bregman divergence generated by the strongly convex regularizer Ω .

Let π' be some given policy (typically π_k , when computing π_{k+1}), the Bregman divergence generated by Ω is

$$\begin{aligned} \Omega_{\pi'}(\pi_s) &= D_\Omega(\pi_s \| \pi') \\ &= \Omega(\pi_s) - \Omega(\pi') - \langle \nabla \Omega(\pi'), \pi_s - \pi' \rangle. \end{aligned}$$

For example, the KL divergence is generated by the negative entropy: $\text{KL}(\pi_s \| \pi') = \sum_a \pi_s(a) \ln \frac{\pi_s(a)}{\pi'(a)}$. With a slight abuse of notation, as before, we will write

$$\Omega_{\pi'}(\pi) = D_\Omega(\pi \| \pi') = \Omega(\pi) - \Omega(\pi') - \langle \nabla \Omega(\pi'), \pi - \pi' \rangle.$$

This divergence is always positive, it satisfies $\Omega_{\pi'}(\pi') = 0$, and it is strongly convex in π (so Prop. 1 applies).

We consider a reg-MPI algorithmic scheme with a Bregman divergence replacing the regularizer. For the greedy step, we simply consider $\pi_{k+1} = \mathcal{G}_{\Omega_{\pi_k}}(v_k)$, that is

$$\pi_{k+1} = \underset{\pi}{\operatorname{argmax}} \langle q_k, \pi \rangle - D_\Omega(\pi \| \pi_k).$$

This is similar to the update of the Mirror Descent (MD) algorithm in its proximal form (Beck & Teboulle, 2003), with $-q_k$ playing the role of the gradient in MD. Therefore, we will call this approach Mirror Descent Modified Policy Iteration (MD-MPI). For the partial evaluation step, we can regularize according to the previous policy π_k , that is $v_{k+1} = (T_{\pi_{k+1}, \Omega_{\pi_k}})^m v_k$, or according to the current policy π_{k+1} , that is $v_{k+1} = (T_{\pi_{k+1}, \Omega_{\pi_{k+1}}})^m v_k$. As

$\Omega_{\pi_{k+1}}(\pi_{k+1}) = 0$, this simplifies as $v_{k+1} = (T_{\pi_{k+1}})^m v_k$, that is a partial unregularized evaluation.

To sum up, we will consider two general algorithmic schemes based on a Bregman divergence, MD-MPI types 1 and 2 respectively defined as

$$\begin{cases} \pi_{k+1} = \mathcal{G}_{\Omega_{\pi_k}}(v_k) \\ v_{k+1} = (T_{\pi_{k+1}, \Omega_{\pi_k}})^m v_k \end{cases}, \begin{cases} \pi_{k+1} = \mathcal{G}_{\Omega_{\pi_k}}(v_k) \\ v_{k+1} = (T_{\pi_{k+1}})^m v_k \end{cases}$$

and both initialized with some v_0 and π_0 .

5.1. Related algorithms

To derive practical algorithms, the recipes provided in Sec. 4.1 still apply, just replacing Ω by Ω_{π_k} . If $m = 1$, greedy and evaluation steps can be concatenated (only for MD-MPI type 1). In the general case ($m \geq 1$) the greedy policy (for MD-MPI types 1 and 2) can be either directly estimated (Eq. (5)) or trained to generalize the analytical solution (Eq. (3)). The partial evaluation can be done using a TD-like approach, either done repeatedly while keeping the policy fixed or considering m -step rollouts. Specifically, in the case of a KL divergence, one could use the fact that $\Omega_{\pi_k}^*(q_k(s, \cdot)) = \ln \sum_a \pi_k(a|s) \exp q_k(s, a)$ and that $\nabla \Omega_{\pi_k}^*(q_k(s, \cdot)) = \frac{\pi_k(\cdot|s) \exp q_k(s, \cdot)}{\sum_a \pi_k(a|s) \exp q_k(s, a)}$.

This general algorithmic scheme allows recovering state of the art algorithms. For example, MD-MPI type 2 with $m = \infty$ and a KL divergence as the regularizer is TRPO (Schulman et al., 2015) (with a direct optimization of the regularized greedy step, as in Eq. (5), up to the use of a constraint instead of a regularization). DPP can be seen as a reparametrization⁴ of MD-MPI type 1 with $m = 1$ (Azar et al., 2012, Appx. A). MPO (Abdolmaleki et al., 2018b) is derived from an expectation-maximization principle, but it can be seen as an instantiation of MD-MPI type 2, with a KL divergence, a greedy step similar to Eq. (3) (up to additional regularization) and an evaluation step similar to Eq. (4) (without regularization, as in type 2, with m -step return and with the Retrace off-policy correction). This also generally applies to the approach proposed by Abdolmaleki et al. (2018a) (up to an additional subtlety in the greedy step consisting in decoupling updates for the mean and variance in the case of a Gaussian policy).

5.2. Analysis

Here, we propose to analyze the error propagation of MD-MPI (and thus, its convergence and rate of convergence). We think this is an important topic, as it has only been partly

⁴Indeed, if one see MD-MPI as a Mirror Descent approach, one can see DPP as a dual averaging approach, somehow updating a kind of cumulative q -functions directly in the dual. However, how to generalize this beyond the specific DPP algorithm is unclear, and we let it for future work.

studied for the special cases discussed in Sec. 5.1. For example, DPP enjoys an error propagation analysis in supremum norm (yet it is a reparametrization of a special case of MD-MPI, so not directly covered here), while TRPO or MPO are only guaranteed to have monotonic improvements, under some assumptions. Notice that we do not claim that our analysis covers all these cases, but it will provide the key technical aspects to analyze similar schemes (much like CBMPI compared to AMPI, as discussed in Sec. 4.2 or by Scherrer et al. (2015); where the error is injected changes the bounds).

In Sec. 4.2, the analysis was a straightforward adaptation of the one of AMPI, thanks to the results of Sec. 3 (the regularized quantities behave like their unregularized counterparts). It is no longer the case here, as the regularizer changes over iterations, depending on what has been computed so far. We will notably need a slightly different notion of approximate regularized greediness.

Definition 5 (Approximate Bregman divergence-regularized greediness). *Write $J_k(\pi)$ the (negative) optimization problem corresponding to the Bregman divergence-regularized greediness (that is, negative regularized Bellman operator of π applied to v_k):*

$$J_k(\pi) = \langle -q_k, \pi \rangle + D_{\Omega}(\pi || \pi_k) = -T_{\pi, \Omega_{\pi_k}} v_k.$$

We write $\pi_{k+1} \in \mathcal{G}_{\Omega_{\pi_k}}^{\epsilon'_{k+1}}(v_k)$ if for any policy π the policy π_{k+1} satisfies

$$\langle \nabla J_k(\pi_{k+1}), \pi - \pi_{k+1} \rangle + \epsilon'_{k+1} \geq 0.$$

In other words, $\pi_{k+1} \in \mathcal{G}_{\Omega_{\pi_k}}^{\epsilon'_{k+1}}(v_k)$ means that π_{k+1} is ϵ'_{k+1} -close to satisfying the optimality condition, which might be slightly stronger than being ϵ'_{k+1} -close to the optimal (as for AMPI or reg-MPI). Given this, we consider MD-MPI with errors in both greedy and evaluation steps, type 1

$$\begin{cases} \pi_{k+1} = \mathcal{G}_{\Omega_{\pi_k}}^{\epsilon'_{k+1}}(v_k) \\ v_{k+1} = (T_{\pi_{k+1}, \Omega_{\pi_k}})^m v_k + \epsilon_{k+1} \end{cases}$$

and type 2

$$\begin{cases} \pi_{k+1} = \mathcal{G}_{\Omega_{\pi_k}}^{\epsilon'_{k+1}}(v_k) \\ v_{k+1} = (T_{\pi_{k+1}})^m v_k + \epsilon_{k+1} \end{cases}.$$

The quantity we are interested in is $v_* - v_{\pi_k}$, that is sub-optimality in the unregularized MDP, while the algorithms compute new policies with a regularized greedy operator. So, we need to relate regularized and unregularized quantities when using a Bregman divergence based on the previous policy. The next lemma is the key technical result that allows analyzing MD-MPI.

Lemma 1. Assume that $\pi_{k+1} \in \mathcal{G}_{\Omega_{\pi_k}}^{\epsilon'_{k+1}}(v_k)$, as defined in Def. 5. Then, the policy π_{k+1} is ϵ'_{k+1} -close to the regularized greedy policy, in the sense that for any policy π

$$T_{\pi, \Omega_{\pi_k}} v_k - T_{\pi_{k+1}, \Omega_{\pi_k}} v_k \leq \epsilon'_{k+1}.$$

Moreover, we can relate the (un)regularized Bellman operators applied to v_k . For any policy π (so notably for the unregularized optimal policy π_*), we have

$$T_{\pi} v_k - T_{\pi_{k+1}, \Omega_{\pi_k}} v_k \leq \epsilon'_{k+1} + D_{\Omega}(\pi || \pi_k) - D_{\Omega}(\pi || \pi_{k+1}),$$

$$T_{\pi} v_k - T_{\pi_{k+1}} v_k \leq \epsilon'_{k+1} + D_{\Omega}(\pi || \pi_k) - D_{\Omega}(\pi || \pi_{k+1}).$$

We're interested in bounding the loss $l_k = v_* - v_{\pi_k}$, or some related quantity, for each type of MD-MPI. To do so, we introduce quantities similar to the ones of the AMPI analysis (Scherrer et al., 2015), defined respectively for types 1 and 2: **1**) The distance between the optimal value function and the value before approximation at the k^{th} iteration, $d_k^1 = v_* - (T_{\pi_k, \Omega_{\pi_{k-1}}})^m v_{k-1} = v_* - (v_k - \epsilon_k)$ and $d_k^2 = v_* - (T_{\pi_k})^m v_{k-1} = v_* - (v_k - \epsilon_k)$; **2**) The shift between the value before approximation and the policy value a iteration k , $s_k^1 = (T_{\pi_k, \Omega_{\pi_{k-1}}})^m v_{k-1} - v_{\pi_k} = (v_k - \epsilon_k) - v_{\pi_k}$ and $s_k^2 = (T_{\pi_k})^m v_{k-1} - v_{\pi_k} = (v_k - \epsilon_k) - v_{\pi_k}$; **3**) the Bellman residual at iteration k , $b_k^1 = v_k - T_{\pi_{k+1}, \Omega_{\pi_k}} v_k$ and $b_k^2 = v_k - T_{\pi_{k+1}} v_k$.

For both types ($h \in \{1, 2\}$), we have that $l_k^h = d_k^h + s_k^h$, so bounding the loss requires bounding these quantities, which is done in the following lemma (quantities related to both types enjoy the same bounds).

Lemma 2. Let $k \geq 1$, define $x_k = (I - \gamma P_{\pi_k}) \epsilon_k + \epsilon'_{k+1}$ and $y_k = -\gamma P_{\pi_*} \epsilon_k + \epsilon'_{k+1}$, as well as $\delta_k(\pi_*) = D_{\Omega}(\pi_* || \pi_k) - D_{\Omega}(\pi_* || \pi_{k+1})$. We have for $h \in \{1, 2\}$:

$$\begin{aligned} b_k^h &\leq (\gamma P_{\pi_k})^m b_{k-1}^h + x_k, \\ s_k^h &\leq (\gamma P_{\pi_k})^m (I - \gamma P_{\pi_k})^{-1} b_{k-1}^h \text{ and} \\ d_{k+1}^h &\leq \gamma P_{\pi_*} d_k^h + y_k + \sum_{j=1}^{m-1} (\gamma P_{\pi_{k+1}})^j b_k^h + \delta_k(\pi_*). \end{aligned}$$

These bounds are almost the same as the ones of AMPI (Scherrer et al., 2015, Lemma 2), up to the additional $\delta_k(\pi_*)$ term in the bound of the distance d_k^h . One can notice that summing these terms gives a telescopic sum: $\sum_{k=0}^{K-1} \delta_k(\pi_*) = D_{\Omega}(\pi_* || \pi_0) - D_{\Omega}(\pi_* || \pi_K) \leq D_{\Omega}(\pi_* || \pi_0) \leq \sup_{\pi} D_{\Omega}(\pi || \pi_0)$. For example, if D_{Ω} is the KL divergence and π_0 the uniform policy, then $\|\sup_{\pi} D_{\Omega}(\pi || \pi_0)\|_{\infty} = \ln |\mathcal{A}|$. This suggests that we must bound the regret L_K defined as

$$L_K = \sum_{k=1}^K l_k = \sum_{k=1}^K (v_* - v_{\pi_k}).$$

Theorem 4. Define $R_{\Omega_{\pi_0}} = \|\sup_{\pi} D_{\Omega}(\pi || \pi_0)\|_{\infty}$, after K iterations of MD-MPI, for $h = 1, 2$, the regret satisfies

$$\begin{aligned} L_K &\leq 2 \sum_{k=2}^K \sum_{i=1}^{k-1} \sum_{j=i}^{\infty} \Gamma^j |\epsilon_{k-i}| + \sum_{k=1}^K \sum_{i=0}^{k-1} \sum_{j=i}^{\infty} \Gamma^j |\epsilon'_{k-i}| \\ &\quad + \sum_{k=1}^K h(k) + \frac{1 - \gamma^K}{(1 - \gamma)^2} R_{\Omega_{\pi_0}} \mathbf{1}. \end{aligned}$$

with $h(k) = 2 \sum_{j=k}^{\infty} \Gamma^j |d_0|$ or $h(k) = 2 \sum_{j=k}^{\infty} \Gamma^j |b_0|$.

From this, we can derive an ℓ_p -bound for the regret.

Corollary 2. Let ρ and μ be distributions over states. Let p , q and q' be such that $\frac{1}{q} + \frac{1}{q'} = 1$. Define the concentrability coefficients C_q^i as in Cor. 1. After K iterations, the regret satisfies

$$\begin{aligned} \|L_K\|_{p, \rho} &\leq 2 \sum_{k=2}^K \sum_{i=1}^{k-1} \frac{\gamma^i}{1 - \gamma} (C_q^i)^{\frac{1}{p}} \|\epsilon_{k-i}\|_{pq', \mu} \\ &\quad + \sum_{k=1}^K \sum_{i=0}^{k-1} \frac{\gamma^i}{1 - \gamma} (C_q^i)^{\frac{1}{p}} \|\epsilon'_{k-i}\|_{pq', \mu} \\ &\quad + g(k) + \frac{1 - \gamma^K}{(1 - \gamma)^2} R_{\Omega_{\pi_0}}. \end{aligned}$$

with $g(k) = 2 \sum_{k=1}^K \frac{\gamma^k}{1 - \gamma} (C_q^k)^{\frac{1}{p}} \min(\|d_0\|_{pq', \mu}, \|b_0\|_{pq', \mu})$.

This result bounds the regret, while it is usually the loss that is bounded. Both can be related as follows.

Proposition 4. For any $p \geq 1$ and distribution ρ , we have $\min_{1 \leq k \leq K} \|v_* - v_{\pi_k}\|_{1, \rho} \leq \frac{1}{K} \|L_K\|_{p, \rho}$.

This means that if we can control the average regret, then we can control the loss of the best policy computed so far. This suggests that practically we should not use the last policy, but this best policy.

From Cor. 2 can be derived the convergence and rate of convergence of MD-MPI in the exact case.

Corollary 3. Both MD-MPI type 1 and 2 enjoy the following rate of convergence, when no approximation is done ($\epsilon_k = \epsilon'_k = 0$),

$$\frac{1}{K} \|L_K\|_{\infty} \leq \frac{1 - \gamma^K}{(1 - \gamma)^2} \frac{2\gamma \|v_* - v_0\|_{\infty} + R_{\Omega_{\pi_0}}}{K}.$$

In classical DP and in regularized DP (see Cor. 1), there is a linear convergence rate (the bound is $\frac{2\gamma^K}{1 - \gamma} \|v_* - v_0\|_{\infty}$), while in this case we only have a logarithmic convergence rate. We also pay an horizon factor (square dependency in $\frac{1}{1 - \gamma}$ instead of linear). This is normal, as we bound the regret instead of the loss. Bounding the regret in classical DP would lead to the bound of Cor. 3 (without the $R_{\Omega_{\pi_0}}$ term).

The convergence rate of the loss of MD-MPI is an open question, but a sublinear rate is quite possible. Compared to classical DP, we slow down greediness by adding the Bregman divergence penalty. Yet, this kind of regularization is used in an approximate setting, where it favors stability empirically (even if studying this further would require much more work regarding the $\|\epsilon'_k\|$ term, as discussed in Sec. 4.2).

As far as we know, the only other approach that studies a DP scheme regularized by a divergence and that offers a convergence rate is DPP, up to the reparameterization we discussed earlier. MD-MPI has the same upper-bound as DPP in the exact case (Azar et al., 2012, Thm. 2). However, DPP bounds the loss, while we bound a regret. This means that if the rate of convergence of our loss can be sublinear, it is superlogarithmic (as the rate of the regret is logarithmic), while the rate of the loss of DPP is logarithmic.

To get more insight on Cor. 2, we can group the terms differently, by grouping the errors.

Corollary 4. *With the same notations as Cor. 2, we have*

$$\frac{1}{K} \|L_k\|_{p,\rho} \leq \sum_{i=1}^{K-1} \frac{\gamma^i}{1-\gamma} (C_q^i)^{\frac{1}{p}} \frac{2E_{K-i} + E'_{K-i}}{K} + \frac{1}{K} \left(g(k) + \frac{1-\gamma^K}{(1-\gamma)^2} R_{\Omega_{\pi_0}} \right),$$

with $E_i = \sum_{j=1}^i \|\epsilon_j\|_{pq',\mu}$ and $E'_i = \sum_{j=1}^i \|\epsilon'_j\|_{pq',\mu}$.

Compared to the bound of AMPI (Scherrer et al., 2015, Thm. 7), instead of propagating the errors, we propagate the sum of errors over previous iterations normalized by the total number of iterations. So, contrary to approximate DP, it is no longer the last iterations that have the highest influence on the regret. Yet, we highlight again the fact that we bound a regret, and bounding the regret of AMPI would provide a similar result.

Our result is similar to the error propagation of DPP (Azar et al., 2012, Thm. 5), except that we sum norms of errors, instead of norming a sum of errors, the later being much better (as it allows the noise to cancel over iterations). Yet, as said before, DPP is not a special case of our framework, but a reparameterization of such one. Consequently, while we estimate value functions, DPP estimate roughly at iteration k a sum of k advantage functions (converging to $-\infty$ for any suboptimal action in the exact case). As explained before, where the error is injected does matter. Knowing if the DPP's analysis can be generalized to our framework (MPI scheme, ℓ_p bounds) remains an open question.

To get further insight, we can express the bound using different concentrability coefficients.

Corollary 5. *Define the concentrability coefficient $C_q^{l,k}$ as*

$C_q^{l,k} = \frac{(1-\gamma)^2}{\gamma^l - \gamma^k} \sum_{i=l}^{k-1} \sum_{j=i}^{\infty} c_q(j)$, the regret then satisfies

$$\|L_K\|_{p,\rho} \leq 2 \sum_{i=1}^{K-1} \frac{\gamma - \gamma^{i+1}}{(1-\gamma)^2} (C_q^{1,i+1})^{\frac{1}{p}} \|\epsilon_{K-i}\|_{pq',\mu} + \sum_{i=0}^{K-1} \frac{1 - \gamma^{i+1}}{(1-\gamma)^2} (C_q^{0,i+1})^{\frac{1}{p}} \|\epsilon'_{K-i}\|_{pq',\mu} + f(k)$$

with $f(k) = \frac{\gamma - \gamma^{K+1}}{(1-\gamma)^2} (C_q^{1,K+1})^{\frac{1}{p}} \min(\|d_0\|_{pq',\mu}, \|b_0\|_{pq',\mu}) + \frac{1-\gamma^K}{(1-\gamma)^2} R_{\Omega_{\pi_0}}$.

We observe again that contrary to ADP, the last iteration does not have the highest influence, and we do not enjoy a decrease of influence at the exponential rate γ towards the initial iterations. However, we bound a different quantity (regret instead of loss), that explains this behavior. Here again, bounding the regret in AMPI would lead to the same bound (up to the term $R_{\Omega_{\pi_0}}$). Moreover, sending p and K to infinity, defining $\epsilon = \sup_j \|\epsilon_j\|_{\infty}$ and $\epsilon' = \sup_j \|\epsilon'_j\|_{\infty}$, we get $\limsup_{K \rightarrow \infty} \frac{1}{K} \|L_K\|_{\infty} \leq \frac{2\gamma\epsilon + \epsilon'}{(1-\gamma)^2}$, which is the classical asymptotical bound for approximate value and policy iterations (Bertsekas & Tsitsiklis, 1996) (usually stated without greedy error). It is generalized here to an approximate MPI scheme regularized with a Bregman divergence.

6. Conclusion

We have introduced a general theory of regularized MDPs, where the usual Bellman evaluation operator is modified by either a fixed convex function or a Bregman divergence between consecutive policies. For both cases, we proposed a general algorithmic scheme based on MPI. We shown how many (variations of) existing algorithms could be derived from this general algorithmic scheme, and also analyzed and discussed the related propagation of errors.

We think that this framework can open many perspectives, among which links between (approximate) DP and proximal convex optimization (going beyond mirror descent), temporal consistency equations (roughly regularized Bellman residuals), regularized policy search (maximizing the expected regularized value function), inverse reinforcement learning (thanks to uniqueness of greediness in this regularized framework) or zero-sum Markov games (regularizing the two-player Bellman operators). We develop more these points in the appendix.

This work also leaves open questions, such as combining the propagation of errors with a finite sample analysis, or what specific regularizer one should choose for what context. Some approaches also combine a fixed regularizer and a divergence (Akrouf et al., 2018), a case not covered here and worth being investigated.

References

- Abdolmaleki, A., Springenberg, J. T., Degraeve, J., Bohez, S., Tassa, Y., Belov, D., Heess, N., and Riedmiller, M. Relative entropy regularized policy iteration. *arXiv preprint arXiv:1812.02256*, 2018a.
- Abdolmaleki, A., Springenberg, J. T., Tassa, Y., Munos, R., Heess, N., and Riedmiller, M. Maximum a posteriori policy optimisation. In *International Conference on Learning Representations (ICLR)*, 2018b.
- Akrou, R., Abdolmaleki, A., Abdulsamad, H., Peters, J., and Neumann, G. Model-free trajectory-based policy optimization with monotonic improvement. *The Journal of Machine Learning Research (JMLR)*, 19(1):565–589, 2018.
- Asadi, K. and Littman, M. L. An alternative softmax operator for reinforcement learning. In *International Conference on Machine Learning (ICML)*, 2017.
- Azar, M. G., Gómez, V., and Kappen, H. J. Dynamic policy programming. *Journal of Machine Learning Research (JMLR)*, 13(Nov):3207–3245, 2012.
- Beck, A. and Teboulle, M. Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters*, 31(3):167–175, 2003.
- Bertsekas, D. P. and Tsitsiklis, J. N. *Neuro-Dynamic Programming*. Athena Scientific, 1st edition, 1996. ISBN 1886529108.
- Dai, B., Shaw, A., Li, L., Xiao, L., He, N., Liu, Z., Chen, J., and Song, L. Sbeed: Convergent reinforcement learning with nonlinear function approximation. In *International Conference on Machine Learning (ICML)*, 2018.
- Finn, C., Levine, S., and Abbeel, P. Guided cost learning: Deep inverse optimal control via policy optimization. In *International Conference on Machine Learning (ICML)*, 2016.
- Fox, R., Pakman, A., and Tishby, N. Taming the noise in reinforcement learning via soft updates. In *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2016.
- Fu, J., Luo, K., and Levine, S. Learning robust rewards with adversarial inverse reinforcement learning. In *International Conference on Representation Learning*, 2018.
- Haarnoja, T., Tang, H., Abbeel, P., and Levine, S. Reinforcement learning with deep energy-based policies. In *International Conference on Machine Learning (ICML)*, 2017.
- Haarnoja, T., Zhou, A., Abbeel, P., and Levine, S. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International Conference on Machine Learning (ICML)*, 2018a.
- Haarnoja, T., Zhou, A., Hartikainen, K., Tucker, G., Ha, S., Tan, J., Kumar, V., Zhu, H., Gupta, A., Abbeel, P., et al. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*, 2018b.
- Hiriart-Urruty, J.-B. and Lemaréchal, C. *Fundamentals of convex analysis*. Springer Science & Business Media, 2012.
- Lee, K., Choi, S., and Oh, S. Sparse markov decision processes with causal sparse tsallis entropy regularization for reinforcement learning. *IEEE Robotics and Automation Letters*, 3(3):1466–1473, 2018.
- Levine, S. Reinforcement Learning and Control as Probabilistic Inference: Tutorial and Review. *arXiv preprint arXiv:1805.00909*, 2018.
- Martins, A. and Astudillo, R. From softmax to sparsemax: A sparse model of attention and multi-label classification. In *International Conference on Machine Learning (ICML)*, 2016.
- Mensch, A. and Blondel, M. Differentiable dynamic programming for structured prediction and attention. In *International Conference on Machine Learning (ICML)*, 2018.
- 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 (ICML)*, 2016.
- Morgenstern, O. and Von Neumann, J. *Theory of games and economic behavior*. Princeton university press, 1953.
- Munos, R., Stepleton, T., Harutyunyan, A., and Bellemare, M. Safe and efficient off-policy reinforcement learning. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 1054–1062, 2016.
- Nachum, O., Norouzi, M., Xu, K., and Schuurmans, D. Bridging the gap between value and policy based reinforcement learning. In *Advances in Neural Information Processing Systems (NIPS)*, 2017.
- Nachum, O., Chow, Y., and Ghavamzadeh, M. Path consistency learning in tsallis entropy regularized mdps. *arXiv preprint arXiv:1802.03501*, 2018.
- Nemirovski, A. Prox-method with rate of convergence $o(1/t)$ for variational inequalities with lipschitz continuous monotone operators and smooth convex-concave

- saddle point problems. *SIAM Journal on Optimization*, 15(1):229–251, 2004.
- Nesterov, Y. Smooth minimization of non-smooth functions. *Mathematical programming*, 103(1):127–152, 2005.
- Nesterov, Y. Primal-dual subgradient methods for convex problems. *Mathematical programming*, 120(1):221–259, 2009.
- Neu, G., Jonsson, A., and Gómez, V. A unified view of entropy-regularized Markov decision processes. *arXiv preprint arXiv:1705.07798*, 2017.
- Ng, A. Y., Harada, D., and Russell, S. Policy invariance under reward transformations: Theory and application to reward shaping. In *International Conference on Machine Learning (ICML)*, 1999.
- Perolat, J., Scherrer, B., Piot, B., and Pietquin, O. Approximate dynamic programming for two-player zero-sum markov games. In *International Conference on Machine Learning (ICML)*, 2015.
- Peters, J., Mulling, K., and Altun, Y. Relative entropy policy search. In *AAAI Conference on Artificial Intelligence*, 2010.
- Puterman, M. L. and Shin, M. C. Modified policy iteration algorithms for discounted markov decision problems. *Management Science*, 24(11):1127–1137, 1978.
- Richemond, P. H. and Maginnis, B. A short variational proof of equivalence between policy gradients and soft q learning. *arXiv preprint arXiv:1712.08650*, 2017.
- Riedmiller, M., Hafner, R., Lampe, T., Neunert, M., Degrave, J., Wiele, T., Mnih, V., Heess, N., and Springenberg, J. T. Learning by playing solving sparse reward tasks from scratch. In *International Conference on Machine Learning (ICML)*, pp. 4341–4350, 2018.
- Scherrer, B., Ghavamzadeh, M., Gabillon, V., Lesner, B., and Geist, M. Approximate modified policy iteration and its application to the game of tetris. *Journal of Machine Learning Research (JMLR)*, 16:1629–1676, 2015.
- Schulman, J., Levine, S., Abbeel, P., Jordan, M., and Moritz, P. Trust region policy optimization. In *International Conference on Machine Learning (ICML)*, 2015.
- Schulman, J., Chen, X., and Abbeel, P. Equivalence between policy gradients and soft q-learning. *arXiv preprint arXiv:1704.06440*, 2017.
- Sutton, R. S., McAllester, D. A., Singh, S. P., and Mansour, Y. Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems (NIPS)*, 2000.
- Williams, R. J. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.
- Ziebart, B. D., Maas, A. L., Bagnell, J. A., and Dey, A. K. Maximum entropy inverse reinforcement learning. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2008.

This appendices provide the proofs for all stated results (Appx. A to C) and discuss in more details the perspectives mentioned in Sec. 5 (Appx. D).

A. Proofs of section 3

In this section, we prove the results of Sec. 3. We start with the properties of the regularized Bellman operators.

Proof of Proposition 2. We can write $T_{\pi,\Omega}v = r_\pi - \Omega(\pi) + \gamma P_\pi v$, it is obviously affine (in v). Then, we show that the operators are monotonous. For the evaluation operator, we have

$$v_1 \geq v_2 \Rightarrow T_\pi v_1 \geq T_\pi v_2 \Leftrightarrow T_{\pi,\Omega}v_1 = T_\pi v_1 - \Omega(\pi) \geq T_\pi v_2 - \Omega(\pi) = T_{\pi,\Omega}v_2.$$

For the optimality operator, we have

$$\begin{aligned} v_1 \geq v_2 &\Rightarrow \forall s \in \mathcal{S}, \quad q_{s,1} \geq q_{s,2} \\ &\Rightarrow \forall s \in \mathcal{S}, \quad \Omega^*(q_{s,1}) \geq \Omega^*(q_{s,2}) && \text{by Prop. 1} \\ &\Leftrightarrow \forall s \in \mathcal{S}, \quad [T_{*,\Omega}v_1](s) \geq [T_{*,\Omega}v_2](s) \\ &\Leftrightarrow T_{*,\Omega}v_1 \geq T_{*,\Omega}v_2. \end{aligned}$$

Then, we show the distributivity property. For the evaluation operator, we have

$$T_{\pi,\Omega}(v + c\mathbf{1}) = T_\pi(v + c\mathbf{1}) - \Omega(\pi) = T_\pi v + \gamma c\mathbf{1} - \Omega(\pi) = T_{\pi,\Omega}v + \gamma c\mathbf{1}.$$

For the optimality operator, for any $s \in \mathcal{S}$, we have

$$\begin{aligned} [T_{*,\Omega}(v + c\mathbf{1})](s) &= \Omega^*(q_s + \gamma c\mathbf{1}) \\ &= \Omega^*(q_s) + \gamma c && \text{by Prop. 1} \\ &= [T_{*,\Omega}v](s) + \gamma c. \end{aligned}$$

Lastly, we study the contraction of both operators. For the evaluation operator, we have

$$T_{\pi,\Omega}v_1 - T_{\pi,\Omega}v_2 = T_\pi v_1 - \Omega(\pi) - (T_\pi v_2 - \Omega(\pi)) = T_\pi v_1 - T_\pi v_2.$$

So, the contraction is the same as the one of the unregularized operator. For the optimality operator, we have that

$$\|T_{*,\Omega}v_1 - T_{*,\Omega}v_2\|_\infty = \max_{s \in \mathcal{S}} |[T_{*,\Omega}v_1](s) - [T_{*,\Omega}v_2](s)|.$$

Pick $s \in \mathcal{S}$, and without loss of generality assume that $[T_{*,\Omega}v_1](s) \geq [T_{*,\Omega}v_2](s)$. Write also $\pi_1 = \mathcal{G}_\Omega(v_1)$ and $\pi_2 = \mathcal{G}_\Omega(v_2)$. We have

$$\begin{aligned} |[T_{*,\Omega}v_1](s) - [T_{*,\Omega}v_2](s)| &= [T_{*,\Omega}v_1](s) - [T_{*,\Omega}v_2](s) \\ &= [T_{\pi_1,\Omega}v_1](s) - [T_{\pi_2,\Omega}v_2](s) \\ &\leq [T_{\pi_1,\Omega}v_1](s) - [T_{\pi_1,\Omega}v_2](s) && \text{as } T_{*,\Omega}v_2 = T_{\pi_2,\Omega}v_2 \geq T_{\pi_1,\Omega}v_2, \\ &\leq \gamma \|v_1 - v_2\|_\infty. \end{aligned}$$

The stated result follows immediately. □

Then, we show that in a regularized MDP, the policy greedy respectively to the optimal value function is indeed the optimal policy, and is unique.

Proof of Theorem 1. The uniqueness of $\pi_{*,\Omega}$ is a consequence of the strong convexity of Ω , see Prop. 1. On the other hand, by definition of the greediness, we have

$$\pi_{*,\Omega} = \mathcal{G}_\Omega(v_{*,\Omega}) \Leftrightarrow T_{\pi_{*,\Omega},\Omega}v_{*,\Omega} = T_{*,\Omega}v_{*,\Omega} = v_{*,\Omega}.$$

This proves that $v_{*,\Omega}$ is the value function of $\pi_{*,\Omega}$. Next, for any function $v \in \mathbb{R}^S$ and any policy π , we have

$$T_{*,\Omega}v \geq T_{\pi,\Omega}v.$$

Using monotonicity, we have that

$$T_{*,\Omega}^2v = T_{*,\Omega}(T_{*,\Omega}v) \geq T_{*,\Omega}(T_{\pi,\Omega}v) \geq T_{\pi,\Omega}(T_{\pi,\Omega}v) = T_{\pi,\Omega}^2v.$$

By direct induction, for any $n \geq 1$, $T_{*,\Omega}^n v \geq T_{\pi,\Omega}^n v$. Taking the limit as $n \rightarrow \infty$, we conclude that $v_{*,\Omega} \geq v_{\pi,\Omega}$. \square

Next, we relate regularized and unregularized value functions (for a given policy, and for the optimal value function).

Proof of Proposition 3. We start by linking (un)regularized values of a given policy. Let $v \in \mathbb{R}^S$. As $T_{\pi,\Omega}v = T_{\pi}v - \Omega(\pi)$, we have that

$$T_{\pi}v - U_{\Omega}\mathbf{1} \leq T_{\pi,\Omega}v \leq T_{\pi}v - L_{\Omega}\mathbf{1}.$$

We work on the left inequality first. We have

$$T_{\pi,\Omega}^2v = T_{\pi,\Omega}(T_{\pi,\Omega}v) \geq T_{\pi,\Omega}(T_{\pi}v - U_{\Omega}\mathbf{1}) \geq T_{\pi}(T_{\pi}v - U_{\Omega}\mathbf{1}) - U_{\Omega}\mathbf{1} = T_{\pi}^2v - \gamma U_{\Omega}\mathbf{1} - U_{\Omega}\mathbf{1}.$$

By direct induction, for any $n \geq 1$,

$$T_{\pi,\Omega}^n v \geq T_{\pi}^n v - \sum_{k=0}^{n-1} \gamma^k U_{\Omega}\mathbf{1}.$$

Taking the limit as $n \rightarrow \infty$ we obtain

$$v_{\pi,\Omega} \geq v_{\pi} - \frac{U_{\Omega}}{1-\gamma}\mathbf{1}.$$

The proof is similar for the right inequality. Next, we link the (un)regularized optimal values. As a direct corollary of Prop. 1, for any $v \in \mathbb{R}^S$ we have

$$T_*v - U_{\Omega}\mathbf{1} \leq T_{*,\Omega}v \leq T_*v - L_{\Omega}\mathbf{1}.$$

Then, the proof is the same as above, switching evaluation and optimality operators. \square

Lastly, we show how good is the optimal policy of the regularized MDP for the original problem (unregularized MDP).

Proof of Theorem 2. The right inequality is obvious, as for any π , $v_* \geq v_{\pi}$. For the left inequality,

$$\begin{aligned} v_* &\leq v_{*,\Omega} + \frac{U_{\Omega}}{1-\gamma} && \text{by Prop. 3} \\ &= v_{\pi_{*,\Omega},\Omega} + \frac{U_{\Omega}}{1-\gamma} && \text{by Thm. 1} \\ &\leq v_{\pi_{*,\Omega}} + \frac{U_{\Omega}}{1-\gamma} - \frac{L_{\Omega}}{1-\gamma} && \text{by Prop. 3.} \end{aligned}$$

\square

B. Proofs of section 4

The results of section 4 do not need to be proven. Indeed, we have shown in Sec. 3 that all involved quantities of regularized MDPs satisfy the same properties as their unregularized counterpart. Therefore, the proofs of these results are identical to the proofs provided by Scherrer et al. (2015), up to the replacement of value functions, Bellman operators, and so on, by their regularized counterparts. The proofs for Mirror Descent Modified Policy Iteration (Sec. 5) are less straightforward.

C. Proofs of section 5

As a prerequisite of Lemma 1, we need the following result.

Lemma 3 (Three-point identity). *Let π be any policy, we have that*

$$\langle \nabla \Omega(\pi_k) - \nabla \Omega(\pi_{k+1}), \pi - \pi_{k+1} \rangle = D_\Omega(\pi || \pi_{k+1}) - D_\Omega(\pi || \pi_k) + D_\Omega(\pi_{k+1} || \pi_k).$$

Proof. This is the classical three-point identity of Bregman divergences, and can be checked by calculus:

$$\begin{aligned} D_\Omega(\pi || \pi_{k+1}) - D_\Omega(\pi || \pi_k) + D_\Omega(\pi_{k+1} || \pi_k) &= \Omega(\pi) - \Omega(\pi_{k+1}) - \langle \nabla \Omega(\pi_{k+1}), \pi - \pi_{k+1} \rangle \\ &\quad - \Omega(\pi) + \Omega(\pi_k) + \langle \nabla \Omega(\pi_k), \pi - \pi_k \rangle \\ &\quad + \Omega(\pi_{k+1}) - \Omega(\pi_k) - \langle \nabla \Omega(\pi_k), \pi_{k+1} - \pi_k \rangle \\ &= \langle \nabla \Omega(\pi_k) - \nabla \Omega(\pi_{k+1}), \pi - \pi_{k+1} \rangle. \end{aligned}$$

□

Now, we can prove the key lemma of MD-MPI.

Proof of Lemma 1. Let J_k be as defined in Def. 5,

$$J_k(\pi) = \langle -q_k, \pi \rangle + D_\Omega(\pi || \pi_k) = -T_{\pi, \Omega_{\pi_k}} v_k,$$

and let $\pi_{k+1} \in \mathcal{G}_{\Omega_{\pi_k}}^{\epsilon'_{k+1}}(v_k)$, that is $\langle \nabla J_k(\pi_{k+1}), \pi - \pi_{k+1} \rangle + \epsilon'_{k+1} \geq 0$. By convexity of J_k , for any policy π , we have

$$\begin{aligned} J_k(\pi) - J_k(\pi_{k+1}) &\geq \langle \nabla J_k(\pi_{k+1}), \pi - \pi_{k+1} \rangle && \text{by convexity of } J_k \\ &\geq -\epsilon'_{k+1} && \text{as } \pi_{k+1} \in \mathcal{G}_{\Omega_{\pi_k}}^{\epsilon'_{k+1}}(v_k) \\ \Leftrightarrow -T_{\pi, \Omega_{\pi_k}} v_k + T_{\pi_{k+1}, \Omega_{\pi_k}} v_k &\geq -\epsilon'_{k+1} && \text{using } J_k(\pi) = -T_{\pi, \Omega_{\pi_k}} v_k \\ \Leftrightarrow T_{\pi, \Omega_{\pi_k}} v_k - T_{\pi_{k+1}, \Omega_{\pi_k}} v_k &\leq \epsilon'_{k+1} \end{aligned}$$

This is the first result stated in Lemma 1.

Next, we relate (un)regularized quantities. We start with the following decomposition

$$\langle -q_k, \pi \rangle = \langle -q_k, \pi_{k+1} \rangle + \langle -q_k, \pi - \pi_{k+1} \rangle. \quad (7)$$

Taking the gradient of J_k (by using the definition of the Bregman divergence), we get

$$\begin{aligned} \nabla J_k(\pi_{k+1}) &= -q_k + \nabla \Omega(\pi_{k+1}) - \nabla \Omega(\pi_k) \\ \Leftrightarrow -q_k &= \nabla J_k(\pi_{k+1}) + \nabla \Omega(\pi_k) - \nabla \Omega(\pi_{k+1}). \end{aligned} \quad (8)$$

Injecting Eq. (8) into Eq. (7), we get

$$\begin{aligned} \langle -q_k, \pi \rangle &= \langle -q_k, \pi_{k+1} \rangle + \langle \nabla J_k(\pi_{k+1}) + \nabla \Omega(\pi_k) - \nabla \Omega(\pi_{k+1}), \pi - \pi_{k+1} \rangle \\ &= \langle -q_k, \pi_{k+1} \rangle + \underbrace{\langle \nabla J_k(\pi_{k+1}), \pi - \pi_{k+1} \rangle}_{\geq -\epsilon'_{k+1}} + \underbrace{\langle \nabla \Omega(\pi_k) - \nabla \Omega(\pi_{k+1}), \pi - \pi_{k+1} \rangle}_{\text{three-point identity (Lemma 3)}} \\ &\geq \langle -q_k, \pi_{k+1} \rangle - \epsilon'_{k+1} + D_\Omega(\pi || \pi_{k+1}) - D_\Omega(\pi || \pi_k) + D_\Omega(\pi_{k+1} || \pi_k) \\ \Leftrightarrow T_\pi v_k &\leq T_{\pi_{k+1}} v_k + \epsilon'_{k+1} + D_\Omega(\pi || \pi_k) - D_\Omega(\pi || \pi_{k+1}) - D_\Omega(\pi_{k+1} || \pi_k), \end{aligned} \quad (9)$$

where we used in the last inequality the fact that $\langle q_k, \pi \rangle = T_\pi v_k$. From the definition of the regularized Bellman operator, we have that $T_{\pi_{k+1}} v_k - D_\Omega(\pi_{k+1} || \pi_k) = T_{\pi_{k+1}, \Omega_{\pi_k}} v_k$, so Eq. (9) is equivalent to the second result of Lemma 1:

$$T_\pi v_k \leq T_{\pi_{k+1}, \Omega_{\pi_k}} v_k + \epsilon'_{k+1} + D_\Omega(\pi || \pi_k) - D_\Omega(\pi || \pi_{k+1}).$$

As the Bregman divergence is positive, $-D_\Omega(\pi || \pi_{k+1}) \leq 0$, and thus Eq. (9) implies the last result of Lemma 1:

$$T_\pi v_k \leq T_{\pi_{k+1}} v_k + \epsilon'_{k+1} + D_\Omega(\pi || \pi_k) - D_\Omega(\pi || \pi_{k+1}).$$

This concludes the proof. □

Next, we prove the bounds for b_k^h , s_k^h and d_k^h , for $h = 1, 2$ (if the bounds are the same, the proofs differ).

Proof of Lemma 2. We start by bounding the quantities for MD-MPI type 1. First, we consider the Bellman residual:

$$\begin{aligned}
 b_k^1 &= v_k - T_{\pi_{k+1}, \Omega_{\pi_k}} v_k \\
 &\stackrel{(a)}{=} v_k - T_{\pi_k} v_k + T_{\pi_k, \Omega_{\pi_k}} v_k - T_{\pi_{k+1}, \Omega_{\pi_k}} v_k \\
 &\stackrel{(b)}{\leq} v_k - T_{\pi_k, \Omega_{\pi_{k-1}}} v_k + \epsilon'_{k+1} \\
 &= v_k - \epsilon_k - T_{\pi_k, \Omega_{\pi_{k-1}}} v_k + \gamma P_{\pi_k} \epsilon_k + \epsilon_k - \gamma P_{\pi_k} \epsilon_k + \epsilon'_{k+1} \\
 &\stackrel{(c)}{=} (v_k - \epsilon_k) - T_{\pi_k, \Omega_{\pi_{k-1}}} (v_k - \epsilon_k) + x_k \\
 &\stackrel{(d)}{=} (T_{\pi_k, \Omega_{\pi_{k-1}}})^m v_{k-1} - T_{\pi_k, \Omega_{\pi_{k-1}}} (T_{\pi_k, \Omega_{\pi_{k-1}}})^m v_{k-1} + x_k \\
 &= (T_{\pi_k, \Omega_{\pi_{k-1}}})^m v_{k-1} - (T_{\pi_k, \Omega_{\pi_{k-1}}})^m T_{\pi_k, \Omega_{\pi_{k-1}}} v_{k-1} + x_k \\
 &= (\gamma P_{\pi_k})^m (v_{k-1} - T_{\pi_k, \Omega_{\pi_{k-1}}} v_{k-1}) + x_k \\
 &= (\gamma P_{\pi_k})^m b_{k-1}^1 + x_k.
 \end{aligned}$$

In the previous equations, we used the following facts:

- (a) $T_{\pi_k} v_k = T_{\pi_k, \Omega_{\pi_k}} v_k$ as $\Omega_{\pi_k}(\pi_k) = 0$.
- (b) We used two facts. First, $T_{\pi_k} v_k \geq T_{\pi_k} v_k - \Omega_{\pi_{k-1}}(\pi_k) = T_{\pi_k, \Omega_{\pi_{k-1}}} v_k$, as $\Omega_{\pi_{k-1}}(\pi_k) \geq 0$. Second, by Lemma 1, $T_{\pi_k, \Omega_{\pi_k}} v_k - T_{\pi_{k+1}, \Omega_{\pi_k}} v_k \leq \epsilon'_{k+1}$.
- (c) We used two facts. First, generally speaking, we have $T_{\pi, \Omega}(v_1 + v_2) = T_{\pi, \Omega} v_1 + \gamma P_{\pi} v_2$ (as $T_{\pi, \Omega}$ is affine). Second, by definition $x_k = (I - \gamma P_{\pi_k}) \epsilon_k + \epsilon'_{k+1}$.
- (d) By definition, $v_k = (T_{\pi_k, \Omega_{\pi_{k-1}}})^m v_{k-1} + \epsilon_k$.

Next, we bound the shift s_k^1 :

$$\begin{aligned}
 s_k^1 &= (T_{\pi_k, \Omega_{\pi_{k-1}}})^m v_{k-1} - v_{\pi_k} \\
 &= (T_{\pi_k, \Omega_{\pi_{k-1}}})^m v_{k-1} - v_{\pi_k, \Omega_{\pi_{k-1}}} + v_{\pi_k, \Omega_{\pi_{k-1}}} - v_{\pi_k} \\
 &\stackrel{(a)}{\leq} (T_{\pi_k, \Omega_{\pi_{k-1}}})^m v_{k-1} - v_{\pi_k, \Omega_{\pi_{k-1}}} \\
 &\stackrel{(b)}{=} (T_{\pi_k, \Omega_{\pi_{k-1}}})^m v_{k-1} - (T_{\pi_k, \Omega_{\pi_{k-1}}})^\infty v_{k-1} \\
 &= (T_{\pi_k, \Omega_{\pi_{k-1}}})^m v_{k-1} - (T_{\pi_k, \Omega_{\pi_{k-1}}})^m (T_{\pi_k, \Omega_{\pi_{k-1}}})^\infty v_{k-1} \\
 &= (\gamma P_{\pi_k})^m (v_{k-1} - (T_{\pi_k, \Omega_{\pi_{k-1}}})^\infty v_{k-1}) \\
 &= (\gamma P_{\pi_k})^m \sum_{j=0}^{\infty} ((T_{\pi_k, \Omega_{\pi_{k-1}}})^j v_{k-1} - (T_{\pi_k, \Omega_{\pi_{k-1}}})^{j+1} v_{k-1}) \\
 &= (\gamma P_{\pi_k})^m \sum_{j=0}^{\infty} ((T_{\pi_k, \Omega_{\pi_{k-1}}})^j v_{k-1} - (T_{\pi_k, \Omega_{\pi_{k-1}}})^j T_{\pi_k, \Omega_{\pi_{k-1}}} v_{k-1}) \\
 &= (\gamma P_{\pi_k})^m \sum_{j=0}^{\infty} (\gamma P_{\pi_k})^j (v_{k-1} - T_{\pi_k, \Omega_{\pi_{k-1}}} v_{k-1}) \\
 &= (\gamma P_{\pi_k})^m (I - \gamma P_{\pi_k})^{-1} b_{k-1}^1.
 \end{aligned}$$

In the previous equations, we used the following facts:

(a) Generally speaking, if $\Omega \geq 0$ then $v_{\pi, \Omega} - v_{\pi} \leq 0$ (see Prop. 3).

(b) With a slight abuse of notation, for any $v \in \mathbb{R}^S$, $v_{\pi, \Omega} = T_{\pi, \Omega}^{\infty} v$.

Then we bound the distance d_k^1 :

$$\begin{aligned}
 d_{k+1}^1 &= v_* - (T_{\pi_{k+1}, \Omega_{\pi_k}})^m v_k \\
 &= T_{\pi_*} v_* - T_{\pi_*} v_k + T_{\pi_*} v_k - T_{\pi_{k+1}, \Omega_{\pi_k}} v_k + T_{\pi_{k+1}, \Omega_{\pi_k}} - (T_{\pi_{k+1}, \Omega_{\pi_k}})^m v_k \\
 &= \gamma P_{\pi_*} (v_* - v_k) + T_{\pi_*} v_k - T_{\pi_{k+1}, \Omega_{\pi_k}} v_k + T_{\pi_{k+1}, \Omega_{\pi_k}} - (T_{\pi_{k+1}, \Omega_{\pi_k}})^m v_k \\
 &\stackrel{(a)}{\leq} \gamma P_{\pi_*} (v_* - v_k) + \epsilon'_{k+1} + \delta_k(\pi_*) + T_{\pi_{k+1}, \Omega_{\pi_k}} - (T_{\pi_{k+1}, \Omega_{\pi_k}})^m v_k \\
 &\stackrel{(b)}{=} \gamma P_{\pi_*} (v_* - v_k) + \epsilon'_{k+1} + \delta_k(\pi_*) + \sum_{j=1}^{m-1} (\gamma P_{\pi_{k+1}})^j b_k^1 \\
 &= \gamma P_{\pi_*} (v_* - v_k) + \gamma P_{\pi_*} \epsilon_k - \gamma P_{\pi_*} \epsilon_k + \epsilon'_{k+1} + \delta_k(\pi_*) + \sum_{j=1}^{m-1} (\gamma P_{\pi_{k+1}})^j b_k^1 \\
 &\stackrel{(c)}{=} \gamma P_{\pi_*} d_k^1 + y_k + \delta_k(\pi_*) + \sum_{j=1}^{m-1} (\gamma P_{\pi_{k+1}})^j b_k^1.
 \end{aligned}$$

In the previous equations, we used the following facts:

(a) By Lemma 1,

$$T_{\pi_*} v_k - T_{\pi_{k+1}, \Omega_{\pi_k}} v_k \leq \epsilon'_{k+1} + \delta_k(\pi_*).$$

(b) For this step, we used:

$$\begin{aligned}
 T_{\pi_{k+1}, \Omega_{\pi_k}} - (T_{\pi_{k+1}, \Omega_{\pi_k}})^m v_k &= \sum_{j=1}^{m-1} ((T_{\pi_{k+1}, \Omega_{\pi_k}})^j v_k - (T_{\pi_{k+1}, \Omega_{\pi_k}})^{j+1} v_k) \\
 &= \sum_{j=1}^{m-1} ((T_{\pi_{k+1}, \Omega_{\pi_k}})^j v_k - (T_{\pi_{k+1}, \Omega_{\pi_k}})^j T_{\pi_{k+1}, \Omega_{\pi_k}} v_k) \\
 &= \sum_{j=1}^{m-1} (\gamma P_{\pi_{k+1}})^j (v_k - T_{\pi_{k+1}, \Omega_{\pi_k}} v_k) \\
 &= \sum_{j=1}^{m-1} (\gamma P_{\pi_{k+1}})^j b_k^1
 \end{aligned}$$

(c) By definition of $d_k^1 = v_* - (v_k - \epsilon_k)$ and $y_k = -\gamma P_{\pi_*} \epsilon_k + \epsilon'_{k+1}$.

The proofs for the quantities involved in MD-MPI type 2 are similar, even if their definition differ. First, we consider the Bellman residual

$$\begin{aligned}
 b_k^2 &= v_k - T_{\pi_{k+1}} v_k \\
 &\stackrel{(a)}{\leq} v_k - T_{\pi_{k+1}} v_k + \Omega_{\pi_k}(\pi_{k+1}) = v_k - T_{\pi_{k+1}, \Omega_{\pi_k}} v_k \\
 &\stackrel{(b)}{=} v_k - T_{\pi_k} v_k + T_{\pi_k, \Omega_{\pi_k}} v_k - T_{\pi_{k+1}, \Omega_{\pi_k}} v_k \\
 &\stackrel{(c)}{\leq} v_k - T_{\pi_k} v_k + \epsilon'_{k+1} \\
 &= v_k - \epsilon_k - T_{\pi_k} v_k + \gamma P_{\pi_k} \epsilon_k + \epsilon_k - \gamma P_{\pi_k} \epsilon_k + \epsilon'_{k+1} \\
 &\stackrel{(d)}{=} (v_k - \epsilon_k) - T_{\pi_k} v_k (v_k - \epsilon_k) + x_k \\
 &\stackrel{(e)}{=} (T_{\pi_k})^m v_{k-1} - T_{\pi_k} (T_{\pi_k})^m v_{k-1} + x_k = (\gamma P_{\pi_k})^m (v_{k-1} - T_{\pi_k} v_{k-1}) + x_k = (\gamma P_{\pi_k})^m b_{k-1}^2 + x_k.
 \end{aligned}$$

In the previous equations, we used the following facts:

- (a) This is because $\Omega_{\pi_k}(\pi_{k+1}) \geq 0$ and by definition of $T_{\pi, \Omega} v = T_{\pi} v - \Omega(\pi)$.
- (b) It uses the fact that $T_{\pi_k} v_k = T_{\pi_k, \Omega_{\pi_k}} v_k$ (as $\Omega_{\pi_k}(\pi_k) = 0$).
- (c) This is by Lemma 1.
- (d) This is by definition of $x_k = \epsilon_k - \gamma P_{\pi_k} \epsilon_k + \epsilon'_{k+1}$.
- (e) This is by definition of $v_k = (T_{\pi_k})^m v_{k-1} + \epsilon_k$.

Then, we bound the shift s_k^2 , the technique being the same as before:

$$\begin{aligned}
 s_k^2 &= (T_{\pi_k})^m v_{k-1} - v_{\pi_k} \\
 &= (T_{\pi_k})^m v_{k-1} - (T_{\pi_k})^{m+\infty} v_{k-1} \\
 &= (\gamma P_{\pi_k})^m \sum_{j=0}^{\infty} ((T_{\pi_k})^j v_{k-1} - (T_{\pi_k})^{j+1} v_{k-1}) \\
 &= (\gamma P_{\pi_k})^m \sum_{j=0}^{\infty} (\gamma P_{\pi_k})^j (v_{k-1} - T_{\pi_k} v_{k-1}) \\
 &= (\gamma P_{\pi_k})^m (I - \gamma P_{\pi_k})^{-1} b_{k-1}^2.
 \end{aligned}$$

To finish with, we prove the bound on the distance

$$\begin{aligned}
 d_{k+1}^2 &= v_* - (T_{\pi_{k+1}})^m v_k \\
 &= \underbrace{T_{\pi_*} v_* - T_{\pi_*} v_k}_{=\gamma P_{\pi_*} (v_* - v_k)} + \underbrace{T_{\pi_*} v_k - T_{\pi_{k+1}} v_k}_{\leq \epsilon'_{k+1} + \delta_k(\pi_*) \text{ by Lemma 1}} + \underbrace{T_{\pi_{k+1}} v_k - (T_{\pi_{k+1}} v_k)^m}_{=\sum_{j=1}^{m-1} (\gamma P_{\pi_{k+1}})^j b_k^2} \\
 &= \underbrace{\gamma P_{\pi_*} (v_* - v_k) + \gamma P_{\pi_*} \epsilon_k}_{=\gamma P_{\pi_*} (v_* - (v_k - \epsilon_k)) = \gamma P_{\pi_*} d_k^2} - \underbrace{\gamma P_{\pi_*} \epsilon_k + \epsilon'_{k+1}}_{=y_k} + \delta_k(\pi_*) + \sum_{j=1}^{m-1} (\gamma P_{\pi_{k+1}})^j b_k^2 \\
 &= \gamma P_{\pi_*} d_k^2 + y_k + \sum_{j=1}^{m-1} (\gamma P_{\pi_{k+1}})^j b_k^2 + \delta_k(\pi_*).
 \end{aligned}$$

□

Now, we will show the component-wise bound on the regret L_k of Thm. 4

Proof of Theorem 4. The proof is similar to the one of AMPI (Scherrer et al., 2015, Lemma 2), up to the additional term $\delta_k(\pi_*)$ and to the different bounded quantity. We will make use of the notation Γ , defined in Def. 4, and we will write Γ_* if only the stochastic kernel induced by the optimal policy π_* is involved. In other words, we write $\Gamma_*^j = (\gamma P_{\pi_*})^j$.

From Lemma 2, we have that

$$d_{k+1}^h \leq \Gamma_* d_k^h + y_k + \sum_{j=1}^{m-1} \Gamma^j b_k^h + \delta_k(\pi_*).$$

As the bound is the same for $h = 1, 2$, we remove the superscript (and reintroduce it only when necessary, that is when going back to the core definition of these quantities). By direct induction, we get

$$d_k \leq \sum_{j=0}^{k-1} \Gamma_*^{k-1-j} \left(y_j + \sum_{l=1}^{m-1} \Gamma^l b_j + \delta_j(\pi_*) \right) + \Gamma^k d_0.$$

Therefore, the loss l_k can be bounded as (defining \mathcal{L}_k at the same time)

$$l_k = d_k + s_k \leq \underbrace{\sum_{j=0}^{k-1} \Gamma_*^{k-1-j} (y_j + \sum_{l=1}^{m-1} \Gamma^l b_j) + \Gamma^k d_0 + s_k}_{=\mathcal{L}_k} + \sum_{j=0}^{k-1} \Gamma_*^{k-1-j} \delta_j(\pi_*).$$

The loss \mathcal{L}_k is exactly of the same form as the one of AMPI, we'll take advantage of this later. From this bound on the loss l_k , we can bound the regret as:

$$L_k = \sum_{k=1}^K l_k \leq \sum_{k=1}^K \mathcal{L}_k + \sum_{k=1}^K \sum_{j=0}^{k-1} \Gamma_*^{k-1-j} \delta_j(\pi_*). \quad (10)$$

We will first work on the last double sum. For this, we define $\Delta_j(\pi_*) = \sum_{k=0}^j \delta_k(\pi_*)$. We have

$$\begin{aligned} \sum_{k=1}^K \sum_{j=0}^{k-1} \Gamma_*^{k-1-j} \delta_j(\pi_*) &= \sum_{k=0}^{K-1} \sum_{j=0}^k \Gamma_*^{k-j} \delta_j(\pi_*) \\ &= \sum_{k=0}^{K-1} \sum_{j=0}^k \Gamma_*^j \delta_{k-j}(\pi_*) \\ &= \sum_{j=0}^{K-1} \sum_{k=j}^{K-1} \Gamma_*^j \delta_{k-j}(\pi_*) \\ &= \sum_{j=0}^{K-1} \Gamma_*^j \sum_{k=0}^{K-1-j} \delta_k(\pi_*) = \sum_{j=0}^{K-1} \Gamma_*^j \Delta_{K-1-j}(\pi_*). \end{aligned}$$

For the last line, we used the fact that Γ_* only involves the P_{π_*} transition kernel, that does not depend on any iteration. Now, we can bound the term $\Delta_k(\pi_*)$, for any $k \geq 0$ as follows. Let write $R_{\Omega_{\pi_0}} = \|\sup_{\pi} D_{\Omega}(\pi || \pi_0)\|_{\infty}$, we have

$$\Delta_k(\pi_*) = \sum_{j=0}^k \delta_j(\pi_*) = \sum_{j=0}^k (D_{\Omega}(\pi_* || \pi_j) - D_{\Omega}(\pi_* || \pi_{j+1})) = D_{\Omega}(\pi_* || \pi_0) - D_{\Omega}(\pi_* || \pi_{k+1}) \leq D_{\Omega}(\pi_* || \pi_0) \leq R_{\Omega_{\pi_0}} \mathbf{1}.$$

Given the definition of Γ , we have that $\Gamma^j \mathbf{1} = \gamma^j \mathbf{1}$, so

$$\sum_{k=1}^K \sum_{j=0}^{k-1} \Gamma_*^{k-1-j} \delta_j(\pi_*) \leq \sum_{j=0}^{K-1} \Gamma_*^j R_{\Omega_{\pi_0}} \mathbf{1} = \sum_{j=0}^{K-1} \gamma^j \mathbf{1} R_{\Omega_{\pi_0}} = \frac{1 - \gamma^K}{1 - \gamma} R_{\Omega_{\pi_0}} \mathbf{1}. \quad (11)$$

Next, we work on the term \mathcal{L}_k . As stated before, it is exactly the same as the one of the AMPI analysis, and the proof of [Scherrer et al. \(2015, Lemma 2\)](#) applies almost readily, we do not repeat it fully here. The only difference that appears and that induces a slight modification of the bound (on \mathcal{L}_k) is how b_0 and d_0 are related, that will modify the η_k term of the original proof. We will link b_0 and d_0 for both types of MD-MPI.

For MD-MPI type 1 (with the natural convention that $\epsilon_0 = 0$), we have that

$$b_0^1 = v_0 - T_{\pi_1, \Omega_{\pi_0}} v_0 \text{ and } d_0^1 = v_* - (v_0 - \epsilon_0) = v_* - v_0.$$

The Bellman residual can be written as

$$\begin{aligned} b_0^1 &= v_0 - T_{\pi_1, \Omega_{\pi_0}} v_0 \\ &= v_0 - v_* + \underbrace{T_{\pi_*} v_* - T_{\pi_*} v_0}_{=(I - \gamma P_{\pi_*})(-d_0^1)} + \underbrace{T_{\pi_*} v_0 - T_{\pi_1, \Omega_{\pi_0}} v_0}_{\leq \epsilon'_1 + \delta_0(\pi_*) \text{ by Lemma 1}} \\ &\leq (I - \gamma P_{\pi_*})(-d_0^1) + \epsilon'_1 + D_{\Omega}(\pi_* || \pi_0) \\ &\leq (I - \gamma P_{\pi_*})(-d_0^1) + \epsilon'_1 + R_{\Omega_{\pi_0}} \mathbf{1}, \end{aligned}$$

where we used in the penultimate line the fact that $\delta_0(\pi_*) = D_\Omega(\pi_*||\pi_0) - D_\Omega(\pi_*||\pi_1) \leq D_\Omega(\pi_*||\pi_0)$ and in the last line the same bounding as before. So, the link between b_0^1 and d_0^1 is the same as for AMPI, up to the additional $R_{\Omega_{\pi_0}} \mathbf{1}$ term.

For MD-MPI type 2, we have (with the same convention)

$$b_0^2 = v_0 - T_{\pi_1} v_0 \text{ and } d_0^2 = v_* - (v_0 - \epsilon_0) = v_* - v_0.$$

Working on the Bellman residual

$$\begin{aligned} b_0^2 &= v_0 - T_{\pi_1} v_0 \\ &= \underbrace{v_0 - v_* + T_{\pi_*} v_* - T_{\pi_*} v_0}_{=(I - \gamma P_{\pi_*})(-d_0^2)} + \underbrace{T_{\pi_*} v_0 - T_{\pi_1} v_0}_{\leq \epsilon'_1 + \delta_0(\pi_*) \text{ by Lemma 1}} \\ &\leq (I - \gamma P_{\pi_*})(-d_0^2) + \epsilon'_1 + D_\Omega(\pi_*||\pi_0) \leq (I - \gamma P_{\pi_*})(-d_0^2) + \epsilon'_1 + R_{\Omega_{\pi_0}} \mathbf{1}. \end{aligned}$$

So, we have the same bound.

Combining this with the part of the proof that does not change, and that we do not repeat, we get the following bound on \mathcal{L}_k :

$$\mathcal{L}_k \leq 2 \sum_{i=1}^{k-1} \sum_{j=i}^{\infty} \Gamma^j |\epsilon_{k-i}| + \sum_{i=0}^{k-1} \sum_{j=i}^{\infty} \Gamma^j |\epsilon'_{k-i}| + h(k) + \underbrace{\sum_{j=k}^{\infty} \Gamma^j \mathbf{1} R_{\Omega_{\pi_0}}}_{= \frac{\gamma^k}{1-\gamma} \mathbf{1} R_{\Omega_{\pi_0}}}, \quad (12)$$

with $h(k)$ being defined as

$$h(k) = 2 \sum_{j=k}^{\infty} \Gamma^j |d_0| \text{ or } h(k) = 2 \sum_{j=k}^{\infty} \Gamma^j |b_0|.$$

Combining Eqs (11) and (12) into Eq. (10), we can bound the regret:

$$\begin{aligned} L_K &\leq \sum_{k=1}^K \mathcal{L}_k + \sum_{k=1}^K \sum_{j=0}^{k-1} \Gamma_*^{k-1-j} \delta_j(\pi_*) \\ &\leq \sum_{k=1}^K \left(2 \sum_{i=1}^{k-1} \sum_{j=i}^{\infty} \Gamma^j |\epsilon_{k-i}| + \sum_{i=0}^{k-1} \sum_{j=i}^{\infty} \Gamma^j |\epsilon'_{k-i}| + h(k) + \frac{\gamma^k}{1-\gamma} \mathbf{1} R_{\Omega_{\pi_0}} \right) + \frac{1-\gamma^K}{1-\gamma} R_{\Omega_{\pi_0}} \mathbf{1} \\ &= 2 \sum_{k=2}^K \sum_{i=1}^{k-1} \sum_{j=i}^{\infty} \Gamma^j |\epsilon_{k-i}| + \sum_{k=1}^K \sum_{i=0}^{k-1} \sum_{j=i}^{\infty} \Gamma^j |\epsilon'_{k-i}| + \sum_{k=1}^K h(k) + \frac{1-\gamma^K}{(1-\gamma)^2} R_{\Omega_{\pi_0}} \mathbf{1}. \end{aligned}$$

This concludes the proof. □

To prove Cor. 2, we will need a result from Scherrer et al. (2015), that we recall first.

Lemma 4 (Lemma 6 of Scherrer et al. (2015)). *Let \mathcal{I} and $(\mathcal{J}_i)_{i \in \mathcal{I}}$ be sets of non-negative integers, $\{\mathcal{I}_1, \dots, \mathcal{I}_n\}$ be a partition of \mathcal{I} , and f and $(g_i)_{i \in \mathcal{I}}$ be functions satisfying*

$$|f| \leq \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}_i} \Gamma^j |g_i| = \sum_{l=1}^n \sum_{i \in \mathcal{I}_l} \sum_{j \in \mathcal{J}_i} \Gamma^j |g_i|.$$

Then, for all p, q such that $\frac{1}{p} + \frac{1}{q} = 1$ and for all distributions ρ and μ , we have

$$\|f\|_{p, \rho} \leq \sum_{l=1}^n (\mathcal{C}_q(l))^{\frac{1}{p}} \sup_{i \in \mathcal{I}_l} \|g_i\|_{pq', \mu} \sum_{i \in \mathcal{I}_l} \sum_{j \in \mathcal{J}_i} \gamma^j$$

with the following concentrability coefficients,

$$\mathcal{C}_q(l) = \frac{\sum_{i \in \mathcal{I}_l} \sum_{j \in \mathcal{J}_i} \gamma^j c_q(j)}{\sum_{i \in \mathcal{I}_l} \sum_{j \in \mathcal{J}_i} \gamma^j}, \text{ where } c_q(j) = \max_{\pi_1, \dots, \pi_j} \left\| \frac{\rho P_{\pi_1} P_{\pi_2} \dots P_{\pi_j}}{\mu} \right\|_{q, \mu}.$$

Now, we can prove the stated result.

Proof of Corollary 2. The proof is an application of Lemma 4 to Thm. 4. We define $\mathcal{I} = \{1, 2, \dots, K^2 + K + 1\}$ and the associated trivial partition (that is, $\mathcal{I}_i = \{i\}$). For each $i \in \mathcal{I}$ we define

$$g_i = \begin{cases} 2\epsilon_{k-i'} & \text{if } i = i' + \frac{(k-2)(k-1)}{2} \text{ with } 2 \leq k \leq K \text{ and } 1 \leq i' \leq k-1 \\ \epsilon'_{k-i'} & \text{if } i = i' + \frac{k(k-1)}{2} + \frac{K(K-1)}{2} + 1 \text{ with } 1 \leq k \leq K \text{ and } 0 \leq i' \leq k-1 \\ 2d_0 \text{ or } 2b_0 & \text{if } i = K^2 + k \text{ with } 1 \leq k \leq K \\ \frac{1-\gamma^K}{(1-\gamma)^2} R_{\Omega_{\pi_0}} \mathbf{1} & \text{if } i = K^2 + K + 1 \end{cases}$$

$$\text{and } \mathcal{J}_i = \begin{cases} \{i', \dots\} & \text{if } i = i' + \frac{(k-2)(k-1)}{2} \text{ with } 2 \leq k \leq K \text{ and } 1 \leq i' \leq k-1 \\ \{i', \dots\} & \text{if } i = i' + \frac{k(k-1)}{2} + \frac{K(K-1)}{2} + 1 \text{ with } 1 \leq k \leq K \text{ and } 0 \leq i' \leq k-1 \\ \{k, \dots\} & \text{if } i = K^2 + k \text{ with } 1 \leq k \leq K \\ \{0\} & \text{if } i = K^2 + K + 1 \end{cases}.$$

With this, Thm. 4 rewrites as

$$L_K \leq \sum_{l=1}^{K^2+K+1} \sum_{i \in \mathcal{I}_l} \sum_{j \in \mathcal{J}_i} \Gamma^j |g_i|.$$

The results follows by applying Lemma 4 and using the fact that $\sum_{j \geq i} \gamma^j = \frac{\gamma^i}{1-\gamma}$ □

The proof of Prop. 4 is a basic application of the Hölder inequality.

Proof of Proposition 4. We write \circ the Hadamard product. Recall that $l_k = v_* - v_{\pi_k} \geq 0$ and that $L_K = \sum_{k=1}^K l_k$. Notice that the sequence v_{π_k} is not necessarily monotone, even without approximation error. On one side, we have that

$$\|L_K\|_{1,\rho} = \rho L_K = \sum_{k=1}^K \rho l_k \geq K \min_{1 \leq k \leq K} \rho l_k = K \min_{1 \leq k \leq K} \|l_k\|_{1,\rho}.$$

On the other side, with q such that $\frac{1}{p} + \frac{1}{q} = 1$ we have that

$$\rho L_K = \langle \rho^{\frac{1}{p} + \frac{1}{q}}, L_K \rangle = \langle \rho^{\frac{1}{q}}, \rho^{\frac{1}{p}} \circ L_K \rangle \leq \|\rho^{\frac{1}{q}}\|_q \|\rho^{\frac{1}{p}} \circ L_K\|_p = \|L_K\|_{p,\rho},$$

where we used the Hölder inequality. Combining both equations provides the stated result. □

Cor. 3 is a direct consequence of Cor. 2.

Proof of Corollary 3. Taking the limit of Cor. 2 as $p \rightarrow \infty$, when the errors are null, gives

$$\|L_K\|_{\infty} \leq 2 \sum_{k=1}^K \frac{\gamma^k}{1-\gamma} \min(\|d_0\|_{\infty}, \|b_0\|_{\infty}) + \frac{1-\gamma^K}{(1-\gamma)^2} R_{\Omega_{\pi_0}} \leq 2\gamma \frac{1-\gamma^K}{(1-\gamma)^2} \|d_0\|_{\infty} + \frac{1-\gamma^K}{(1-\gamma)^2} R_{\Omega_{\pi_0}},$$

where we used the fact that $\sum_{k=1}^K \gamma^k = \gamma \frac{1-\gamma^K}{1-\gamma}$. The result follows by grouping terms and using $d_0 = v_* - v_0$. □

The proof of Cor. 4 is mainly a manipulation of sums.

Proof of Corollary 4. From Cor. 2, we have that

$$\|L_K\|_{p,\rho} \leq 2 \sum_{k=2}^K \sum_{i=1}^{k-1} \frac{\gamma^i}{1-\gamma} (C_q^i)^{\frac{1}{p}} \|\epsilon_{k-i}\|_{pq',\mu} + \sum_{k=1}^K \sum_{i=0}^{k-1} \frac{\gamma^i}{1-\gamma} (C_q^i)^{\frac{1}{p}} \|\epsilon'_{k-i}\|_{pq',\mu} + g(k) + \frac{1-\gamma^K}{(1-\gamma)^2} R_{\Omega_{\pi_0}}.$$

We only need to work on the first two sums. To shorten the notations, write $\alpha_i = \frac{\gamma^i}{1-\gamma} (C_q^i)^{\frac{1}{p}}$, $\beta_i = \|\epsilon_{k-i}\|_{pq', \mu}$ and $\beta'_i = \|\epsilon'_{k-i}\|_{pq', \mu}$. We have:

$$\begin{aligned}
 \sum_{k=2}^K \sum_{i=1}^{k-1} \alpha_i \beta_{k-i} + \sum_{k=1}^K \sum_{i=0}^{k-1} \alpha_i \beta'_{k-i} &= \sum_{k=1}^{K-1} \sum_{i=1}^k \alpha_i \beta_{k+1-i} + \sum_{k=0}^{K-1} \sum_{i=0}^k \alpha_i \beta'_{k+1-i} \\
 &= \sum_{i=1}^{K-1} \alpha_i \sum_{k=i}^{K-1} \beta_{k+1-i} + \sum_{i=0}^{K-1} \alpha_i \sum_{k=i}^{K-1} \beta'_{k+1-i} \\
 &= \sum_{i=1}^{K-1} \alpha_i \sum_{k=1}^{K-i} \beta_k + \sum_{i=0}^{K-1} \alpha_i \sum_{k=1}^{K-i} \beta'_k.
 \end{aligned}$$

The result follows by reinjecting this in Cor. 2 after having replaced α_i , β_i and β'_i . □

The proof of Cor. 5 basically consists in the application of Lemma 4 to a rewriting of Thm. 4.

Proof of Corollary 5. By Thm. 4, we have

$$L_K \leq 2 \sum_{k=2}^K \sum_{i=1}^{k-1} \sum_{j=i}^{\infty} \Gamma^j |\epsilon_{k-i}| + \sum_{k=1}^K \sum_{i=0}^{k-1} \sum_{j=i}^{\infty} \Gamma^j |\epsilon'_{k-i}| + \sum_{k=1}^K h(k) + \frac{1-\gamma^K}{(1-\gamma)^2} R_{\Omega_{\pi_0}} \mathbf{1}. \quad (13)$$

We start by rewriting the two first sums. For the first one, we have

$$\begin{aligned}
 \sum_{k=2}^K \sum_{i=1}^{k-1} \sum_{j=i}^{\infty} \Gamma^j |\epsilon_{k-i}| &= \sum_{k=2}^K \sum_{i=1}^{k-1} \sum_{j=k-i}^{\infty} \Gamma^j |\epsilon_i| \\
 &= \sum_{k=1}^{K-1} \sum_{i=1}^k \sum_{j=k+1-i}^{\infty} \Gamma^j |\epsilon_i| \\
 &= \sum_{i=1}^{K-1} \left(\sum_{k=i}^{K-1} \sum_{j=k+1-i}^{\infty} \Gamma^j \right) |\epsilon_i| \\
 &= \sum_{i=1}^{K-1} \left(\sum_{k=1}^{K-i} \sum_{j=k}^{\infty} \Gamma^j \right) |\epsilon_i| \\
 &= \sum_{i=1}^{K-1} \left(\sum_{k=1}^i \sum_{j=k}^{\infty} \Gamma^j \right) |\epsilon_{K-i}|.
 \end{aligned}$$

Similarly, we have that

$$\sum_{k=1}^K \sum_{i=0}^{k-1} \sum_{j=i}^{\infty} \Gamma^j |\epsilon'_{k-i}| = \sum_{i=0}^{K-1} \left(\sum_{k=0}^i \sum_{j=k}^{\infty} \Gamma^j \right) |\epsilon'_{K-i}|.$$

Thus, the bound on the loss can be written as

$$L_K \leq 2 \sum_{i=1}^{K-1} \left(\sum_{k=1}^i \sum_{j=k}^{\infty} \Gamma^j \right) |\epsilon_{K-i}| + \sum_{i=0}^{K-1} \left(\sum_{k=0}^i \sum_{j=k}^{\infty} \Gamma^j \right) |\epsilon'_{K-i}| + 2 \left(\sum_{k=1}^K \sum_{j=k}^{\infty} \Gamma^j \right) (|d_0| \text{ or } |b_0|) + \frac{1-\gamma^K}{(1-\gamma)^2} R_{\Omega_{\pi_0}} \mathbf{1}.$$

In order to apply Lemma 4 to this bound, we consider $\mathcal{I} = \{1, 2, \dots, 2K+1\}$ and the associated trivial partition $\mathcal{I}_i = \{i\}$.

For each $i \in \mathcal{I}$, we define:

$$g_i = \begin{cases} 2\epsilon_{K-i} & \text{if } 1 \leq i \leq K-1 \\ \epsilon'_{2K-i} & \text{if } K \leq i \leq 2K-1 \\ 2|d_0| \text{ or } 2|b_0| & \text{if } i = 2K \\ \frac{1-\gamma^K}{(1-\gamma)^2} R_{\Omega_{\pi_0}} \mathbf{1} & \text{if } i = 2K+1 \end{cases}$$

$$\text{and } \mathcal{J}_i = \begin{cases} \bigcup_{k=1}^i \{k, k+1, \dots\} & \text{if } 1 \leq i \leq K-1 \\ \bigcup_{k=0}^{i-K} \{k, k+1, \dots\} & \text{if } K \leq i \leq 2K-1 \\ \bigcup_{k=1}^K \{k, k+1, \dots\} & \text{if } i = 2K \\ \{0\} & \text{if } i = 2K+1 \end{cases}.$$

With this, Eq. (13) rewrites as

$$L_K \leq \sum_{l=1}^{2K+1} \sum_{i \in \mathcal{I}_l} \sum_{j \in \mathcal{J}_i} \Gamma^j |g_i|.$$

With the c_q term defined in Lemma 4, using the fact that $\sum_{i=l}^{k-1} \sum_{j=i}^{\infty} \gamma^j = \frac{\gamma^l - \gamma^k}{(1-\gamma)^2}$, as well as the following concentrability coefficient,

$$C_q^{l,k} = \frac{(1-\gamma)^2}{\gamma^l - \gamma^k} \sum_{i=l}^{k-1} \sum_{j=i}^{\infty} c_q(j),$$

we get the following bound on the regret:

$$\begin{aligned} \|L_K\|_{p,\rho} &\leq 2 \sum_{i=1}^{K-1} \frac{\gamma - \gamma^{i+1}}{(1-\gamma)^2} (C_q^{1,i+1})^{\frac{1}{p}} \|\epsilon_{k-i}\|_{pq',\mu} + \sum_{i=0}^{K-1} \frac{1 - \gamma^{i+1}}{(1-\gamma)^2} (C_q^{0,i+1})^{\frac{1}{p}} \|\epsilon'_{k-i}\|_{pq',\mu} \\ &\quad + \frac{\gamma - \gamma^{K+1}}{(1-\gamma)^2} (C_q^{1,K+1})^{\frac{1}{p}} \min(\|d_0\|_{pq',\mu}, \|b_0\|_{pq',\mu}) + \frac{1 - \gamma^K}{(1-\gamma)^2} R_{\Omega_{\pi_0}}. \end{aligned}$$

□

D. Perspectives on regularized MDPs

Here, we discuss in more details the perspectives briefly mentioned in Sec. 6.

D.1. Dynamic programming and optimization

We have shown how MPI regularized by a Bregman divergence is related to Mirror Descent. The computation of the regularized greedy policy is similar to a mirror descent step, where the q -function plays the role of the negative subgradient. In this sense, the policy lives in the primal space while the q -function lives in the dual space. It would be interesting to take inspiration from proximal convex optimization to derive new dynamic programming approaches, for example based on Dual Averaging (Nesterov, 2009) or Mirror Prox (Nemirovski, 2004).

For example, consider the case of a fixed regularizer. We have seen that it leads to a different solution than the original one (see Thm. 2). Usually, one considers a scaled negative entropy as such a regularizer⁵, $\Omega(\pi(\cdot|s)) = \alpha \sum_a \pi(a|s) \ln \pi(a|s)$. The choice of this parameter is important practically and problem-dependent: too high and the solution of the regularized problem will be very different from the original one, too low and the algorithm will not benefit from the regularization. A natural idea is to vary the weight of the regularizer over iterations (Peters et al., 2010; Abdolmaleki et al., 2018a; Haarnoja et al., 2018b), much like a learning rate in a gradient descent approach.

More formally, in our framework, write $\Omega_k = \alpha_k \Omega$ and consider the following weighted reg-MPI scheme,

$$\begin{cases} \pi_{k+1} = \mathcal{G}_{\Omega_k}^{\epsilon'_{k+1}}(v_k) \\ v_{k+1} = (T_{\pi_{k+1}, \Omega_k})^m v_k + \epsilon_{k+1} \end{cases},$$

⁵This is the case for SAC or soft Q-learning, for example. Sometimes, it is the reward that is scaled, but both are equivalent.

with $\mathcal{G}_{\Omega_k}^{\epsilon'_{k+1}}(v_k)$ as defined in Sec. 4.2: for any policy π , $T_{\pi, \Omega_k} v_k \leq T_{\pi_{k+1}, \Omega_k} v_k + \epsilon'_{k+1}$. By applying the proof techniques developed previously, one can obtain easily the following result.

Theorem 5. Define $R_\Omega = \|\sup_\pi \Omega(\pi)\|_\infty$. Assume that the series $(\alpha_k)_{k \geq 0}$ is positive and decreasing, and that the regularizer Ω is positive (without loss of generality). After K iterations of the preceding weighted reg-MPI scheme, the regret satisfies

$$L_K \leq 2 \sum_{k=2}^K \sum_{i=1}^{k-1} \sum_{j=i}^{\infty} \Gamma^j |\epsilon_{k-i}| + \sum_{k=1}^K \sum_{i=0}^{k-1} \sum_{j=i}^{\infty} \Gamma^j |\epsilon'_{k-i}| + \sum_{k=1}^K h(k) + \frac{1-\gamma^K}{(1-\gamma)^2} R_\Omega \sum_{k=0}^{K-1} \alpha_k \mathbf{1}.$$

with $h(k) = 2 \sum_{j=k}^{\infty} \Gamma^j |d_0|$ or $h(k) = 2 \sum_{j=k}^{\infty} \Gamma^j |b_0|$.

Proof. The proof is similar to the one of MD-MPI, as it bounds the regret, but also simpler as it does not require Lemma 1. The principle is still to bound the distance d_k and the shift s_k , that both require bounding the Bellman residual b_k . From this, one can bound the loss $l_k = d_k + s_k$, and then the regret $L_K = \sum_{k=1}^K l_k$. We only specify what changes compared to the previous proofs.

Bounding the shift s_k is done as before, and one get the same bound:

$$s_k = (T_{\pi_k, \Omega_{k-1}})^m v_{k-1} - v_{\pi_k} \leq (\gamma P_{\pi_k})^m (I - \gamma P_{\pi_k}) b_{k-1}.$$

For the distance d_k , we have

$$\begin{aligned} d_{k+1} &= v_* - (T_{\pi_{k+1}, \Omega_k})^m v_k \\ &= \underbrace{T_{\pi_*} v_* - T_{\pi_*} v_k}_{=\gamma P_{\pi_*}(v_* - v_k)} + \underbrace{T_{\pi_*} v_k - T_{\pi_*, \Omega_k} v_k}_{=\Omega_k(\pi_*) \leq \alpha_k R_\Omega \mathbf{1}} + \underbrace{T_{\pi_*, \Omega_k} v_k - T_{\pi_{k+1}, \Omega_k} v_k}_{\leq \epsilon'_{k+1}} + T_{\pi_{k+1}, \Omega_k} v_k - (T_{\pi_{k+1}, \Omega_k})^m v_k. \end{aligned}$$

The rest of the bounding is similar to MD-MPI, and we get

$$d_{k+1} \leq \gamma P_{\pi_*} d_k + y_k + \alpha_k R_\Omega \mathbf{1} + \sum_{j=1}^{m-1} (\gamma P_{\pi_{k+1}})^j b_k.$$

So, this is similar to the bound of MD-MPI, with the term $\alpha_k R_\Omega \mathbf{1}$ replacing the term $\delta_k(\pi_*)$. For the Bellman residual, we have

$$b_k = v_k - T_{\pi_{k+1}, \Omega_k} v_k = \underbrace{v_k - T_{\pi_k, \Omega_k} v_k}_{\leq v_k - T_{\pi_k, \Omega_{k-1}} v_k} + \underbrace{T_{\pi_k, \Omega_k} v_k - T_{\pi_{k+1}, \Omega_k} v_k}_{\epsilon'_{k+1}}$$

where we used the facts that $\alpha_k \leq \alpha_{k-1}$ and that $\Omega \geq 0$ for bounding the first term:

$$-T_{\pi_k, \Omega_k} v_k = -T_{\pi_k} v_k + \alpha_k \Omega(\pi_k) \leq -T_{\pi_k} v_k + \alpha_{k-1} \Omega(\pi_k) = -T_{\pi_k, \Omega_{k-1}} v_k.$$

The rest of the bounding is as before and gives $b_k \leq (\gamma P_{\pi_k})^m b_{k-1} + x_k$. From these bounds, one can bound L_k as previously. \square

From Thm. 5, we can obtain an ℓ_p -bound on the regret, and from this the rate of convergence of the average regret. For MD-MPI, the rate of convergence was in $\mathcal{O}(\frac{1}{K})$. Here, it depends on the weighting of the regularizer. For example, if α_k is in $\mathcal{O}(\frac{1}{k})$, then the average regret will be in $\mathcal{O}(\frac{\ln K}{K})$. If α_k is in $\mathcal{O}(\frac{1}{\sqrt{k}})$, then the average regret will be in $\mathcal{O}(\frac{1}{\sqrt{K}})$. This illustrates the kind of things that can be done with the proposed framework of regularized MDPs.

D.2. Temporal consistency equations

Thanks to Ω , the regularized greedy policies are unique, and thus is the regularized optimal policy. The pair of optimal policy and optimal value function can be characterized as follows.

Corollary 6. The optimal policy and optimal value function in a regularized MDP are the unique functions satisfying

$$\forall s \in \mathcal{S}, \quad \begin{cases} v_{*, \Omega}(s) = \Omega^*(r(s, \cdot) + \gamma \mathbb{E}_{s'|s} [v_{*, \Omega}(s')]) \\ \pi_{*, \Omega}(\cdot | s) = \nabla \Omega^*(r(s, \cdot) + \gamma \mathbb{E}_{s'|s} [v_{*, \Omega}(s')]) \end{cases}.$$

Proof. This is a direct consequence of Thm. 1 □

This provides a general way to estimate the optimal policy. For example, if Ω is the negative entropy, this set of equations simplifies to

$$\forall (s, a) \in \mathcal{S} \times \mathcal{A} \quad v_{*,\Omega}(s) = r(s, a) + \gamma \mathbb{E}_{s'|s,a} [v_{*,\Omega}(s')] - \ln \pi_{*,\Omega}(a|s).$$

This has been used by [Nachum et al. \(2017\)](#) or [Dai et al. \(2018\)](#), where it is called “temporal consistency equation”, to estimate the optimal value-policy pair by minimizing the related residual,

$$\mathbb{E}_{s,a} \left[\left(r(s, a) + \gamma \mathbb{E}_{s'|s,a} [v_{*,\Omega}(s')] - \ln \pi_{*,\Omega}(a|s) - v_{*,\Omega}(s) \right)^2 \right].$$

This idea has also been extended to Tsallis entropy ([Nachum et al., 2018](#)). In this case, the set of equations does not simplify as nicely as with the Shannon entropy. Instead, the approach consists in considering the Lagrangian derived from the Legendre-Fenchel transform (and the resulting temporal consistency equation involves Lagrange multipliers, that have to be learnt too). This idea could be extended to other regularizers. One could also replace the regularizer Ω by a Bregman divergence, and estimate the optimal policy by solving a sequence of temporal consistency equations.

D.3. Regularized policy gradient

Policy search approaches often combine policy gradient with an entropic regularization, typically to prevent the policy from becoming too quickly deterministic ([Williams, 1992](#); [Mnih et al., 2016](#)). The policy gradient theorem ([Sutton et al., 2000](#)) can easily be extended to the proposed framework.

Let ν be a (user-defined) state distribution, the classical policy search approach consists in maximizing $J(\pi) = \mathbb{E}_{s \sim \nu} [v_\pi(s)] = \nu v_\pi$. This principle can easily be extended to regularized MDPs, by maximizing

$$J_\Omega(\pi) = \nu v_{\pi,\Omega}.$$

Write $d_{\nu,\pi}$ the γ -weighted occupancy measure induced by the policy π when the initial state is sampled from ν , defined as $d_{\nu,\pi} = (1 - \gamma)\nu(I - \gamma P_\pi)^{-1} \in \Delta_{\mathcal{S}}$. Slightly abusing notations, we’ll also write $d_{\nu,\pi}(s, a) = d_{\nu,\pi}(s)\pi(a|s)$.

Theorem 6 (Policy gradient for regularized MDPs). *The gradient of J_Ω is*

$$\nabla J_\Omega(\pi) = \frac{1}{1 - \gamma} \mathbb{E}_{s,a \sim d_{\nu,\pi}} \left[\left(q_{\pi,\Omega}(s, a) - \frac{\partial \Omega(\pi(\cdot|s))}{\partial \pi(a|s)} \right) \nabla \ln \pi(a|s) \right].$$

Proof. We have that

$$\nabla J_\Omega(\pi) = \sum_{s \in \mathcal{S}} \nu(s) \nabla v_{\pi,\Omega}(s).$$

We have to study the gradient of the value function. For this, the nabla-log trick is useful: $\nabla \pi = \pi \nabla \ln \pi$.

$$\begin{aligned} \nabla v_{\pi,\Omega}(s) &= \nabla \left(\sum_{a \in \mathcal{A}} \pi(a|s) (r(s, a) + \gamma \mathbb{E}_{s'|s,a} [v_{\pi,\Omega}(s')]) - \Omega(\pi(\cdot|s)) \right) \\ &= \sum_{a \in \mathcal{A}} (\nabla \pi(a|s) q_{\pi,\Omega}(s, a) + \pi(a|s) \gamma \mathbb{E}_{s'|s,a} [\nabla v_{\pi,\Omega}(s')]) - \nabla \Omega(\pi(\cdot|s)) \\ &= \sum_{a \in \mathcal{A}} \pi(a|s) (q_{\pi,\Omega}(s, a) \nabla \ln \pi(a|s) - \nabla \Omega(\pi(\cdot|s)) + \gamma \mathbb{E}_{s'|s,a} [\nabla v_{\pi,\Omega}(s')]). \end{aligned}$$

So, the components of $\nabla v_{\pi,\Omega}$ are the (unregularized) value functions corresponding to the rewards being the components of $q_{\pi,\Omega}(s, a) \nabla \ln \pi(a|s) - \nabla \Omega(\pi(\cdot|s))$. Consequently,

$$J_\Omega(\pi) = \frac{1}{1 - \gamma} \sum_{s \in \mathcal{S}} d_{\nu,\pi}(s) \sum_{a \in \mathcal{A}} \pi(a|s) (q_{\pi,\Omega}(s, a) \nabla \ln \pi(a|s) - \nabla \Omega(\pi(\cdot|s))).$$

Using the chain-rule, we have that

$$\nabla \Omega(\pi(\cdot|s)) = \sum_{a \in \mathcal{A}} \frac{\partial \Omega(\pi(\cdot|s))}{\partial \pi(a|s)} \nabla \pi(a|s) = \sum_{a \in \mathcal{A}} \frac{\partial \Omega(\pi(\cdot|s))}{\partial \pi(a|s)} \pi(a|s) \nabla \ln \pi(a|s).$$

Injecting this in the previous result concludes the proof. \square

Even in the entropic case, it might be not exactly the same as the usual regularized policy gradient, notably because our result involves the regularized q -function. Again, the regularizer Ω could be replaced by a Bregman divergence, to ultimately estimate the optimal policy of the original MDP. It would be interesting to compare empirically the different resulting policy gradients approaches, we left this for future work.

D.4. Regularized inverse reinforcement learning

Inverse reinforcement learning (IRL) consists in finding a reward function that explains the behavior of an expert which is assumed to act optimally. It is often said that it is an ill-posed problem. The classical example is the null reward function that explains any behavior (as all policies are optimal in this case).

We argue that in this regularized framework, the problem is not ill-posed, because the optimal policy is unique, thanks to the regularization. For example, if one consider the negative entropy as the regularizer, with a null reward, the optimal policy will be that of maximum entropy, so the uniform policy, and it is unique.

Notice that if for a reward, the associated regularized optimal policy is unique, the converse is not true. For example, the uniform policy is optimal for any constant reward. More generally, reward shaping (Ng et al., 1999) still holds true for regularized MDPs (this being thanks to the results of Sec. 3, again).

This being said, assume that the model (dynamic, discount factor, regularizer) and that the optimal regularized policy $\pi_{*,\Omega}$ are known. It is possible to retrieve a reward function such that $\pi_{*,\Omega}$ is the unique optimal policy.

Proposition 5. *Let $\hat{q} \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$ be any function satisfying*

$$\forall s \in \mathcal{S}, \quad \pi_{*,\Omega}(\cdot|s) = \nabla \Omega^*(\hat{q}(s, \cdot)),$$

then the reward $\hat{r}(s, a)$ defined as

$$\begin{aligned} \forall (s, a) \in \mathcal{S} \times \mathcal{A}, \quad \hat{r}(s, a) &= \hat{q}(s, a) - \gamma \mathbb{E}_{s'|s,a} [\Omega^*(\hat{q}(s, \cdot))] \\ &= \hat{q}(s, a) - \gamma \mathbb{E}_{s'|s,a} [\mathbb{E}_{a' \sim \pi_{*,\Omega}(\cdot|s')} [\hat{q}(s, a')] - \Omega(\pi_{*,\Omega}(\cdot|s))] \end{aligned}$$

has $\pi_{,\Omega}$ as the unique corresponding optimal policy.*

Proof. First, recall (see Prop. 1) that if for any q , there is a unique regularized greedy policy, the converse is not true (simply by the fact that for any $v \in \mathbb{R}^{\mathcal{S}}$, we have $\nabla \Omega^*(q(s, \cdot) + v(s)) = \nabla \Omega^*(q(s, \cdot))$). By assumption and by uniqueness of regularized greediness, $\pi_{*,\Omega}$ is the unique regularized policy corresponding to \hat{q} . Then, with the above defined reward function, \hat{q} is unique the solution of the regularized Bellman optimality equation. This shows the stated result. \square

If this result shows that IRL is well-defined in a regularized framework, it is not very practical (for example, in the entropic case, it tells that $\hat{r}(s, a) = \ln \pi_{*,\Omega}(s, a)$ is such a reward function). Yet, we think that the proposed general framework could lead to more practical algorithm.

For example, many IRL algorithms are based on the maximum-entropy principle, eg. (Ziebart et al., 2008; Finn et al., 2016; Fu et al., 2018). This maximum-entropy IRL framework can be linked to probabilistic inference, that can itself be shown to be equivalent, in some specific cases (deterministic dynamics), to entropy-regularized reinforcement learning (Levine, 2018). We think this to be an interesting connection, and maybe that our proposed regularized framework could allow to generalize or analyze some of these approaches.

D.5. Regularized zero-sum Markov games

A zero-sum Markov game can be seen as a generalization of MDPs. It is a tuple $\{\mathcal{S}, \mathcal{A}^1, \mathcal{A}^2, P, r, \gamma\}$ with \mathcal{S} the state space common to both players, \mathcal{A}^j the action space of player j , $P \in \Delta_{\mathcal{S}}^{\mathcal{S} \times \mathcal{A}^1 \times \mathcal{A}^2}$ the transition kernel ($P(s'|s, a^1, a^2)$ is the probability of transiting to state s' when player 1 played a^1 and player 2 played a^2 in s), $r \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}^1 \times \mathcal{A}^2}$ the reward function of both players (one tries to maximize it, the other one to minimize it) and γ the discount factor. We write $\mu \in \Delta_{\mathcal{A}^1}^{\mathcal{S}}$ a policy of the maximizer, and $\nu \in \Delta_{\mathcal{A}^2}^{\mathcal{S}}$ a policy of the minimizer.

As in the case of classical MDPs, everything can be constructed from an evaluation operator, defined as

$$[T_{\mu, \nu}v](s) = \sum_{a^1} \sum_{a^2} \mu(a^1|s) \nu(a^2|s) \left(r(s, a^1, a^2) + \gamma \sum_{s'} p(s'|s, a^1, a^2) v(s') \right),$$

of fixed point $v_{\mu, \nu}$. From this, the following operators are defined:

$$\begin{aligned} T_{\mu}v &= \min_{\nu} T_{\mu, \nu}v & T_{\nu}v &= \max_{\mu} T_{\mu, \nu}v \\ T v &= \max_{\mu} T_{\mu}v & \hat{T}v &= \min_{\nu} T_{\nu}v. \end{aligned}$$

We also define the greedy operator as $\mu \in \mathcal{G}(v) \Leftrightarrow T v = T_{\mu}v = \min_{\nu} T_{\mu, \nu}v$. Thanks to the Von Neumann's minimax theorem (Morgenstern & Von Neumann, 1953), we have that

$$\hat{T}v = T v$$

and the optimal value function satisfies

$$v_* = \min_{\nu} \max_{\mu} v_{\mu, \nu} = \max_{\mu} \min_{\nu} v_{\mu, \nu}.$$

The modified policy iteration for this kind of games is

$$\begin{cases} \mu_{k+1} \in \mathcal{G}(v_k) \\ v_{k+1} = T_{\mu_{k+1}}^m v_k \end{cases} . \quad (14)$$

As in MDPs, we can regularize the evaluation operator, and construct regularized zero-sum Markov games from this. Let Ω_1 and Ω_2 be two strongly convex regularizers on $\Delta_{\mathcal{A}^1}$ and $\Delta_{\mathcal{A}^2}$, and define the regularized evaluation operator as

$$[T_{\mu, \nu, \Omega}v](s) = [T_{\mu, \nu}v](s) - \Omega_1(\mu(\cdot|s)) + \Omega_2(\nu(\cdot|s)).$$

From this, we can construct a theory of regularized zero-sum Markov games as it was done for MDPs. The Von Neumann's minimax theorem does not only hold for affine operators, but for convex-concave operators, so we're fine.

Notably, the unregularized error propagation analysis of (14) by Perolat et al. (2015) could be easily adapted to the regularized case (much like how the analysis of AMPI directly led to the analysis of Sec. 4.2, thanks to the results of Sec. 3). We left its extension to regularization with a Bregman divergence as future work.