

A Proofs

Proof of Lemma 1. Let us start by writing the left hand side of (13) as

$$\int_{\mathcal{S} \times \mathcal{A}} |\rho(s, a) - \rho_\theta(s, a)| dsda = (1 - \gamma) \int_{\mathcal{S} \times \mathcal{A}} \left| \sum_{t=0}^{\infty} \gamma^t (p_\pi^t(s, a) - p_\theta^t(s, a)) \right| dsda \quad (24)$$

Using the triangle inequality, we upper bound the previous expression as

$$\int_{\mathcal{S} \times \mathcal{A}} |\rho(s, a) - \rho_\theta(s, a)| dsda \leq (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t \int_{\mathcal{S} \times \mathcal{A}} |p_\pi^t(s, a) - p_\theta^t(s, a)| dsda. \quad (25)$$

Notice that to complete the proof it suffices to show that the right hand side of the previous expression is bounded by $\epsilon/(1 - \gamma)$. We next work towards that end, and we start by bounding the difference $|p_\pi^t(s, a) - p_\theta^t(s, a)|$. Notice that this difference can be upper bounded using the triangle inequality as

$$|p_\pi^t(s, a) - p_\theta^t(s, a)| \leq p_\pi^t(s) |\pi(a|s) - \pi_\theta(a|s)| + \pi_\theta(a|s) |p_\pi^t(s) - p_\theta^t(s)|. \quad (26)$$

Since π_θ is an ϵ -approximation of π , it follows from Definition 1 that

$$\int_{\mathcal{S} \times \mathcal{A}} p_\pi^t(s) |\pi(a|s) - \pi_\theta(a|s)| dsda \leq \epsilon \int_{\mathcal{S}} p_\pi^t(s) ds = \epsilon, \quad (27)$$

where the last equality follows from the fact that $p_\pi^t(s)$ is a density. We next work towards bounding the integral of the second term in (26). Using the fact that $\pi_\theta(a|s)$ is a density, it follows that

$$\int_{\mathcal{S} \times \mathcal{A}} \pi_\theta(a|s) |p_\pi^t(s) - p_\theta^t(s)| dsda = \int_{\mathcal{S}} |p_\pi^t(s) - p_\theta^t(s)| ds. \quad (28)$$

Notice that the previous difference is zero for $t = 0$ and for any $t > 0$ it can be upper bounded by

$$\begin{aligned} \int_{\mathcal{S}} |p_\pi^t(s) - p_\theta^t(s)| ds &\leq \int_{\mathcal{S}} \int_{\mathcal{S} \times \mathcal{A}} p(s|s', a') |p_\pi^{t-1}(s', a') - p_\theta^{t-1}(s', a')| ds ds' da' \\ &= \int_{\mathcal{S} \times \mathcal{A}} |p_\pi^{t-1}(s', a') - p_\theta^{t-1}(s', a')| ds' da' \end{aligned} \quad (29)$$

Combining the bounds derived in (25), (27), (29) we have that

$$\begin{aligned} (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t \int_{\mathcal{S} \times \mathcal{A}} |p_\pi^t(s, a) - p_\theta^t(s, a)| dsda &\leq \\ (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t \epsilon + (1 - \gamma) \sum_{t=1}^{\infty} \gamma^t \int_{\mathcal{S} \times \mathcal{A}} |p_\pi^{t-1}(s, a) - p_\theta^{t-1}(s, a)| dsda. \end{aligned} \quad (30)$$

Notice that the first term on the right hand side of the previous expression is the sum of the geometric multiplied by $1 - \gamma$. Hence we have that $(1 - \gamma) \sum_{t=0}^{\infty} \gamma^t \epsilon = \epsilon$. The second term on the right hand side of the previous expression is in fact the same as the term on the left hand side of the expression multiplied by the discount factor γ . Thus, rearranging the terms, the previous expression implies that

$$(1 - \gamma) \sum_{t=0}^{\infty} \gamma^t \int_{\mathcal{S} \times \mathcal{A}} |p_\pi^t(s, a) - p_\theta^t(s, a)| dsda \leq \frac{\epsilon}{1 - \gamma}. \quad (31)$$

This completes the proof of the Lemma. \square

Proof of Theorem 2. Notice that the dual functions $d(\lambda)$ and $d_\theta(\lambda)$ associated to the problems (PI) and (PII) respectively are such that for every λ we have that $d_\theta(\lambda) \leq d(\lambda)$. The latter follows from the fact that the set of maximizers of the Lagrangian for the parametrized policies is contained in the set of maximizers of the non-parametrized policies. In particular, this holds for λ^* the solution of the dual problem associated to (PI). Hence we have the following sequence of inequalities

$$D^* = d(\lambda^*) \geq d_\theta(\lambda^*) \geq D_\theta^*, \quad (32)$$

where the last inequality follows from the fact that D_θ^* is the minimum of (DII). The zero duality gap established in Theorem 1 completes the proof of the upper bound for D_θ^* . We next work towards proving the lower bound for D_θ^* . Let us next write the dual function of the parametrized problem (DII) as

$$d_\theta(\lambda) = d(\lambda) - \left(\max_{\pi \in \mathcal{P}(\mathcal{S})} \mathcal{L}(\pi, \lambda) - \max_{\theta \in \mathbb{R}^p} \mathcal{L}_\theta(\theta, \lambda) \right) \quad (33)$$

Let $\pi^* \triangleq \operatorname{argmax}_{\pi \in \mathcal{P}(\mathcal{S})} \mathcal{L}(\pi, \lambda)$ and let θ^* be an ϵ -approximation of π^* . Then, by definition of the maximum it follows that

$$d_\theta(\lambda) \geq d(\lambda) - (\mathcal{L}(\pi^*, \lambda) - \mathcal{L}_\theta(\theta^*, \lambda)) \quad (34)$$

We next work towards a bound for $\mathcal{L}(\pi^*, \lambda) - \mathcal{L}_\theta(\theta^*, \lambda)$. To do so, notice that we can write the difference in terms of the occupation measures where ρ^* and ρ_θ^* are the occupation measures associated to the policies π^* and the policy π_{θ^*}

$$\mathcal{L}(\pi^*, \lambda) - \mathcal{L}_\theta(\theta^*, \lambda) = \int_{\mathcal{S} \times \mathcal{A}} (r_0 + \lambda^\top r) (d\rho^*(\lambda) - d\rho_\theta^*(\lambda)). \quad (35)$$

Since π_{θ^*} is by definition an ϵ approximation of π^* it follows from Lemma 1 that

$$\int_{\mathcal{S} \times \mathcal{A}} |d\rho^*(\lambda) - d\rho_\theta^*(\lambda)| \leq \frac{\epsilon}{1 - \gamma}. \quad (36)$$

Using the bounds on the the reward functions we can upper bound the difference $\mathcal{L}(\pi^*, \lambda) - \mathcal{L}_\theta(\theta^*, \lambda)$ by

$$\mathcal{L}(\pi^*, \lambda) - \mathcal{L}_\theta(\theta^*, \lambda) \leq (B_{r_0} + \|\lambda\|_1 B_r) \frac{\epsilon}{1 - \gamma}. \quad (37)$$

Combining the previous bound with (34) we can lower bound $d_\theta(\lambda)$ as

$$d_\theta(\lambda) \geq d(\lambda) - (B_{r_0} + \|\lambda\|_1 B_r) \frac{\epsilon}{1 - \gamma} \quad (38)$$

Let us next define $d_\epsilon(\lambda) = d(\lambda) - B_r \epsilon / (1 - \gamma) \|\lambda\|_1$, and notice that in fact $d_\epsilon(\lambda)$ is the dual function associated to Problem (PI') with $\xi_i = B_r \epsilon / (1 - \gamma)$ for all $i = 1, \dots, m$. With this definition, (38) reduces to

$$d_\theta(\lambda) \geq d_\epsilon(\lambda) - B_{r_0} \frac{\epsilon}{1 - \gamma}. \quad (39)$$

Since the previous expression holds for every λ , in particular it holds for λ_θ^* , the dual solution of the parametrized problem (DII). Thus, we have that

$$D_\theta^* \geq d_\epsilon(\lambda_\theta^*) - B_{r_0} \frac{\epsilon}{1 - \gamma}. \quad (40)$$

Recall that $\lambda_\epsilon^* = \operatorname{argmin} d_\epsilon(\lambda)$, and use the definition of the dual function to lower bound D_θ^* by

$$D_\theta^* \geq \max_{\pi \in \mathcal{P}(\mathcal{S})} V_0(\pi) + \sum_{i=1}^m \lambda_{\epsilon,i}^* \left(V_i(\pi) - c_i - B_r \frac{\epsilon}{1 - \gamma} \right) - B_{r_0} \frac{\epsilon}{1 - \gamma}. \quad (41)$$

By definition of maximum, we can lower bound the previous expression by substituting by any $\pi \in \mathcal{P}(\mathcal{S})$. In particular, we select π^* the solution to (PI)

$$D_\theta^* \geq V_0(\pi^*) + \sum_{i=1}^m \lambda_{\epsilon,i}^* \left(V_i(\pi^*) - c_i - B_r \frac{\epsilon}{1 - \gamma} \right) - B_{r_0} \frac{\epsilon}{1 - \gamma}. \quad (42)$$

Since π^* is the optimal solution to (PI) it follows that $V_i(\pi^*) - c_i \geq 0$ and since $\lambda_{\epsilon,i}^* \geq 0$ the previous expression reduces to

$$D_\theta^* \geq V_0(\pi^*) - (B_{r_0} + B_r \|\lambda_\epsilon^*\|) \frac{\epsilon}{1 - \gamma} = P^* - (B_{r_0} + B_r \|\lambda_\epsilon^*\|) \frac{\epsilon}{1 - \gamma} \quad (43)$$

Which completes the proof of th result \square

Proof of Proposition 2. Let λ_θ^* be the solution of the parametrized dual problem (DII). Then we can write the difference of the dual function evaluated at an arbitrary $\lambda \in \mathbb{R}_+^m$ and λ_θ^* as

$$d_\theta(\lambda) - d_\theta(\lambda_\theta^*) = \max_\theta \mathcal{L}_\theta(\theta, \lambda) - \max_\theta \mathcal{L}_\theta(\theta, \lambda_\theta^*) \leq \mathcal{L}_\theta(\theta^*(\lambda), \lambda) - \mathcal{L}_\theta(\theta^\dagger(\lambda), \lambda_\theta^*). \quad (44)$$

It follows from Assumption 1 that there exists $\delta > 0$ such that $\mathcal{L}_\theta(\theta^*(\lambda), \lambda) \leq \mathcal{L}_\theta(\theta^\dagger(\lambda), \lambda) + \delta$, thus we can upper bound the right hand side of the previous inequality by

$$\mathcal{L}_\theta(\theta^*(\lambda), \lambda) - \mathcal{L}_\theta(\theta^\dagger(\lambda), \lambda_\theta^*) \leq \mathcal{L}_\theta(\theta^\dagger(\lambda), \lambda) + \delta - \mathcal{L}_\theta(\theta^\dagger(\lambda), \lambda_\theta^*) = (\lambda - \lambda_\theta^*)^\top (V(\theta^\dagger(\lambda)) - c) + \delta. \quad (45)$$

Combining the two upper bounds completes the proof of the proposition. \square

Proof of Theorem 3. We start by showing the lower bound, which in fact holds for any λ . Notice that for any λ and by definition of the dual problem it follows that $d_\theta(\lambda) \geq D_\theta^*$. Combining this bound with the result of Theorem 2 it follows that

$$d_\theta(\lambda) \geq P^* - (B_{r_0} + \|\lambda_\epsilon^*\|_1 B_r) \frac{\epsilon}{1 - \gamma}. \quad (46)$$

To show the upper bound we start by writing the difference between the dual multiplier $k + 1$ and the solution of (DII) in terms of the iteration at time k . Since $\lambda_\theta^* \in \mathbb{R}_+^m$ and using the non-expansive property of the projection it follows that

$$\|\lambda_{k+1} - \lambda_\theta^*\|^2 \leq \|\lambda_k - \eta (V(\theta^\dagger(\lambda_k)) - c) - \lambda_\theta^*\|^2 \quad (47)$$

Expanding the square and using that $B = \sum_{i=1}^m (B_{r_i}/(1 - \gamma) - c_i)^2$ is a bound on the norm squared of $V(\theta) - s$ it follows that

$$\|\lambda_{k+1} - \lambda_\theta^*\|^2 \leq \|\lambda_k - \lambda_\theta^*\|^2 - 2\eta (\lambda_k - \lambda_\theta^*)^\top (V(\theta^\dagger(\lambda_k)) - c) + \eta^2 B. \quad (48)$$

Using the result of Proposition 2 we can further upper bound the inner product in the previous expression by the difference of the dual function evaluated at λ_k and λ_θ^* plus δ , the error in the solution of the primal maximization,

$$\|\lambda_{k+1} - \lambda_\theta^*\|^2 \leq \|\lambda_k - \lambda_\theta^*\|^2 + 2\eta (\delta + d_\theta(\lambda_\theta^*) - d_\theta(\lambda_k)) + \eta^2 B. \quad (49)$$

Defining $\alpha_k = 2(\delta + d_\theta(\lambda_\theta^*) - d_\theta(\lambda_k)) + \eta B$ and writing recursively the previous expression yields

$$\|\lambda_{k+1} - \lambda_\theta^*\|^2 \leq \|\lambda_0 - \lambda_\theta^*\|^2 + \eta \sum_{j=0}^k \alpha_j. \quad (50)$$

Since $d_\theta(\lambda_\theta^*)$ is the minimum of the dual function, the difference $d_\theta(\lambda_\theta^*) - d_\theta(\lambda_k)$ is always negative. Thus, when λ_k is not close to the solution of the dual problem α_k is negative. The latter implies that the distance between λ_k and λ_θ^* is reduced by virtue of (50). To be formal, for any $\epsilon > 0$, when $\alpha_j > -2\epsilon$ we have that

$$d_\theta(\lambda_j) - d_\theta(\lambda_\theta^*) \leq \eta \frac{B}{2} + \delta + \epsilon. \quad (51)$$

Using the result of Theorem 2 we can upper bound D_θ^* by P^* which establishes the neighborhood defined in (23). We are left to show that the number of iterations required to do so is bounded by

$$K \leq \frac{\|\lambda_0 - \lambda_\theta^*\|^2}{2\eta\epsilon}. \quad (52)$$

To do so, let $K > 0$ be the first iterate in the neighborhood (23). Formally, $K = \min_{j \in \mathbb{N}} \alpha_j > -2\epsilon$. Then it follows from the recursion that

$$\|\lambda_K - \lambda_\theta^*\|^2 \leq \|\lambda_0 - \lambda_\theta^*\|^2 - 2K\eta\epsilon. \quad (53)$$

Since $\|\lambda_K - \lambda_\theta^*\|^2$ is positive the previous expression reduces to $2K\eta\epsilon \leq \|\lambda_0 - \lambda_\theta^*\|^2$. Which completes the proof of the result. \square