

On the Theory of Policy Gradient Methods: Optimality, Approximation, and Distribution Shift

Alekh Agarwal* Sham M. Kakade† Jason D. Lee‡ Gaurav Mahajan§

Abstract

Policy gradient methods are among the most effective methods in challenging reinforcement learning problems with large state and/or action spaces. However, little is known about even their most basic theoretical convergence properties, including: if and how fast they converge to a globally optimal solution or how they cope with approximation error due to using a restricted class of parametric policies. This work provides provable characterizations of the computational, approximation, and sample size properties of policy gradient methods in the context of discounted Markov Decision Processes (MDPs). We focus on both: “tabular” policy parameterizations, where the optimal policy is contained in the class and where we show global convergence to the optimal policy; and parametric policy classes (considering both log-linear and neural policy classes), which may not contain the optimal policy and where we provide agnostic learning results. One central contribution of this work is in providing approximation guarantees that are average case — which avoid explicit worst-case dependencies on the size of state space — by making a formal connection to supervised learning under *distribution shift*. This characterization shows an important interplay between estimation error, approximation error, and exploration (as characterized through a precisely defined condition number).

1 Introduction

Policy gradient methods have a long history in the reinforcement learning (RL) literature [Williams, 1992, Sutton et al., 1999, Konda and Tsitsiklis, 2000, Kakade, 2001] and are an attractive class of algorithms as they are applicable to any differentiable policy parameterization; admit easy extensions to function approximation; easily incorporate structured state and action spaces; are easy to implement in a simulation based, model-free manner. Owing to their flexibility and generality, there has also been a flurry of improvements and refinements to make these ideas work robustly with deep neural network based approaches (see e.g. Schulman et al. [2015, 2017]).

*Microsoft Research, Redmond, WA 98052. Email: alekha@microsoft.com

†University of Washington, Seattle, WA 98195 & Microsoft Research. Email: sham@cs.washington.edu

‡Princeton University, Princeton, NJ 08540. Email: jasonlee@princeton.edu

§University of California San Diego, La Jolla, CA 92093. Email: gmahajan@eng.ucsd.edu

Despite the large body of empirical work around these methods, their convergence properties are only established at a relatively coarse level; in particular, the folklore guarantee is that these methods converge to a stationary point of the objective, assuming adequate smoothness properties hold and assuming either exact or unbiased estimates of a gradient can be obtained (with appropriate regularity conditions on the variance). However, this local convergence viewpoint does not address some of the most basic theoretical convergence questions, including: 1) if and how fast they converge to a globally optimal solution (say with a sufficiently rich policy class); 2) how they cope with approximation error due to using a restricted class of parametric policies; or 3) their finite sample behavior. These questions are the focus of this work.

Overall, the results of this work place policy gradient methods under a solid theoretical footing, analogous to the global convergence guarantees of iterative value function based algorithms.

1.1 Our Contributions

This work focuses on first-order and quasi second-order policy gradient methods which directly work in the space of some parameterized policy class (rather than value-based approaches). We characterize the computational, approximation, and sample size properties of these methods in the context of a discounted Markov Decision Process (MDP). We focus on: 1) *tabular policy parameterizations*, where there is one parameter per state-action pair so the policy class is complete in that it contains the optimal policy, and 2) *function approximation*, where we have a restricted class of parametric policies which may not contain the globally optimal policy. Note that policy gradient methods for discrete action MDPs work in the space of stochastic policies, which permits the policy class to be differentiable. We now discuss our contributions in the both of these contexts.

Tabular case: We consider three algorithms: two of which are first order methods, projected gradient ascent (on the simplex) and gradient ascent (with a softmax policy parameterization); and the third algorithm, natural policy gradient ascent, can be viewed as a quasi second-order method (or preconditioned first-order method). Table 1 summarizes our main results in this case: upper bounds on the number of iterations taken by these algorithms to find an ϵ -optimal policy, when we have access to exact policy gradients.

Arguably, the most natural starting point for an analysis of policy gradient methods is to consider directly doing gradient ascent on the policy simplex itself and then to project back onto the simplex if the constraint is violated after a gradient update; we refer to this algorithm as projected gradient ascent on the simplex. Using a notion of gradient domination [Polyak, 1963], our results provably show that any first-order stationary point of the value function results in an approximately optimal policy, under certain regularity assumptions; this allows for a global convergence analysis by directly appealing to standard results in the non-convex optimization literature.

A more practical and commonly used parameterization is the softmax parameterization, where the simplex constraint is explicitly enforced by the exponential parameterization, thus avoiding projections. This work provides the first global convergence guarantees using only first-order gradient information for the widely-used softmax parameterization. Our first result for this parameterization establishes the asymptotic convergence of the policy gradient algorithm; the analysis

Algorithm	Iteration complexity
Projected Gradient Ascent on Simplex (Thm 4.1)	$O\left(\frac{D_\infty^2 \mathcal{S} \mathcal{A} }{(1-\gamma)^6 \epsilon^2}\right)$
Policy Gradient, softmax parameterization (Thm 5.1)	asymptotic
Policy Gradient + log barrier regularization, softmax parameterization (Cor 5.1)	$O\left(\frac{D_\infty^2 \mathcal{S} ^2 \mathcal{A} ^2}{(1-\gamma)^6 \epsilon^2}\right)$
Natural Policy Gradient (NPG), softmax parameterization (Thm 5.3)	$\frac{2}{(1-\gamma)^2 \epsilon}$

Table 1: **Iteration Complexities with Exact Gradients for the Tabular Case:** A summary of the number of iterations required by different algorithms to find a policy π such that $V^*(s_0) - V^\pi(s_0) \leq \epsilon$ for some fixed s_0 , assuming access to *exact policy gradients*. The first three algorithms optimize the objective $\mathbb{E}_{s \sim \mu}[V^\pi(s)]$, where μ is the starting state distribution for the algorithms. The MDP has $|\mathcal{S}|$ states, $|\mathcal{A}|$ actions, and discount factor $0 \leq \gamma < 1$. The quantity $D_\infty := \max_s \left(\frac{d_{s_0}^{\pi^*}(s)}{\mu(s)}\right)$ is termed the *distribution mismatch coefficient*, where, roughly speaking, $d_{s_0}^{\pi^*}(s)$ is the fraction of time spent in state s when executing an optimal policy π^* , starting from the state s_0 (see (4)). The NPG algorithm directly optimizes $V^\pi(s_0)$ for any state s_0 . In contrast to the complexities of the previous three algorithms, NPG has no dependence on the coefficient D_∞ , nor does it depend on the choice of s_0 . Both the MDP Experts Algorithm [Even-Dar et al., 2009] and MD-MPI algorithm [Geist et al., 2019] (see Corollary 3 of their paper) also yield guarantees for the same update rule as NPG for the softmax parameterization, though at a worse rate. See Section 2 for further discussion.

challenge here is that the optimal policy (which is deterministic) is attained by sending the softmax parameters to infinity.

In order to establish a finite time, convergence rate to optimality for the softmax parameterization, we then consider a *log barrier* regularizer and provide an iteration complexity bound that is polynomial in all relevant quantities. The use of our log barrier regularizer is critical to avoiding the issue of gradients becoming vanishingly small at suboptimal near-deterministic policies, an issue of significant practical relevance. The log barrier regularizer can also be viewed as using a *relative* entropy regularizer; here, we note the general approach of entropy based regularization is common in practice (e.g. see [Williams and Peng, 1991, Mnih et al., 2016, Peters et al., 2010, Abdolmaleki et al., 2018, Ahmed et al., 2019]). One notable distinction, which we discuss later, is that our analysis is for the log barrier regularization rather than the entropy regularization.

For these aforementioned algorithms, our convergence rates depend on the optimization measure having coverage over the state space, as measured by the *distribution mismatch coefficient* D_∞ (see Table 1 caption). In particular, for the convergence rates shown in Table 1 (for the afore-

mentioned algorithms), we assume that the optimization objective is the expected (discounted) cumulative value where the initial state is sampled under some distribution, and D_∞ is a measure of the coverage of this initial distribution. Furthermore, we provide a lower bound that shows such a dependence is unavoidable for first-order methods, even when exact gradients are available.

We then consider the Natural Policy Gradient (NPG) algorithm [Kakade, 2001] (also see Bagnell and Schneider [2003], Peters and Schaal [2008]), which can be considered a quasi second-order method due to the use of its particular preconditioner, and provide an iteration complexity to achieve an ϵ -optimal policy that is at most $\frac{2}{(1-\gamma)^2\epsilon}$ iterations, improving upon the previous related results of [Even-Dar et al., 2009, Geist et al., 2019] (see Section 2). Note the convergence rate has *no* dependence on the number of states or the number of actions, nor does it depend on the distribution mismatch coefficient D_∞ . We provide a simple and concise proof for the convergence rate analysis by extending the approach developed in [Even-Dar et al., 2009], which uses a mirror descent style of analysis [Nemirovsky and Yudin, 1983, Cesa-Bianchi and Lugosi, 2006] and also handles the non-concavity of the policy optimization problem.

This fast and dimension free convergence rate shows how the variable preconditioner in the natural gradient method improves over the standard gradient ascent algorithm. The dimension free aspect of this convergence rate is worth reflecting on, especially given the widespread use of the natural policy gradient algorithm along with variants such as the Trust Region Policy Optimization (TRPO) algorithm [Schulman et al., 2015]; our results may help to provide analysis of a more general family of entropy based algorithms (see for example Neu et al. [2017]).

Function Approximation: We now summarize our results with regards to policy gradient methods in the setting where we work with a restricted policy class, which may not contain the optimal policy. In this sense, these methods can be viewed as approximate methods. Table 2 provides a summary along with the comparisons to some relevant approximate dynamic programming methods.

A long line of work in the function approximation setting focuses on mitigating the worst-case “ ℓ_∞ ” guarantees that are inherent to approximate dynamic programming methods [Bertsekas and Tsitsiklis, 1996] (see the first row in Table 2). The reason to focus on average case guarantees is that it supports the applicability of *supervised machine learning* methods to solve the underlying approximation problem. This is because supervised learning methods, like classification and regression, typically have bounds on the expected error under a distribution, as opposed to worst-case guarantees over all possible inputs.

The existing literature largely consists of two lines of provable guarantees that attempt to mitigate the explicit ℓ_∞ error conditions of approximate dynamic programming: those methods which utilize a problem dependent parameter (the concentrability coefficient [Munos, 2005]) to provide more refined dynamic programming guarantees (e.g. see Munos [2005], Szepesvári and Munos [2005], Antos et al. [2008], Farahmand et al. [2010]) and those which work with a restricted policy class, making incremental updates, such as Conservative Policy Iteration (CPI) [Kakade and Langford, 2002, Scherrer and Geist, 2014], Policy Search by Dynamic Programming (PSDP) [Bagnell et al., 2004], and MD-MPI Geist et al. [2019]. Both styles of approaches give guarantees based on worst-case density ratios, i.e. they depend on a maximum ratio between two different densities over the state space. As discussed in [Scherrer, 2014], the assumptions in the latter class of algorithms

are substantially weaker, in that the worst-case density ratio only depends on the state visitation distribution of an optimal policy (also see Table 2 caption and Section 2).

With regards to function approximation, our main contribution is in providing performance bounds that, in some cases, have milder dependence on these density ratios. We precisely quantify an *approximation/estimation* error decomposition relevant for the analysis of the natural gradient method; this decomposition is stated in terms of the *compatible function approximation error* as introduced in Sutton et al. [1999]. More generally, we quantify our function approximation results in terms of a precisely quantified transfer error notion, based on approximation error under *distribution shift*. Table 2 shows a special case of our convergence rates of NPG, which is governed by four quantities: ϵ_{stat} , ϵ_{approx} , κ , and D_∞ .

Let us discuss the important special case of log-linear policies (i.e. policies that take the softmax of linear functions in a given feature space) where the relevant quantities are as follows: ϵ_{stat} is a bound on the excess risk (the estimation error) in fitting linearly parameterized value functions, which can be driven to 0 with more samples (at the usual statistical rate of $O(1/\sqrt{N})$ where N is the number of samples); ϵ_{approx} is the usual notion of average squared approximation error where the target function may not be perfectly representable by a linear function; κ can be upper bounded with an inverse dependence on the minimal eigenvalue of the feature covariance matrix of the fitting measure (as such it can be viewed as a dimension dependent quantity but not necessarily state dependent); and D_∞ is as before.

For the realizable case, where all policies have values which are linear in the given features (such as in linear MDP models of [Jin et al., 2019, Yang and Wang, 2019, Jiang et al., 2017]), we have that the approximation error ϵ_{approx} is 0. Here, our guarantees yield a fully polynomial and sample efficient convergence guarantee, provided the condition number κ is bounded. Importantly, there always exists a good (universal) initial measure that ensures κ is bounded by a quantity that is only polynomial in the dimension of the features, d , as opposed to an explicit dependence on the size of the (infinite) state space (see Remark 6.3). Such a guarantee would not be implied by algorithms which depend on the coefficients C_∞ or D_∞ .¹

Our results are also suggestive that a broader class of incremental algorithms — such as CPI [Kakade and Langford, 2002], PSDP [Bagnell et al., 2004], and MD-MPI Geist et al. [2019] which make small changes to the policy from one iteration to the next — may also permit a sharper analysis, where the dependence of worst-case density ratios can be avoided through an appropriate approximation/estimation decomposition; this is an interesting direction for future work (a point which we return to in Section 7). One significant advantage of NPG is that the explicit parametric policy representation in NPG (and other policy gradient methods) leads to a succinct policy representation in comparison to CPI, PSDP, or related boosting-style methods [Scherrer and Geist, 2014], where the representation complexity of the policy of the latter class of methods grows linearly in the number of iterations (since these methods add one policy to the ensemble per iteration). This representation complexity is likely why the latter class of algorithms are less widely used in practice.

¹Bounding C_∞ would require a restriction on the dynamics of the MDP (see Chen and Jiang [2019] and Section 2). Bounding D_∞ would require an initial state distribution that is constructed using knowledge of π^* , through d^{π^*} . In contrast, κ can be made $O(d)$, with an initial state distribution that only depends on the geometry of the features (and does not depend on any other properties of the MDP). See Remark 6.3.

Algorithm	Suboptimality after T Iterations	Relevant Quantities
Approx. Value/Policy Iteration [Bertsekas and Tsitsiklis, 1996]	$\frac{\epsilon_\infty}{(1-\gamma)^2} + \frac{\gamma^T}{(1-\gamma)^2}$	ϵ_∞ : ℓ_∞ error of values
Approx. Policy Iteration, with concentrability [Munos, 2005, Antos et al., 2008]	$\frac{C_\infty \epsilon_1}{(1-\gamma)^2} + \frac{\gamma^T}{(1-\gamma)^2}$	ϵ_1 : an ℓ_1 average error C_∞ : concentrability (max density ratio)
Conservative Policy Iteration [Kakade and Langford, 2002] Related: PSDP [Bagnell et al., 2004], MD-MPI Geist et al. [2019]	$\frac{D_\infty \epsilon_1}{(1-\gamma)^2} + \frac{1}{(1-\gamma)\sqrt{T}}$	ϵ_1 : an ℓ_1 average error D_∞ : max density ratio to opt., $D_\infty \leq C_\infty$
Natural Policy Gradient (Cor. 6.1 and Thm. 6.2)	$\sqrt{\frac{\kappa \epsilon_{\text{stat}} + D_\infty \epsilon_{\text{approx}}}{(1-\gamma)^3}} + \frac{1}{(1-\gamma)\sqrt{T}}$	ϵ_{stat} : excess risk ϵ_{approx} : approx. error κ : a condition number D_∞ : max density ratio to opt., $D_\infty \leq C_\infty$

Table 2: **Overview of Approximate Methods:** The suboptimality, $V^*(s_0) - V^\pi(s_0)$, after T iterations for various approximate algorithms, which use different notions of approximation error (sample complexities are not directly considered but instead may be thought of as part of ϵ_1 and ϵ_{stat} . See Section 2 for further discussion). Order notation is used to drop constants, and we assume $|\mathcal{A}| = 2$ for ease of exposition. For approximate dynamic programming methods, the relevant error is the worst case, ℓ_∞ -error in approximating a value function, e.g. $\epsilon_\infty = \max_{s,a} |Q^\pi(s,a) - \widehat{Q}^\pi(s,a)|$, where \widehat{Q}^π is what an estimation oracle returns during the course of the algorithm. The second row (see Lemma 12 in Antos et al. [2008]) is a refinement of this approach, where ϵ_1 is an ℓ_1 -average error in fitting the value functions under the fitting (state) distribution μ , and, roughly, C_∞ is a worst case density ratio between the state visitation distribution of any non-stationary policy and the fitting distribution μ . For Conservative Policy Iteration, ϵ_1 is a related ℓ_1 -average case fitting error with respect to a fitting distribution μ , and D_∞ is as defined as before, in the caption of Table 1 (see also [Kakade and Langford, 2002]); here, $D_\infty \leq C_\infty$ (e.g. see Scherrer [2014]). For NPG, ϵ_{stat} and ϵ_{approx} measure the excess risk (the regret) and approximation errors in fitting the values. Roughly speaking, ϵ_{stat} is the excess squared loss relative to the best fit (among an appropriately defined parametric class) under our fitting distribution (defined with respect to the state distribution μ). Here, ϵ_{approx} is the approximation error: the minimal possible error (in our parametric class) under our fitting distribution. The condition number κ is a relative eigenvalue condition between appropriately defined feature covariances with respect to the state visitation distribution of an optimal policy, $d_{s_0}^{\pi^*}$, and the state fitting distribution μ . See text for further discussion, and Section 6 for precise statements as well as a more general result not explicitly dependent on D_∞ .

2 Related Work

We now discuss related work, roughly in the order which reflects our presentation of results in the previous section.

For the direct policy parameterization in the tabular case, we make use of a gradient domination-like property, namely any first-order stationary point of the policy value is approximately optimal up to a distribution mismatch coefficient. A variant of this result also appears in Theorem 2 of Scherrer and Geist [2014], which itself can be viewed as a generalization of the approach in Kakade and Langford [2002]. In contrast to CPI [Kakade and Langford, 2002] and the more general boosting-based approach in Scherrer and Geist [2014], we phrase this approach as a Polyak-like gradient domination property [Polyak, 1963] in order to directly allow for the transfer of any advances in non-convex optimization to policy optimization in RL. More broadly, it is worth noting the global convergence of policy gradients for Linear Quadratic Regulators [Fazel et al., 2018] also goes through a similar proof approach of gradient domination.

Empirically, the recent work of Ahmed et al. [2019] studies entropy based regularization and shows the value of regularization in policy optimization, even with exact gradients. This is related to our use of the log barrier regularization.

For our convergence results of the natural policy gradient algorithm in the tabular setting, there are close connections between our results and the works of Even-Dar et al. [2009], Geist et al. [2019]. Even-Dar et al. [2009] provides provable online regret guarantees in changing MDPs utilizing experts algorithms (also see Neu et al. [2010], Abbasi-Yadkori et al. [2019a]); as a special case, their MDP Experts Algorithm is equivalent to the natural policy gradient algorithm with the softmax policy parameterization. While the convergence result due to Even-Dar et al. [2009] was not specifically designed for this setting, it is instructive to see what it implies due to the close connections between optimization and regret [Cesa-Bianchi and Lugosi, 2006, Shalev-Shwartz et al., 2012]. The Mirror Descent-Modified Policy Iteration (MD-MPI) algorithm [Geist et al., 2019] with negative entropy as the Bregman divergence results is an identical algorithm as NPG for softmax parameterization in the tabular case; Corollary 3 [Geist et al., 2019] applies to our updates, leading to a bound worse by a $1/(1 - \gamma)$ factor and also has logarithmic dependence on $|\mathcal{A}|$. Our proof for this case is concise and may be of independent interest. Also worth noting is the Dynamic Policy Programming of Azar et al. [2012], which is an actor-critic algorithm with a softmax parameterization; this algorithm, even though not identical, comes with similar guarantees in terms of its rate (it is weaker in terms of an additional $1/(1 - \gamma)$ factor) than the NPG algorithm.

We now turn to function approximation, starting with a discussion of iterative algorithms which make incremental updates in which the next policy is effectively constrained to be close to the previous policy, such as in CPI and PSDP [Bagnell et al., 2004]. Here, the work in Scherrer and Geist [2014] show how CPI is part of broader family of boosting-style methods. Also, with regards to PSDP, the work in Scherrer [2014] shows how PSDP actually enjoys an improved iteration complexity over CPI, namely $O(\log 1/\epsilon_{\text{opt}})$ vs. $O(1/\epsilon_{\text{opt}}^2)$. It is worthwhile to note that both NPG and projected gradient ascent are also both incremental algorithms.

We now discuss the approximate dynamic programming results characterized in terms of the concentrability coefficient. Broadly we use the term approximate dynamic programming to refer to fitted value iteration, fitted policy iteration and more generally generalized policy iteration

schemes such as classification-based policy iteration as well, in addition to the classical approximate value/policy iteration works. While the approximate dynamic programming results typically require ℓ_∞ bounded errors, which is quite stringent, the notion of concentrability (originally due to [Munos, 2003, 2005]) permits sharper bounds in terms of average case function approximation error, provided that the concentrability coefficient is bounded (e.g. see Munos [2005], Szepesvári and Munos [2005], Antos et al. [2008], Lazaric et al. [2016]). Chen and Jiang [2019] provide a more detailed discussion on this quantity. Based on this problem dependent constant being bounded, Munos [2005], Szepesvári and Munos [2005], Antos et al. [2008] and Lazaric et al. [2016] provide meaningful sample size and error bounds for approximate dynamic programming methods, where there is a data collection policy (under which value-function fitting occurs) that induces a concentrability coefficient. In terms of the concentrability coefficient C_∞ and the “distribution mismatch coefficient” D_∞ in Table 2, we have that $D_\infty \leq C_\infty$, as discussed in [Scherrer, 2014] (also see the table caption). Also, as discussed in Chen and Jiang [2019], a finite concentrability coefficient is a restriction on the MDP dynamics itself, while a bounded D_∞ does not require any restrictions on the MDP dynamics. The more refined quantities defined by Farahmand et al. [2010] (for the approximate policy iteration result) partially alleviate some of these concerns, but their assumptions still implicitly constrain the MDP dynamics, like the finiteness of the concentrability coefficient.

Assuming bounded concentrability coefficient, there are a notable set of provable average case guarantees for the MD-MPI algorithm [Geist et al., 2019] (see also [Azar et al., 2012, Scherrer et al., 2015]), which are stated in terms of various norms of function approximation error. MD-MPI is a class of algorithms for approximate planning under regularized notions of optimality in MDPs. Specifically, Geist et al. [2019] analyze a family of actor-critic style algorithms, where there are both approximate value functions updates and approximate policy updates. As a consequence of utilizing approximate value function updates for the critic, the guarantees of Geist et al. [2019] are stated with dependencies on concentrability coefficients.

When dealing with function approximation, computational and statistical complexities are relevant because they determine the effectiveness of approximate updates with finite samples. With regards to sample complexity, the work in Szepesvári and Munos [2005], Antos et al. [2008] provide finite sample rates (as discussed above), further generalized to actor-critic methods in Azar et al. [2012], Scherrer et al. [2015]. In our policy optimization approach, the analysis of both computational and statistical complexities are straightforward, since we can leverage known statistical and computational results from the stochastic approximation literature; in particular, we use the stochastic projected gradient ascent to obtain a simple, linear time method for the critic estimation step in the natural policy gradient algorithm.

In terms of the algorithmic updates for the function approximation setting, our development of NPG bears similarity to the natural actor-critic algorithm Peters and Schaal [2008], for which some asymptotic guarantees under finite concentrability coefficients are obtained in Bhatnagar et al. [2009]. While both updates seek to minimize the compatible function approximation error, we perform streaming updates based on stochastic optimization using Monte Carlo estimates for values. In contrast Peters and Schaal [2008] utilize Least Squares Temporal Difference methods [Boyan, 1999] to minimize the loss. As a consequence, their updates additionally make linear approximations to

the value functions in order to estimate the advantages; our approach is flexible in allowing for wide family of smoothly differentiable policy classes (including neural policies).

Finally, we remark on some concurrent works. The work of Bhandari and Russo [2019] provides gradient domination-like conditions under which there is (asymptotic) global convergence to the optimal policy. Their results are applicable to the projected gradient ascent algorithm; they are not applicable to gradient ascent with the softmax parameterization (see the discussion in Section 5 herein for the analysis challenges). Bhandari and Russo [2019] also provide global convergence results beyond MDPs. Also, Liu et al. [2019] provide an analysis of the TRPO algorithm [Schulman et al., 2015] with neural network parameterizations, which bears resemblance to our natural policy gradient analysis. In particular, Liu et al. [2019] utilize ideas from both Even-Dar et al. [2009] (with a mirror descent style of analysis) along with Cai et al. [2019] (to handle approximation with neural networks) to provide conditions under which TRPO returns a near optimal policy. Liu et al. [2019] do not explicitly consider the case where the policy class is not complete (i.e when there is approximation). Another related work of Shani et al. [2019] considers the TRPO algorithm and provides theoretical guarantees in the tabular case; their convergence rates with exact updates are $O(1/\sqrt{T})$ for the (unregularized) objective function of interest; they also provide faster rates on a modified (regularized) objective function. They do not consider the case of infinite state spaces and function approximation. The closely related recent papers [Abbasi-Yadkori et al., 2019a,b] also consider closely related algorithms to the Natural Policy Gradient approach studied here, in an infinite horizon, average reward setting. Specifically, the EE-POLITEX algorithm is closely related to the Q-NPG algorithm which we study in Section 6.2, though our approach is in the discounted setting. We adopt the name Q-NPG to capture its close relationship with the NPG algorithm, with the main difference being the use of function approximation for the Q -function instead of advantages. We refer the reader to Section 6.2 (and Remark 6.5) for more discussion of the technical differences between the two works.

3 Setting

A (finite) Markov Decision Process (MDP) $M = (\mathcal{S}, \mathcal{A}, P, r, \gamma, \rho)$ is specified by: a finite state space \mathcal{S} ; a finite action space \mathcal{A} ; a transition model P where $P(s'|s, a)$ is the probability of transitioning into state s' upon taking action a in state s ; a reward function $r : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ where $r(s, a)$ is the immediate reward associated with taking action a in state s ; a discount factor $\gamma \in [0, 1)$; a starting state distribution ρ over \mathcal{S} .

A deterministic, stationary policy $\pi : \mathcal{S} \rightarrow \mathcal{A}$ specifies a decision-making strategy in which the agent chooses actions adaptively based on the current state, i.e., $a_t = \pi(s_t)$. The agent may also choose actions according to a stochastic policy $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$ (where $\Delta(\mathcal{A})$ is the probability simplex over \mathcal{A}), and, overloading notation, we write $a_t \sim \pi(\cdot|s_t)$.

A policy induces a distribution over trajectories $\tau = (s_t, a_t, r_t)_{t=0}^{\infty}$, where s_0 is drawn from the starting state distribution ρ , and, for all subsequent timesteps t , $a_t \sim \pi(\cdot|s_t)$ and $s_{t+1} \sim P(\cdot|s_t, a_t)$. The value function $V^\pi : \mathcal{S} \rightarrow \mathbb{R}$ is defined as the discounted sum of future rewards starting at state

s and executing π , i.e.

$$V^\pi(s) := \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \mid \pi, s_0 = s \right],$$

where the expectation is with respect to the randomness of the trajectory τ induced by π in M . Since we assume that $r(s, a) \in [0, 1]$, we have $0 \leq V^\pi(s) \leq \frac{1}{1-\gamma}$. We overload notation and define $V^\pi(\rho)$ as the expected value under the initial state distribution ρ , i.e.

$$V^\pi(\rho) := \mathbb{E}_{s_0 \sim \rho} [V^\pi(s_0)].$$

The action-value (or Q-value) function $Q^\pi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ and the *advantage* function $A^\pi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ are defined as:

$$Q^\pi(s, a) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \mid \pi, s_0 = s, a_0 = a \right], \quad A^\pi(s, a) := Q^\pi(s, a) - V^\pi(s).$$

The goal of the agent is to find a policy π that maximizes the expected value from the initial state, i.e. the optimization problem the agent seeks to solve is:

$$\max_{\pi} V^\pi(\rho), \tag{1}$$

where the \max is over all policies. The famous theorem of Bellman and Dreyfus [1959] shows there exists a policy π^* which simultaneously maximizes $V^\pi(s_0)$, for all states $s_0 \in \mathcal{S}$.

Policy Parameterizations. This work studies ascent methods for the optimization problem:

$$\max_{\theta \in \Theta} V^{\pi_\theta}(\rho),$$

where $\{\pi_\theta \mid \theta \in \Theta\}$ is some class of parametric (stochastic) policies. We consider a number of different policy classes. The first two are *complete* in the sense that any stochastic policy can be represented in the class. The final class may be restrictive. These classes are as follows:

- *Direct parameterization:* The policies are parameterized by

$$\pi_\theta(a|s) = \theta_{s,a}, \tag{2}$$

where $\theta \in \Delta(\mathcal{A})^{|\mathcal{S}|}$, i.e. θ is subject to $\theta_{s,a} \geq 0$ and $\sum_{a \in \mathcal{A}} \theta_{s,a} = 1$ for all $s \in \mathcal{S}$ and $a \in \mathcal{A}$.

- *Softmax parameterization:* For unconstrained $\theta \in \mathbb{R}^{|\mathcal{S}| \cdot |\mathcal{A}|}$,

$$\pi_\theta(a|s) = \frac{\exp(\theta_{s,a})}{\sum_{a' \in \mathcal{A}} \exp(\theta_{s,a'})}. \tag{3}$$

The softmax parameterization is also complete.




Figure 1: (Non-concavity example) A deterministic MDP corresponding to Lemma 3.1 where $V^{\pi_\theta}(s)$ is not concave. Numbers on arrows represent the rewards for each action.




Figure 2: (Vanishing gradient example) A deterministic, chain MDP of length $H + 2$. We consider a policy where $\pi(a|s_i) = \theta_{s_i,a}$ for $i = 1, 2, \dots, H$. Rewards are 0 everywhere other than $r(s_{H+1}, a_1) = 1$. See Proposition 4.1.

- *Restricted parameterizations:* We also study parametric classes $\{\pi_\theta | \theta \in \Theta\}$ that may not contain all stochastic policies. In particular, we pay close attention to both log-linear policy classes and neural policy classes (see Section 6). Here, the best we may hope for is an agnostic result where we do as well as the best policy in this class.

While the softmax parameterization is the more natural parametrization among the two complete policy classes, it is also informative to consider the direct parameterization.

It is worth explicitly noting that $V^{\pi_\theta}(s)$ is non-concave in θ for both the direct and the softmax parameterizations, so the standard tools of convex optimization are not applicable. For completeness, we formalize this as follows (with a proof in Appendix A, along with an example in Figure 1):

Lemma 3.1. *There is an MDP M (described in Figure 1) such that the optimization problem $V^{\pi_\theta}(s)$ is not concave for both the direct and softmax parameterizations.*

Policy gradients. In order to introduce these methods, it is useful to define the discounted state visitation distribution $d_{s_0}^\pi$ of a policy π as:

$$d_{s_0}^\pi(s) := (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t \Pr^\pi(s_t = s | s_0), \quad (4)$$

where $\Pr^\pi(s_t = s | s_0)$ is the state visitation probability that $s_t = s$, after we execute π starting at state s_0 . Again, we overload notation and write:

$$d_\rho^\pi(s) = \mathbb{E}_{s_0 \sim \rho} [d_{s_0}^\pi(s)],$$

where d_ρ^π is the discounted state visitation distribution under initial distribution ρ .

The policy gradient functional form (see e.g. Williams [1992], Sutton et al. [1999]) is then:

$$\nabla_\theta V^{\pi_\theta}(s_0) = \frac{1}{1 - \gamma} \mathbb{E}_{s \sim d_{s_0}^\pi} \mathbb{E}_{a \sim \pi_\theta(\cdot|s)} [\nabla_\theta \log \pi_\theta(a|s) Q^{\pi_\theta}(s, a)]. \quad (5)$$

Furthermore, if we are working with a differentiable parameterization of $\pi_\theta(\cdot|s)$ that explicitly constrains $\pi_\theta(\cdot|s)$ to be in the simplex, i.e. $\pi_\theta \in \Delta(\mathcal{A})^{|\mathcal{S}|}$ for all θ , then we also have:

$$\nabla_\theta V^{\pi_\theta}(s_0) = \frac{1}{1-\gamma} \mathbb{E}_{s \sim d_{s_0}^{\pi_\theta}} \mathbb{E}_{a \sim \pi_\theta(\cdot|s)} [\nabla_\theta \log \pi_\theta(a|s) A^{\pi_\theta}(s, a)]. \quad (6)$$

Note the above gradient expression (Equation 6) does not hold for the direct parameterization, while Equation 5 is valid.²

The performance difference lemma. The following lemma is helpful throughout:

Lemma 3.2. (*The performance difference lemma [Kakade and Langford, 2002]*) For all policies π, π' and states s_0 ,

$$V^\pi(s_0) - V^{\pi'}(s_0) = \frac{1}{1-\gamma} \mathbb{E}_{s \sim d_{s_0}^\pi} \mathbb{E}_{a \sim \pi(\cdot|s)} [A^{\pi'}(s, a)].$$

For completeness, we provide a proof in Appendix A.

The distribution mismatch coefficient. We often characterize the difficulty of the exploration problem faced by our policy optimization algorithms when maximizing the objective $V^\pi(\mu)$ through the following notion of *distribution mismatch coefficient*.

Definition 3.1 (Distribution mismatch coefficient). Given a policy π and measures $\rho, \mu \in \Delta(\mathcal{S})$, we refer to $\left\| \frac{d_\rho^\pi}{\mu} \right\|_\infty$ as the *distribution mismatch coefficient* of π relative to μ . Here, $\frac{d_\rho^\pi}{\mu}$ denotes componentwise division.

We often instantiate this coefficient with μ as the initial state distribution used in a policy optimization algorithm, ρ as the distribution to measure the sub-optimality of our policy (this is the start state distribution of interest), and where π above is often chosen to be $\pi^* \in \operatorname{argmax}_{\pi \in \Pi} V^\pi(\rho)$, given a policy class Π .

Notation. Following convention, we use V^* and Q^* to denote V^{π^*} and Q^{π^*} respectively. For iterative algorithms which obtain policy parameters $\theta^{(t)}$ at iteration t , we let $\pi^{(t)}$, $V^{(t)}$ and $A^{(t)}$ denote the corresponding quantities parameterized by $\theta^{(t)}$, i.e. $\pi_{\theta^{(t)}}$, $V^{\theta^{(t)}}$ and $A^{\theta^{(t)}}$, respectively. For vectors u and v , we use $\frac{u}{v}$ to denote the componentwise ratio; $u \geq v$ denotes a componentwise inequality; we use the standard convention where $\|v\|_2 = \sqrt{\sum_i v_i^2}$, $\|v\|_1 = \sum_i |v_i|$, and $\|v\|_\infty = \max_i |v_i|$.

²This is due to $\sum_a \nabla_\theta \pi_\theta(a|s) = 0$ not explicitly being maintained by the direct parameterization.

4 Warmup: Constrained Tabular Parameterization

Our starting point is, arguably, the simplest first-order method: we directly take gradient ascent updates on the policy simplex itself and then project back onto the simplex if the constraints are violated after a gradient update. This algorithm is projected gradient ascent on the direct policy parametrization of the MDP, where the parameters are the state-action probabilities, i.e. $\theta_{s,a} = \pi_\theta(a|s)$ (see (2)). As noted in Lemma 3.1, $V^{\pi_\theta}(s)$ is non-concave in the parameters π_θ . Here, we first prove that $V^{\pi_\theta}(\mu)$ satisfies a Polyak-like gradient domination condition [Polyak, 1963], and this tool helps in providing convergence rates. The basic approach was also used in the analysis of CPI [Kakade and Langford, 2002]; related gradient domination-like lemmas also appeared in Scherrer and Geist [2014].

It is instructive to consider this special case due to the connections it makes to the non-convex optimization literature. We also provide a lower bound that rules out algorithms whose runtime appeals to the curvature of saddle points (e.g. [Nesterov and Polyak, 2006, Ge et al., 2015, Jin et al., 2017]).

For the direct policy parametrization where $\theta_{s,a} = \pi_\theta(a|s)$, the gradient is:

$$\frac{\partial V^\pi(\mu)}{\partial \pi(a|s)} = \frac{1}{1-\gamma} d_\mu^\pi(s) Q^\pi(s, a), \quad (7)$$

using (5). In particular, for this parameterization, we may write $\nabla_\pi V^\pi(\mu)$ instead of $\nabla_\theta V^{\pi_\theta}(\mu)$.

4.1 Gradient Domination

Informally, we say a function $f(\theta)$ satisfies a gradient domination property if for all $\theta \in \Theta$,

$$f(\theta^*) - f(\theta) = O(G(\theta)),$$

where $\theta^* \in \operatorname{argmax}_{\theta' \in \Theta} f(\theta')$ and where $G(\theta)$ is some suitable scalar notion of first-order stationarity, which can be considered a measure of how large the gradient is (see [Karimi et al., 2016, Bolte et al., 2007, Attouch et al., 2010]). Thus if one can find a θ that is (approximately) a first-order stationary point, then the parameter θ will be near optimal (in terms of function value). Such conditions are a standard device to establishing global convergence in non-convex optimization, as they effectively rule out the presence of bad critical points. In other words, given such a condition, quantifying the convergence rate for a specific algorithm, like say projected gradient ascent, will require quantifying the rate of its convergence to a first-order stationary point, for which one can invoke standard results from the optimization literature.

The following lemma shows that the direct policy parameterization satisfies a notion of gradient domination. This is the basic approach used in the analysis of CPI [Kakade and Langford, 2002]; a variant of this lemma also appears in Scherrer and Geist [2014]. We give a proof for completeness.

Even though we are interested in the value $V^\pi(\rho)$, it is helpful to consider the gradient with respect to another state distribution $\mu \in \Delta(\mathcal{S})$.

Lemma 4.1 (Gradient domination). *For the direct policy parameterization (as in (2)), for all state distributions $\mu, \rho \in \Delta(\mathcal{S})$, we have*

$$\begin{aligned} V^*(\rho) - V^\pi(\rho) &\leq \left\| \frac{d_\rho^{\pi^*}}{d_\mu^\pi} \right\|_\infty \max_{\bar{\pi}} (\bar{\pi} - \pi)^\top \nabla_\pi V^\pi(\mu) \\ &\leq \frac{1}{1-\gamma} \left\| \frac{d_\rho^{\pi^*}}{\mu} \right\|_\infty \max_{\bar{\pi}} (\bar{\pi} - \pi)^\top \nabla_\pi V^\pi(\mu), \end{aligned}$$

where the max is over the set of all policies, i.e. $\bar{\pi} \in \Delta(\mathcal{A})^{|\mathcal{S}|}$.

Before we provide the proof, a few comments are in order with regards to the performance measure ρ and the optimization measure μ . Subtly, note that although the gradient is with respect to $V^\pi(\mu)$, the final guarantee applies to *all* distributions ρ . The significance is that even though we may be interested in our performance under ρ , it may be helpful to optimize under the distribution μ . To see this, note the lemma shows that a sufficiently small gradient magnitude in the feasible directions implies the policy is nearly optimal in terms of its value, but only if the state distribution of π , i.e. d_μ^π , adequately covers the state distribution of some optimal policy π^* . Here, it is also worth recalling the theorem of Bellman and Dreyfus [1959] which shows there exists a single policy π^* that is simultaneously optimal for all starting states s_0 . Note that the hardness of the exploration problem is captured through the distribution mismatch coefficient (Definition 3.1).

Proof:[of Lemma 4.1] By the performance difference lemma (Lemma 3.2),

$$\begin{aligned} V^*(\rho) - V^\pi(\rho) &= \frac{1}{1-\gamma} \sum_{s,a} d_\rho^{\pi^*}(s) \pi^*(a|s) A^\pi(s, a) \\ &\leq \frac{1}{1-\gamma} \sum_{s,a} d_\rho^{\pi^*}(s) \max_{\bar{a}} A^\pi(s, \bar{a}) \\ &= \frac{1}{1-\gamma} \sum_s \frac{d_\rho^{\pi^*}(s)}{d_\mu^\pi(s)} \cdot d_\mu^\pi(s) \max_{\bar{a}} A^\pi(s, \bar{a}) \\ &\leq \frac{1}{1-\gamma} \left(\max_s \frac{d_\rho^{\pi^*}(s)}{d_\mu^\pi(s)} \right) \sum_s d_\mu^\pi(s) \max_{\bar{a}} A^\pi(s, \bar{a}), \end{aligned} \tag{8}$$

where the last inequality follows since $\max_{\bar{a}} A^\pi(s, \bar{a}) \geq 0$ for all states s and policies π . We wish

to upper bound (8). We then have:

$$\begin{aligned}
\sum_s \frac{d_\mu^\pi(s)}{1-\gamma} \max_{\bar{a}} A^\pi(s, \bar{a}) &= \max_{\bar{\pi} \in \Delta(\mathcal{A})^{|\mathcal{S}|}} \sum_{s,a} \frac{d_\mu^\pi(s)}{1-\gamma} \bar{\pi}(a|s) A^\pi(s, a) \\
&= \max_{\bar{\pi} \in \Delta(\mathcal{A})^{|\mathcal{S}|}} \sum_{s,a} \frac{d_\mu^\pi(s)}{1-\gamma} (\bar{\pi}(a|s) - \pi(a|s)) A^\pi(s, a) \\
&= \max_{\bar{\pi} \in \Delta(\mathcal{A})^{|\mathcal{S}|}} \sum_{s,a} \frac{d_\mu^\pi(s)}{1-\gamma} (\bar{\pi}(a|s) - \pi(a|s)) Q^\pi(s, a) \\
&= \max_{\bar{\pi} \in \Delta(\mathcal{A})^{|\mathcal{S}|}} (\bar{\pi} - \pi)^\top \nabla_\pi V^\pi(\mu)
\end{aligned}$$

where the first step follows since $\max_{\bar{a}}$ is attained at an action which maximizes $A^\pi(s, \cdot)$ (per state); the second step follows as $\sum_a \pi(a|s) A^\pi(s, a) = 0$; the third step uses $\sum_a (\bar{\pi}(a|s) - \pi(a|s)) V^\pi(s) = 0$ for all s ; and the final step follows from the gradient expression (see (7)). Using this in (8),

$$\begin{aligned}
V^*(\rho) - V^\pi(\rho) &\leq \left\| \frac{d_\rho^{\pi^*}}{d_\mu^\pi} \right\|_\infty \max_{\bar{\pi} \in \Delta(\mathcal{A})^{|\mathcal{S}|}} (\bar{\pi} - \pi)^\top \nabla_\pi V^\pi(\mu) \\
&\leq \frac{1}{1-\gamma} \left\| \frac{d_\rho^{\pi^*}}{\mu} \right\|_\infty \max_{\bar{\pi} \in \Delta(\mathcal{A})^{|\mathcal{S}|}} (\bar{\pi} - \pi)^\top \nabla_\pi V^\pi(\mu).
\end{aligned}$$

where the last step follows due to $\max_{\bar{\pi} \in \Delta(\mathcal{A})^{|\mathcal{S}|}} (\bar{\pi} - \pi)^\top \nabla_\pi V^\pi(\mu) \geq 0$ for any policy π and $d_\mu^\pi(s) \geq (1-\gamma)\mu(s)$ (see (4)). \blacksquare

In a sense, the use of an appropriate μ circumvents the issues of strategic exploration. It is natural to ask whether this additional term is necessary, a question which we return to. First, we provide a convergence rate for the projected gradient ascent algorithm.

4.2 Convergence Rates for Projected Gradient Ascent

Using this notion of gradient domination, we now give an iteration complexity bound for projected gradient ascent over the space of stochastic policies, i.e. over $\Delta(\mathcal{A})^{|\mathcal{S}|}$. The projected gradient ascent algorithm updates

$$\pi^{(t+1)} = P_{\Delta(\mathcal{A})^{|\mathcal{S}|}}(\pi^{(t)} + \eta \nabla_\pi V^{(t)}(\mu)), \tag{9}$$

where $P_{\Delta(\mathcal{A})^{|\mathcal{S}|}}$ is the projection onto $\Delta(\mathcal{A})^{|\mathcal{S}|}$ in the Euclidean norm.

Theorem 4.1. *The projected gradient ascent algorithm (9) on $V^\pi(\mu)$ with stepsize $\eta = \frac{(1-\gamma)^3}{2\gamma|\mathcal{A}|}$ satisfies for all distributions $\rho \in \Delta(\mathcal{S})$,*

$$\min_{i < T} \{V^*(\rho) - V^{(i)}(\rho)\} \leq \epsilon \quad \text{whenever} \quad T > \frac{64\gamma|\mathcal{S}||\mathcal{A}|}{(1-\gamma)^6\epsilon^2} \left\| \frac{d_\rho^{\pi^*}}{\mu} \right\|_\infty^2.$$

A proof is provided in Appendix B.1. The proof first invokes a standard iteration complexity result of projected gradient ascent to show that the gradient magnitude with respect to all feasible directions is small. More concretely, we show the policy is ϵ -stationary³, that is, for all $\pi_\theta + \delta \in \Delta(\mathcal{A})^{|\mathcal{S}|}$ and $\|\delta\|_2 \leq 1$, $\delta^\top \nabla_\pi V^{\pi_\theta}(\mu) \leq \epsilon$. We then use Lemma 4.1 to complete the proof.

Note that the guarantee we provide is for the best policy found over the T rounds, which we obtain from a bound on the average norm of the gradients. This type of a guarantee is standard in the non-convex optimization literature, where an average regret bound cannot be used to extract a single good solution, e.g. by averaging. In the context of policy optimization, this is not a serious limitation as we collect on-policy trajectories for each policy in doing sample-based gradient estimation, and these samples can be also used to estimate the policy’s value. Note that the evaluation step is not required for every policy, and can also happen on a schedule, though we still need to evaluate $O(T)$ policies to obtain the convergence rates described here.

4.3 A Lower Bound: Vanishing Gradients and Saddle Points

To understand the necessity of the distribution mismatch coefficient in Lemma 4.1 and Theorem 4.1, let us first give an informal argument that some condition on the state distribution of π , or equivalently μ , is necessary for stationarity to imply optimality. For example, in a sparse-reward MDP (where the agent is only rewarded upon visiting some small set of states), a policy that does not visit *any* rewarding states will have zero gradient, even though it is arbitrarily suboptimal in terms of values. Below, we give a more quantitative version of this intuition, which demonstrates that even if π chooses all actions with reasonable probabilities (and hence the agent will visit all states if the MDP is connected), then there is an MDP where a large fraction of the policies π have vanishingly small gradients, and yet these policies are highly suboptimal in terms of their value.

Concretely, consider the chain MDP of length $H + 2$ shown in Figure 2. The starting state of interest is state s_0 and the discount factor $\gamma = H/(H + 1)$. Suppose we work with the direct parameterization, where $\pi_\theta(a|s) = \theta_{s,a}$ for $a = a_1, a_2, a_3$ and $\pi_\theta(a_4|s) = 1 - \theta_{s,a_1} - \theta_{s,a_2} - \theta_{s,a_3}$. Note we do not over-parameterize the policy. For this MDP and policy structure, if we were to initialize the probabilities over actions, say deterministically, then there is an MDP (obtained by permuting the actions) where all the probabilities for a_1 will be less than $1/4$.

The following result not only shows that the gradient is exponentially small in H , it also shows that many higher order derivatives, up to $O(H/\log H)$, are also exponentially small in H .

Proposition 4.1 (Vanishing gradients at suboptimal parameters). Consider the chain MDP of Figure 2, with $H + 2$ states, $\gamma = H/(H + 1)$, and with the direct policy parameterization (with $3|\mathcal{S}|$ parameters, as described in the text above). Suppose θ is such that $0 < \theta < 1$ (componentwise) and $\theta_{s,a_1} < 1/4$ (for all states s). For all $k \leq \frac{H}{40 \log(2H)} - 1$, we have $\|\nabla_\theta^k V^{\pi_\theta}(s_0)\| \leq (1/3)^{H/4}$, where $\nabla_\theta^k V^{\pi_\theta}(s_0)$ is a tensor of the k_{th} order derivatives of $V^{\pi_\theta}(s_0)$ and the norm is the operator norm of the tensor.⁴ Furthermore, $V^*(s_0) - V^{\pi_\theta}(s_0) \geq (H + 1)/8 - (H + 1)^2/3^H$.

³See Appendix B.1 for discussion on this definition.

⁴The operator norm of a k_{th} -order tensor $J \in \mathbb{R}^{d^{\otimes k}}$ is defined as $\sup_{u_1, \dots, u_k \in \mathbb{R}^d : \|u_i\|_2=1} \langle J, u_1 \otimes \dots \otimes u_d \rangle$.

This lemma also suggests that results in the non-convex optimization literature, on escaping from saddle points, e.g. [Nesterov and Polyak, 2006, Ge et al., 2015, Jin et al., 2017], do not directly imply global convergence due to that the higher order derivatives are small.

Remark 4.1. (Exact vs. Approximate Gradients) The chain MDP of Figure 2, is a common example where *sample* based estimates of gradients will be 0 under random exploration strategies; there is an exponentially small in H chance of hitting the goal state under a random exploration strategy. Note that this lemma is with regards to *exact* gradients. This suggests that even with exact computations (along with using exact higher order derivatives) we might expect numerical instabilities.

Remark 4.2. (Comparison with the upper bound) The lower bound does not contradict the upper bound of Theorem 4.1 (where a small gradient is turned into a small policy suboptimality bound), as the distribution mismatch coefficient, as defined in Definition 3.1, could be infinite in the chain MDP of Figure 2, since the start-state distribution is concentrated on one state only. More generally, for any policy with $\theta_{s,a_1} < 1/4$ in all states s , $\left\| \frac{d\pi^*}{d\theta} \right\|_{\infty} = \Omega(4^H)$.

Remark 4.3. (Comparison with information-theoretic lower bounds) The lower bound here is *not information theoretic*, in that it does not present a hard problem instance for all algorithms. Indeed, exploration algorithms for tabular MDPs starting from E^3 [Kearns and Singh, 2002], RMAX [Brafman and Tennensky, 2003] and several subsequent works yield polynomial sample complexities for the chain MDP. Proposition 4.1 should be interpreted as a hardness result for the specific class of policy gradient like approaches that search for a policy with a small policy gradient, as these methods will find the initial parameters to be valid in terms of the size of (several orders of) gradients. In particular, it precludes any meaningful claims on global optimality, based just on the size of the policy gradients, without additional assumptions as discussed in the previous remark.

The proof is provided in Appendix B.2. The lemma illustrates that lack of good exploration can indeed be detrimental in policy gradient algorithms, since the gradient can be small either due to π being near-optimal, or, simply because π does not visit advantageous states often enough. In this sense, it also demonstrates the necessity of the distribution mismatch coefficient in Lemma 4.1.

5 The Softmax Tabular Parameterization

We now consider the softmax policy parameterization (3). Here, we still have a non-concave optimization problem in general, as shown in Lemma 3.1, though we do show that global optimality can be reached under certain regularity conditions. From a practical perspective, the softmax parameterization of policies is preferable to the direct parameterization, since the parameters θ are unconstrained and standard unconstrained optimization algorithms can be employed. However, optimization over this policy class creates other challenges as we study in this section, as the optimal policy (which is deterministic) is attained by sending the parameters to infinity.

We study three algorithms for this problem. The first performs direct policy gradient ascent on the objective without modification, while the second adds a log barrier regularizer to keep the

parameters from becoming too large, as a means to ensure adequate exploration. Finally, we study the natural policy gradient algorithm and establish a global optimality result with no dependence on the distribution mismatch coefficient or dimension-dependent factors.

For the softmax parameterization, the gradient takes the form:

$$\frac{\partial V^{\pi_\theta}(\mu)}{\partial \theta_{s,a}} = \frac{1}{1-\gamma} d_\mu^{\pi_\theta}(s) \pi_\theta(a|s) A^{\pi_\theta}(s, a) \quad (10)$$

(see Lemma C.1 for a proof).

5.1 Asymptotic Convergence, without Regularization

Due to the exponential scaling with the parameters θ in the softmax parameterization, *any* policy that is nearly deterministic will have gradients close to 0. In spite of this difficulty, we provide a positive result that gradient ascent asymptotically converges to the global optimum for the softmax parameterization.

The update rule for gradient ascent is:

$$\theta^{(t+1)} = \theta^{(t)} + \eta \nabla_\theta V^{(t)}(\mu). \quad (11)$$

Theorem 5.1 (Global convergence for softmax parameterization). *Assume we follow the gradient ascent update rule as specified in Equation (11) and that the distribution μ is strictly positive i.e. $\mu(s) > 0$ for all states s . Suppose $\eta \leq \frac{(1-\gamma)^3}{8}$, then we have that for all states s , $V^{(t)}(s) \rightarrow V^*(s)$ as $t \rightarrow \infty$.*

Remark 5.1. (Strict positivity of μ and exploration) Theorem 5.1 assumed that optimization distribution μ was *strictly* positive, i.e. $\mu(s) > 0$ for all states s . We leave it is an open question of whether or not gradient ascent will globally converge if this condition is not met. The concern is that if this condition is not met, then gradient ascent may not globally converge due to that $d_\mu^{\pi_\theta}(s)$ effectively scales down the learning rate for the parameters associated with state s (see (10)).

The complete proof is provided in the Appendix C.1. We now discuss the subtleties in the proof and show why the softmax parameterization precludes a direct application of the gradient domination lemma. In order to utilize the gradient domination property (in Lemma 4.1), we would desire to show that: $\nabla_\pi V^\pi(\mu) \rightarrow 0$. However, using the functional form of the softmax parameterization (see Lemma C.1) and (7), we have that:

$$\frac{\partial V^{\pi_\theta}(\mu)}{\partial \theta_{s,a}} = \frac{1}{1-\gamma} d_\mu^{\pi_\theta}(s) \pi_\theta(a|s) A^{\pi_\theta}(s, a) = \pi_\theta(a|s) \frac{\partial V^{\pi_\theta}(\mu)}{\partial \pi_\theta(a|s)}.$$

Hence, we see that even if $\nabla_\theta V^{\pi_\theta}(\mu) \rightarrow 0$, we are not guaranteed that $\nabla_\pi V^{\pi_\theta}(\mu) \rightarrow 0$.

We now briefly discuss the main technical challenges in the proof. The proof first shows that the sequence $V^{(t)}(s)$ is monotone increasing pointwise, i.e. for *every* state s , $V^{(t+1)}(s) \geq V^{(t)}(s)$ (Lemma C.2). This implies the existence of a limit $V^{(\infty)}(s)$ by the monotone convergence theorem

(Lemma C.3). Based on the limiting quantities $V^{(\infty)}(s)$ and $Q^{(\infty)}(s, a)$, which we show exist, define the following limiting sets for each state s :

$$\begin{aligned} I_0^s &:= \{a | Q^{(\infty)}(s, a) = V^{(\infty)}(s)\} \\ I_+^s &:= \{a | Q^{(\infty)}(s, a) > V^{(\infty)}(s)\} \\ I_-^s &:= \{a | Q^{(\infty)}(s, a) < V^{(\infty)}(s)\}. \end{aligned}$$

The challenge is to then show that, for all states s , the set I_+^s is the empty set, which would immediately imply $V^{(\infty)}(s) = V^*(s)$. The proof proceeds by contradiction, assuming that I_+^s is non-empty. Using that I_+^s is non-empty and that the gradient tends to zero in the limit, i.e. $\nabla_{\theta} V^{\pi_{\theta}}(\mu) \rightarrow 0$, we have that for all $a \in I_+^s$, $\pi^{(t)}(a|s) \rightarrow 0$ (see (10)). This, along with the functional form of the softmax parameterization, implies that there must be divergence (in magnitude) among the set of parameters associated with *some* action a at state s , i.e. that $\max_{a \in \mathcal{A}} |\theta_{s,a}^{(t)}| \rightarrow \infty$. The primary technical challenge in the proof is to then use this divergence, along with the dynamics of gradient ascent, to show that I_+^s is empty via a contradiction.

We leave it as a question for future work as to characterizing the convergence rate, which we conjecture is exponentially slow in some of the relevant quantities, such as in terms of the size of state space. Here, we turn to a regularization based approach to ensure convergence at a polynomial rate in all relevant quantities.

5.2 Polynomial Convergence with Log Barrier Regularization

Due to the exponential scaling with the parameters θ , policies can rapidly become near deterministic, when optimizing under the softmax parameterization, which can result in slow convergence. Indeed a key challenge in the asymptotic analysis in the previous section was to handle the growth of the absolute values of parameters as they tend to infinity. A common practical remedy for this is to use entropy-based regularization to keep the probabilities from getting too small [Williams and Peng, 1991, Mnih et al., 2016], and we study gradient ascent on a similarly regularized objective in this section. Recall that the relative-entropy for distributions p and q is defined as: $\text{KL}(p, q) := \mathbb{E}_{x \sim p}[-\log q(x)/p(x)]$. Denote the uniform distribution over a set \mathcal{X} by $\text{Unif}_{\mathcal{X}}$, and define the following log barrier regularized objective as:

$$\begin{aligned} L_{\lambda}(\theta) &:= V^{\pi_{\theta}}(\mu) - \lambda \mathbb{E}_{s \sim \text{Unif}_{\mathcal{S}}} \left[\text{KL}(\text{Unif}_{\mathcal{A}}, \pi_{\theta}(\cdot|s)) \right] \\ &= V^{\pi_{\theta}}(\mu) + \frac{\lambda}{|\mathcal{S}| |\mathcal{A}|} \sum_{s,a} \log \pi_{\theta}(a|s) + \lambda \log |\mathcal{A}|, \end{aligned} \quad (12)$$

where λ is a regularization parameter. The constant (i.e. the last term) is not relevant with regards to optimization. This regularizer is different from the more commonly utilized entropy regularizer as in Mnih et al. [2016], a point which we return to in Remark 5.2.

The policy gradient ascent updates for $L_{\lambda}(\theta)$ are given by:

$$\theta^{(t+1)} = \theta^{(t)} + \eta \nabla_{\theta} L_{\lambda}(\theta^{(t)}). \quad (13)$$

Our next theorem shows that approximate first-order stationary points of the entropy-regularized objective are approximately globally optimal, provided the regularization is sufficiently small.

Theorem 5.2. (*Log barrier regularization*) Suppose θ is such that:

$$\|\nabla_{\theta} L_{\lambda}(\theta)\|_2 \leq \epsilon_{\text{opt}}$$

and $\epsilon_{\text{opt}} \leq \lambda/(2|\mathcal{S}||\mathcal{A}|)$. Then we have that for all starting state distributions ρ :

$$V^{\pi_{\theta}}(\rho) \geq V^*(\rho) - \frac{2\lambda}{1-\gamma} \left\| \frac{d_{\rho}^{\pi^*}}{\mu} \right\|_{\infty}.$$

Proof: The proof consists of showing that $\max_a A^{\pi_{\theta}}(s, a) \leq 2\lambda/(\mu(s)|\mathcal{S}|)$ for all states. To see that this is sufficient, observe that by the performance difference lemma (Lemma 3.2),

$$\begin{aligned} V^*(\rho) - V^{\pi_{\theta}}(\rho) &= \frac{1}{1-\gamma} \sum_{s,a} d_{\rho}^{\pi^*}(s) \pi^*(a|s) A^{\pi_{\theta}}(s, a) \\ &\leq \frac{1}{1-\gamma} \sum_s d_{\rho}^{\pi^*}(s) \max_{a \in \mathcal{A}} A^{\pi_{\theta}}(s, a) \\ &\leq \frac{1}{1-\gamma} \sum_s 2d_{\rho}^{\pi^*}(s) \lambda/(\mu(s)|\mathcal{S}|) \\ &\leq \frac{2\lambda}{1-\gamma} \max_s \left(\frac{d_{\rho}^{\pi^*}(s)}{\mu(s)} \right). \end{aligned}$$

which would then complete the proof.

We now proceed to show that $\max_a A^{\pi_{\theta}}(s, a) \leq 2\lambda/(\mu(s)|\mathcal{S}|)$. For this, it suffices to bound $A^{\pi_{\theta}}(s, a)$ for any state-action pair s, a where $A^{\pi_{\theta}}(s, a) \geq 0$ else the claim is trivially true. Consider an (s, a) pair such that $A^{\pi_{\theta}}(s, a) > 0$. Using the policy gradient expression for the softmax parameterization (see Lemma C.1),

$$\frac{\partial L_{\lambda}(\theta)}{\partial \theta_{s,a}} = \frac{1}{1-\gamma} d_{\mu}^{\pi_{\theta}}(s) \pi_{\theta}(a|s) A^{\pi_{\theta}}(s, a) + \frac{\lambda}{|\mathcal{S}|} \left(\frac{1}{|\mathcal{A}|} - \pi_{\theta}(a|s) \right). \quad (14)$$

The gradient norm assumption $\|\nabla_{\theta} L_{\lambda}(\theta)\|_2 \leq \epsilon_{\text{opt}}$ implies that:

$$\begin{aligned} \epsilon_{\text{opt}} &\geq \frac{\partial L_{\lambda}(\theta)}{\partial \theta_{s,a}} = \frac{1}{1-\gamma} d_{\mu}^{\pi_{\theta}}(s) \pi_{\theta}(a|s) A^{\pi_{\theta}}(s, a) + \frac{\lambda}{|\mathcal{S}|} \left(\frac{1}{|\mathcal{A}|} - \pi_{\theta}(a|s) \right) \\ &\geq \frac{\lambda}{|\mathcal{S}|} \left(\frac{1}{|\mathcal{A}|} - \pi_{\theta}(a|s) \right), \end{aligned}$$

where we have used $A^{\pi_{\theta}}(s, a) \geq 0$. Rearranging and using our assumption $\epsilon_{\text{opt}} \leq \lambda/(2|\mathcal{S}||\mathcal{A}|)$,

$$\pi_{\theta}(a|s) \geq \frac{1}{|\mathcal{A}|} - \frac{\epsilon_{\text{opt}}|\mathcal{S}|}{\lambda} \geq \frac{1}{2|\mathcal{A}|}.$$

Solving for $A^{\pi_\theta}(s, a)$ in (14), we have:

$$\begin{aligned}
A^{\pi_\theta}(s, a) &= \frac{1-\gamma}{d_\mu^{\pi_\theta}(s)} \left(\frac{1}{\pi_\theta(a|s)} \frac{\partial L_\lambda(\theta)}{\partial \theta_{s,a}} + \frac{\lambda}{|\mathcal{S}|} \left(1 - \frac{1}{\pi_\theta(a|s)|\mathcal{A}|} \right) \right) \\
&\leq \frac{1-\gamma}{d_\mu^{\pi_\theta}(s)} \left(2|\mathcal{A}|\epsilon_{\text{opt}} + \frac{\lambda}{|\mathcal{S}|} \right) \\
&\leq 2 \frac{1-\gamma}{d_\mu^{\pi_\theta}(s)} \frac{\lambda}{|\mathcal{S}|} \\
&\leq 2\lambda/(\mu(s)|\mathcal{S}|),
\end{aligned}$$

where the penultimate step uses $\epsilon_{\text{opt}} \leq \lambda/(2|\mathcal{S}||\mathcal{A}|)$ and the final step uses $d_\mu^{\pi_\theta}(s) \geq (1-\gamma)\mu(s)$. This completes the proof. \blacksquare

By combining the above theorem with standard results on the convergence of gradient ascent (to first order stationary points), we obtain the following corollary.

Corollary 5.1. (*Iteration complexity with log barrier regularization*) Let $\beta_\lambda := \frac{8\gamma}{(1-\gamma)^3} + \frac{2\lambda}{|\mathcal{S}|}$. Starting from any initial $\theta^{(0)}$, consider the updates (13) with $\lambda = \frac{\epsilon(1-\gamma)}{2\left\|\frac{d_\rho^{\pi^*}}{\mu}\right\|_\infty}$ and $\eta = 1/\beta_\lambda$. Then for all starting state distributions ρ , we have

$$\min_{t < T} \{V^*(\rho) - V^{(t)}(\rho)\} \leq \epsilon \quad \text{whenever} \quad T \geq \frac{320|\mathcal{S}|^2|\mathcal{A}|^2}{(1-\gamma)^6 \epsilon^2} \left\|\frac{d_\rho^{\pi^*}}{\mu}\right\|_\infty^2.$$

See Appendix C.2 for the proof. The corollary shows the importance of balancing how the regularization parameter λ is set relative to the desired accuracy ϵ , as well as the importance of the initial distribution μ to obtain global optimality.

Remark 5.2. (*Entropy vs. log barrier regularization*) The more commonly considered regularizer is the entropy [Mnih et al., 2016] (also see Ahmed et al. [2019] for a more detailed empirical investigation), where the regularizer would be:

$$\frac{1}{|\mathcal{S}|} \sum_s H(\pi_\theta(\cdot|s)) = \frac{1}{|\mathcal{S}|} \sum_s \sum_a -\pi_\theta(a|s) \log \pi_\theta(a|s).$$

Note the entropy is far less aggressive in penalizing small probabilities, in comparison to the log barrier, which is equivalent to the relative entropy. In particular, the entropy regularizer is always bounded between 0 and $\log |\mathcal{A}|$, while the relative entropy (against the uniform distribution over actions), is bounded between 0 and infinity, where it tends to infinity as probabilities tend to 0. We leave it as an open question if a polynomial convergence rate⁵ is achievable with the more common entropy regularizer; our polynomial convergence rate using the KL regularizer crucially relies on the aggressive nature in which the relative entropy prevents small probabilities (the proof shows that any action, with a positive advantage, has a significant probability for any near-stationary policy of the regularized objective).

⁵Here, ideally we would like to be poly in $|\mathcal{S}|$, $|\mathcal{A}|$, $1/(1-\gamma)$, $1/\epsilon$, and the distribution mismatch coefficient, which we conjecture may not be possible.

5.3 Dimension-free Convergence of Natural Policy Gradient Ascent

We now show the Natural Policy Gradient algorithm, with the softmax parameterization (3), obtains an improved iteration complexity. The NPG algorithm defines a Fisher information matrix (induced by π), and performs gradient updates in the geometry induced by this matrix as follows:

$$\begin{aligned} F_\rho(\theta) &= \mathbb{E}_{s \sim d_\rho^\pi} \mathbb{E}_{a \sim \pi_\theta(\cdot|s)} \left[\nabla_\theta \log \pi_\theta(a|s) \left(\nabla_\theta \log \pi_\theta(a|s) \right)^\top \right] \\ \theta^{(t+1)} &= \theta^{(t)} + \eta F_\rho(\theta^{(t)})^\dagger \nabla_\theta V^{(t)}(\rho), \end{aligned} \quad (15)$$

where M^\dagger denotes the Moore-Penrose pseudoinverse of the matrix M . Throughout this section, we restrict to using the initial state distribution $\rho \in \Delta(\mathcal{S})$ in our update rule in (15) (so our optimization measure μ and the performance measure ρ are identical). Also, we restrict attention to states $s \in \mathcal{S}$ reachable from ρ , since, without loss of generality, we can exclude states that are not reachable under this start state distribution⁶.

We leverage a particularly convenient form the update takes for the softmax parameterization (see Kakade [2001]). For completeness, we provide a proof in Appendix C.3.

Lemma 5.1. (*NPG as soft policy iteration*) *For the softmax parameterization (3), the NPG updates (15) take the form:*

$$\theta^{(t+1)} = \theta^{(t)} + \frac{\eta}{1-\gamma} A^{(t)} \quad \text{and} \quad \pi^{(t+1)}(a|s) = \pi^{(t)}(a|s) \frac{\exp(\eta A^{(t)}(s, a)/(1-\gamma))}{Z_t(s)},$$

where $Z_t(s) = \sum_{a \in \mathcal{A}} \pi^{(t)}(a|s) \exp(\eta A^{(t)}(s, a)/(1-\gamma))$.

The updates take a strikingly simple form in this special case; they are identical to the classical multiplicative weights updates [Freund and Schapire, 1997, Cesa-Bianchi and Lugosi, 2006] for online linear optimization over the probability simplex, where the linear functions are specified by the advantage function of the current policy at each iteration. Notably, there is no dependence on the state distribution $d_\rho^{(t)}$, since the pseudoinverse of the Fisher information cancels out the effect of the state distribution in NPG. We now provide a dimension free convergence rate of this algorithm.

Theorem 5.3 (Global convergence for NPG). *Suppose we run the NPG updates (15) using $\rho \in \Delta(\mathcal{S})$ and with $\theta^{(0)} = 0$. Fix $\eta > 0$. For all $T > 0$, we have:*

$$V^{(T)}(\rho) \geq V^*(\rho) - \frac{\log |\mathcal{A}|}{\eta T} - \frac{1}{(1-\gamma)^2 T}.$$

In particular, setting $\eta \geq (1-\gamma)^2 \log |\mathcal{A}|$, we see that NPG finds an ϵ -optimal policy in a number of iterations that is at most:

$$T \leq \frac{2}{(1-\gamma)^2 \epsilon},$$

⁶Specifically, we restrict the MDP to the set of states $\{s \in \mathcal{S} : \exists \pi \text{ such that } d_\rho^\pi(s) > 0\}$.

which has no dependence on the number of states or actions, despite the non-concavity of the underlying optimization problem.

The proof strategy we take borrows ideas from the online regret framework in changing MDPs (in [Even-Dar et al., 2009]); here, we provide a faster rate of convergence than the analysis implied by Even-Dar et al. [2009] or by Geist et al. [2019]. We also note that while this proof is obtained for the NPG updates, it is known in the literature that in the limit of small stepsizes, NPG and TRPO updates are closely related (e.g. see Schulman et al. [2015], Neu et al. [2017], Rajeswaran et al. [2017]).

First, the following improvement lemma is helpful:

Lemma 5.2 (Improvement lower bound for NPG). *For the iterates $\pi^{(t)}$ generated by the NPG updates (15), we have for all starting state distributions μ*

$$V^{(t+1)}(\mu) - V^{(t)}(\mu) \geq \frac{(1-\gamma)}{\eta} \mathbb{E}_{s \sim \mu} \log Z_t(s) \geq 0.$$

Proof: First, let us show that $\log Z_t(s) \geq 0$. To see this, observe:

$$\begin{aligned} \log Z_t(s) &= \log \sum_a \pi^{(t)}(a|s) \exp(\eta A^{(t)}(s, a)/(1-\gamma)) \\ &\geq \sum_a \pi^{(t)}(a|s) \log \exp(\eta A^{(t)}(s, a)/(1-\gamma)) = \frac{\eta}{1-\gamma} \sum_a \pi^{(t)}(a|s) A^{(t)}(s, a) = 0. \end{aligned}$$

where the inequality follows by Jensen's inequality on the concave function $\log x$ and the final equality uses $\sum_a \pi^{(t)}(a|s) A^{(t)}(s, a) = 0$. Using $d^{(t+1)}$ as shorthand for $d_\mu^{(t+1)}$, the performance difference lemma implies:

$$\begin{aligned} V^{(t+1)}(\mu) - V^{(t)}(\mu) &= \frac{1}{1-\gamma} \mathbb{E}_{s \sim d^{(t+1)}} \sum_a \pi^{(t+1)}(a|s) A^{(t)}(s, a) \\ &= \frac{1}{\eta} \mathbb{E}_{s \sim d^{(t+1)}} \sum_a \pi^{(t+1)}(a|s) \log \frac{\pi^{(t+1)}(a|s) Z_t(s)}{\pi^{(t)}(a|s)} \\ &= \frac{1}{\eta} \mathbb{E}_{s \sim d^{(t+1)}} \text{KL}(\pi_s^{(t+1)} || \pi_s^{(t)}) + \frac{1}{\eta} \mathbb{E}_{s \sim d^{(t+1)}} \log Z_t(s) \\ &\geq \frac{1}{\eta} \mathbb{E}_{s \sim d^{(t+1)}} \log Z_t(s) \geq \frac{1-\gamma}{\eta} \mathbb{E}_{s \sim \mu} \log Z_t(s), \end{aligned}$$

where the last step uses that $d^{(t+1)} = d_\mu^{(t+1)} \geq (1-\gamma)\mu$, componentwise (by (4)), and that $\log Z_t(s) \geq 0$. \blacksquare

With this lemma, we now prove Theorem 5.3.

Proof:[of Theorem 5.3] Since ρ is fixed, we use d^* as shorthand for $d_\rho^{\pi^*}$; we also use π_s as

shorthand for the vector of $\pi(\cdot|s)$. By the performance difference lemma (Lemma 3.2),

$$\begin{aligned}
V^{\pi^*}(\rho) - V^{(t)}(\rho) &= \frac{1}{1-\gamma} \mathbb{E}_{s \sim d^*} \sum_a \pi^*(a|s) A^{(t)}(s, a) \\
&= \frac{1}{\eta} \mathbb{E}_{s \sim d^*} \sum_a \pi^*(a|s) \log \frac{\pi^{(t+1)}(a|s) Z_t(s)}{\pi^{(t)}(a|s)} \\
&= \frac{1}{\eta} \mathbb{E}_{s \sim d^*} \left(\text{KL}(\pi_s^* || \pi_s^{(t)}) - \text{KL}(\pi_s^* || \pi_s^{(t+1)}) + \sum_a \pi^*(a|s) \log Z_t(s) \right) \\
&= \frac{1}{\eta} \mathbb{E}_{s \sim d^*} \left(\text{KL}(\pi_s^* || \pi_s^{(t)}) - \text{KL}(\pi_s^* || \pi_s^{(t+1)}) + \log Z_t(s) \right),
\end{aligned}$$

where we have used the closed form of our updates from Lemma 5.1 in the second step.

By applying Lemma 5.2 with d^* as the starting state distribution, we have:

$$\frac{1}{\eta} \mathbb{E}_{s \sim d^*} \log Z_t(s) \leq \frac{1}{1-\gamma} \left(V^{(t+1)}(d^*) - V^{(t)}(d^*) \right)$$

which gives us a bound on $\mathbb{E}_{s \sim d^*} \log Z_t(s)$.

Using the above equation and that $V^{(t+1)}(\rho) \geq V^{(t)}(\rho)$ (as $V^{(t+1)}(s) \geq V^{(t)}(s)$ for all states s by Lemma 5.2), we have:

$$\begin{aligned}
V^{\pi^*}(\rho) - V^{(T-1)}(\rho) &\leq \frac{1}{T} \sum_{t=0}^{T-1} (V^{\pi^*}(\rho) - V^{(t)}(\rho)) \\
&\leq \frac{1}{\eta T} \sum_{t=0}^{T-1} \mathbb{E}_{s \sim d^*} (\text{KL}(\pi_s^* || \pi_s^{(t)}) - \text{KL}(\pi_s^* || \pi_s^{(t+1)})) + \frac{1}{\eta T} \sum_{t=0}^{T-1} \mathbb{E}_{s \sim d^*} \log Z_t(s) \\
&\leq \frac{\mathbb{E}_{s \sim d^*} \text{KL}(\pi_s^* || \pi^{(0)})}{\eta T} + \frac{1}{(1-\gamma)T} \sum_{t=0}^{T-1} \left(V^{(t+1)}(d^*) - V^{(t)}(d^*) \right) \\
&= \frac{\mathbb{E}_{s \sim d^*} \text{KL}(\pi_s^* || \pi^{(0)})}{\eta T} + \frac{V^{(T)}(d^*) - V^{(0)}(d^*)}{(1-\gamma)T} \\
&\leq \frac{\log |\mathcal{A}|}{\eta T} + \frac{1}{(1-\gamma)^2 T}.
\end{aligned}$$

The proof is completed using that $V^{(T)}(\rho) \geq V^{(T-1)}(\rho)$. ■

6 Function Approximation and Distribution Shift

We now analyze the case of using parametric policy classes:

$$\Pi = \{ \pi_\theta \mid \theta \in \mathbb{R}^d \},$$

where Π may not contain all stochastic policies (and it may not even contain an optimal policy). In contrast with the tabular results in the previous sections, the policy classes that we are often interested in are not fully expressive, e.g. $d \ll |\mathcal{S}||\mathcal{A}|$ (indeed $|\mathcal{S}|$ or $|\mathcal{A}|$ need not even be finite for the results in this section); in this sense, we are in the regime of function approximation.

We focus on obtaining *agnostic* results, where we seek to do as well as the best policy in this class (or as well as some other comparator policy). While we are interested in a solution to the (unconstrained) policy optimization problem

$$\max_{\theta \in \mathbb{R}^d} V^{\pi_\theta}(\rho),$$

(for a given initial distribution ρ), we will see that optimization with respect to a different distribution will be helpful, just as in the tabular case,

We will consider variants of the NPG update rule (15):

$$\theta \leftarrow \theta + \eta F_\rho(\theta)^\dagger \nabla_\theta V^\theta(\rho). \quad (16)$$

Our analysis will leverage a close connection between the NPG update rule (15) with the notion of *compatible function approximation* [Sutton et al., 1999], as formalized in Kakade [2001]. Specifically, it can be easily seen that:

$$F_\rho(\theta)^\dagger \nabla_\theta V^\theta(\rho) = \frac{1}{1-\gamma} w^*, \quad (17)$$

where w^* is a minimizer of the following regression problem:

$$w^* \in \operatorname{argmin}_w \mathbb{E}_{s \sim d_\rho^{\pi_\theta}, a \sim \pi_\theta(\cdot|s)} [(w^\top \nabla_\theta \log \pi_\theta(\cdot|s) - A^{\pi_\theta}(s, a))^2].$$

The above is a straightforward consequence of the first order optimality conditions (see (50)). The above regression problem can be viewed as “compatible” function approximation: we are approximating $A^{\pi_\theta}(s, a)$ using the $\nabla_\theta \log \pi_\theta(\cdot|s)$ as features. We also consider a variant of the above update rule, *Q-NPG*, where instead of using advantages in the above regression we use the *Q*-values.

This viewpoint provides a methodology for approximate updates, where we can solve the relevant regression problems with samples. Our main results establish the effectiveness of NPG updates where there is error both due to statistical estimation (where we may not use exact gradients) and approximation (due to using a parameterized function class); in particular, we provide a novel estimation/approximation decomposition relevant for the NPG algorithm. For these algorithms, we will first consider log linear policies classes (as a special case) and then move on to more general policy classes (such as neural policy classes). Finally, it is worth remarking that the results herein provide one of the first provable approximation guarantees where the error conditions required do not have explicit worst case dependencies over the state space.

6.1 NPG and Q-NPG Examples

In practice, the most common policy classes are of the form:

$$\Pi = \left\{ \pi_\theta(a|s) = \frac{\exp(f_\theta(s, a))}{\sum_{a' \in \mathcal{A}} \exp(f_\theta(s, a'))} \mid \theta \in \mathbb{R}^d \right\}, \quad (18)$$

where f_θ is a differentiable function. For example, the tabular softmax policy class is one where $f_\theta(s, a) = \theta_{s,a}$. Typically, f_θ is either a linear function or a neural network. Let us consider the NPG algorithm, and a variant Q -NPG, in each of these two cases.

6.1.1 Log-linear Policy Classes and Soft Policy Iteration

For any state-action pair (s, a) , suppose we have a feature mapping $\phi_{s,a} \in \mathbb{R}^d$. Each policy in the log-linear policy class is of the form:

$$\pi_\theta(a|s) = \frac{\exp(\theta \cdot \phi_{s,a})}{\sum_{a' \in \mathcal{A}} \exp(\theta \cdot \phi_{s,a'})},$$

with $\theta \in \mathbb{R}^d$. Here, we can take $f_\theta(s, a) = \theta \cdot \phi_{s,a}$.

With regards to compatible function approximation for the log-linear policy class, we have:

$$\nabla_\theta \log \pi_\theta(a|s) = \bar{\phi}_{s,a}^\theta, \text{ where } \bar{\phi}_{s,a}^\theta = \phi_{s,a} - \mathbb{E}_{a' \sim \pi_\theta(\cdot|s)}[\phi_{s,a'}],$$

that is, $\bar{\phi}_{s,a}^\theta$ is the centered version of $\phi_{s,a}$. With some abuse of notation, we accordingly also define $\bar{\phi}^\pi$ for any policy π . Here, using (17), the NPG update rule (16) is equivalent to:

$$\text{NPG: } \theta \leftarrow \theta + \eta w_\star, \quad w_\star \in \operatorname{argmin}_w \mathbb{E}_{s \sim d_\rho^{\pi_\theta}, a \sim \pi_\theta(\cdot|s)} \left[(A^{\pi_\theta}(s, a) - w \cdot \bar{\phi}_{s,a}^\theta)^2 \right].$$

(We have rescaled the learning rate η in comparison to (16)). Note that we recompute w_\star for every update of θ . Here, the compatible function approximation error measures the expressivity of our parameterization in how well linear functions of the parameterization can capture the policy's advantage function.

We also consider a variant of the NPG update rule (16), termed Q -NPG, where:

$$Q\text{-NPG: } \theta \leftarrow \theta + \eta w_\star, \quad w_\star \in \operatorname{argmin}_w \mathbb{E}_{s \sim d_\rho^{\pi_\theta}, a \sim \pi_\theta(\cdot|s)} \left[(Q^{\pi_\theta}(s, a) - w \cdot \phi_{s,a})^2 \right].$$

Note we do not center the features for Q -NPG; observe that $Q^\pi(s, a)$ is also not 0 in expectation under $\pi(\cdot|s)$, unlike the advantage function.

Remark 6.1. (NPG/ Q -NPG and Soft-Policy Iteration) We now see how we can view both NPG and Q -NPG as an incremental (soft) version of policy iteration, just as in Lemma 5.1 for the tabular case. Rather than writing the update rule in terms of the parameter θ , we can write an equivalent update rule directly in terms of the (log-linear) policy π :

$$\text{NPG: } \pi(a|s) \leftarrow \pi(a|s) \exp(w_\star \cdot \phi_{s,a}) / Z_s, \quad w_\star \in \operatorname{argmin}_w \mathbb{E}_{s \sim d_\rho^\pi, a \sim \pi(\cdot|s)} \left[(A^\pi(s, a) - w \cdot \bar{\phi}_{s,a}^\pi)^2 \right],$$

where Z_s is normalization constant. While the policy update uses the original features ϕ instead of $\bar{\phi}^\pi$, whereas the quadratic error minimization is terms of the centered features $\bar{\phi}^\pi$, this distinction is not relevant due to that we may also instead use $\bar{\phi}^\pi$ (in the policy update) which would result in an equivalent update; the normalization makes the update invariant to (constant) translations of the features. Similarly, an equivalent update for Q -NPG, where we update π directly rather than θ , is:

$$Q\text{-NPG: } \pi(a|s) \leftarrow \pi(a|s) \exp(w_\star \cdot \phi_{s,a}) / Z_s, \quad w_\star \in \operatorname{argmin}_w \mathbb{E}_{s \sim d_\rho^\pi, a \sim \pi(\cdot|s)} \left[(Q^\pi(s, a) - w \cdot \phi_{s,a})^2 \right].$$

Remark 6.2. (On the equivalence of NPG and Q -NPG) If it is the case that the compatible function approximation error is 0, then it is straightforward to verify that the NPG and Q -NPG are equivalent algorithms, in that their corresponding policy updates will be equivalent to each other.

6.1.2 Neural Policy Classes

Now suppose $f_\theta(s, a)$ is a neural network parameterized by $\theta \in \mathbb{R}^d$, where the policy class Π is of form in (18). Observe:

$$\nabla_\theta \log \pi_\theta(a|s) = g_\theta(s, a), \text{ where } g_\theta(s, a) = \nabla_\theta f_\theta(s, a) - \mathbb{E}_{a' \sim \pi_\theta(\cdot|s)}[\nabla_\theta f_\theta(s, a')],$$

and, using (17), the NPG update rule (16) is equivalent to:

$$\text{NPG: } \theta \leftarrow \theta + \eta w_\star, \quad w_\star \in \operatorname{argmin}_w \mathbb{E}_{s \sim d_\rho^{\pi_\theta}, a \sim \pi_\theta(\cdot|s)} \left[\left(A^{\pi_\theta}(s, a) - w \cdot g_\theta(s, a) \right)^2 \right]$$

(Again, we have rescaled the learning rate η in comparison to (16)).

The Q -NPG variant of this update rule is:

$$\text{Q-NPG: } \theta \leftarrow \theta + \eta w_\star, \quad w_\star \in \operatorname{argmin}_w \mathbb{E}_{s \sim d_\rho^{\pi_\theta}, a \sim \pi_\theta(\cdot|s)} \left[\left(Q^{\pi_\theta}(s, a) - w \cdot \nabla_\theta f_\theta(s, a) \right)^2 \right].$$

6.2 Q -NPG: Performance Bounds for Log-Linear Policies

For a state-action distribution ν , define:

$$L(w; \theta, \nu) := \mathbb{E}_{s, a \sim \nu} \left[\left(Q^{\pi_\theta}(s, a) - w \cdot \phi_{s, a} \right)^2 \right].$$

The iterates of the Q -NPG algorithm can be viewed as minimizing this loss under some (changing) distribution ν .

We now specify an approximate version of Q -NPG. It is helpful to consider a slightly more general version of the algorithm in the previous section, where instead of optimizing under a starting state distribution ρ , we have a different starting *state-action* distribution ν . Analogous to the definition of the state visitation measure, d_μ^π , we can define a visitation measure over states *and* actions induced by following π after $s_0, a_0 \sim \nu$. We overload notation using d_ν^π to also refer to the state-action visitation measure; precisely,

$$d_\nu^\pi(s, a) := (1 - \gamma) \mathbb{E}_{s_0, a_0 \sim \nu} \sum_{t=0}^{\infty} \gamma^t \Pr^\pi(s_t = s, a_t = a | s_0, a_0) \quad (19)$$

where $\Pr^\pi(s_t = s, a_t = a | s_0, a_0)$ is the probability that $s_t = s$ and $a_t = a$, after starting at state s_0 , taking action a_0 , and following π thereafter. While we overload notation for visitation distributions ($d_\mu^\pi(s)$ and $d_\nu^\pi(s, a)$) for notational convenience, note that the state-action measure d_ν^π uses the subscript ν , which is a state-action measure.

Q -NPG will be defined with respect to the *on-policy* state action measure starting with $s_0, a_0 \sim \nu$. As per our convention, we define

$$d^{(t)} := d_{\nu}^{\pi^{(t)}}.$$

The approximate version of this algorithm is:

$$\text{Approx. Q-NPG: } \theta^{(t+1)} = \theta^{(t)} + \eta w^{(t)}, \quad w^{(t)} \approx \operatorname{argmin}_{\|w\|_2 \leq W} L(w; \theta^{(t)}, d^{(t)}), \quad (20)$$

where the above update rule also permits us to constrain the norm of the update direction $w^{(t)}$ (alternatively, we could use ℓ_2 regularization as is also common in practice). The exact minimizer is denoted as:

$$w_{\star}^{(t)} \in \operatorname{argmin}_{\|w\|_2 \leq W} L(w; \theta^{(t)}, d^{(t)}).$$

Note that $w_{\star}^{(t)}$ depends on the current parameter $\theta^{(t)}$.

Our analysis will take into account both the *excess risk* (often also referred to as estimation error) and the *transfer error*. Here, the excess risk will be due to that $w^{(t)}$ may not be equal $w_{\star}^{(t)}$, and the approximation error will be due to that even the best linear fit using $w_{\star}^{(t)}$ may not perfectly match the Q -values, i.e. $L(w_{\star}^{(t)}; \theta^{(t)}; d^{(t)})$ is unlikely to be 0 in practical applications.

We now formalize these concepts in the following assumption:

Assumption 6.1 (Estimation/Transfer errors). *Fix a state distribution ρ ; a state-action distribution ν ; an arbitrary comparator policy π^* (not necessarily an optimal policy). With respect to π^* , define the state-action measure d^* as*

$$d^*(s, a) = d_{\rho}^{\pi^*}(s) \circ \operatorname{Unif}_{\mathcal{A}}(a)$$

i.e. d^* samples states from the comparators state visitation measure, $d_{\rho}^{\pi^*}$ and actions from the uniform distribution. Let us permit the sequence of iterates $w^{(0)}, w^{(1)}, \dots, w^{(T-1)}$ used by the Q -NPG algorithm to be random, where the randomness could be due to sample-based, estimation error. Suppose the following holds for all $t < T$:

1. (Excess risk) Assume that the estimation error is bounded as follows:

$$\mathbb{E} \left[L(w^{(t)}; \theta^{(t)}, d^{(t)}) - L(w_{\star}^{(t)}; \theta^{(t)}, d^{(t)}) \right] \leq \epsilon_{\text{stat}}$$

Note that using a sample based approach we would expect $\epsilon_{\text{stat}} = O(1/\sqrt{N})$ or better, where N is the number of samples used to estimate. $w_{\star}^{(t)}$ We formalize this in Corollary 6.2.

2. (Transfer error) Suppose that the best predictor $w_{\star}^{(t)}$ has an error bounded by ϵ_{bias} , in expectation, with respect to the comparator's measure of d^* . Specifically, assume:

$$\mathbb{E} \left[L(w_{\star}^{(t)}; \theta^{(t)}, d^*) \right] \leq \epsilon_{\text{bias}}.$$

We refer to ϵ_{bias} as the transfer error (or transfer bias); it is the error where relevant distribution is shifted to d^* . For the softmax policy parameterization for tabular MDPs, $\epsilon_{\text{bias}} = 0$ (see remark 6.4 for another example).

In both conditions, the expectations are with respect to the randomness in the sequence of iterates $w^{(0)}, w^{(1)}, \dots, w^{(T-1)}$, e.g. the approximate algorithm may be sample based.

Shortly, we discuss how the transfer error relates to the more standard approximation-estimation decomposition. Importantly, with the transfer error, it is always defined with respect to a single, fixed measure, d^* .

Assumption 6.2 (Relative condition number). *Consider the same ρ , ν , and π^* as in Assumption 6.1. With respect to any state-action distribution ν , define:*

$$\Sigma_\nu = \mathbb{E}_{s,a \sim \nu} [\phi_{s,a} \phi_{s,a}^\top],$$

and define

$$\sup_{w \in \mathbb{R}^d} \frac{w^\top \Sigma_{d^*} w}{w^\top \Sigma_\nu w} = \kappa.$$

Assume that κ is finite.

Remark 6.3 discusses why it is reasonable to expect that κ is not a quantity related to the size of the state space.⁷

Our main theorem below shows how the approximation error, the excess risk, and the conditioning, determine the final performance. Note that both the transfer error ϵ_{bias} and κ are defined with respect to the comparator policy π^* .

Theorem 6.1. (Agnostic learning with Q-NPG) *Fix a state distribution ρ ; a state-action distribution ν ; an arbitrary comparator policy π^* (not necessarily an optimal policy). Suppose Assumption 6.2 holds and $\|\phi_{s,a}\|_2 \leq B$ for all s, a . Suppose the Q-NPG update rule (in (20)) starts with $\theta^{(0)} = 0$, $\eta = \sqrt{2 \log |\mathcal{A}| / (B^2 W^2 T)}$, and the (random) sequence of iterates satisfies Assumption 6.1. We have that*

$$\mathbb{E} \left[\min_{t < T} \{V^{\pi^*}(\rho) - V^{(t)}(\rho)\} \right] \leq \frac{BW}{1-\gamma} \sqrt{\frac{2 \log |\mathcal{A}|}{T}} + \sqrt{\frac{4|\mathcal{A}|\kappa\epsilon_{\text{stat}}}{(1-\gamma)^3}} + \frac{\sqrt{4|\mathcal{A}|\epsilon_{\text{bias}}}}{1-\gamma}.$$

The proof is provided in Section 6.4.

Note when $\epsilon_{\text{bias}} = 0$, our convergence rate is $O(\sqrt{1/T})$ plus a term that depends on the excess risk; hence, provided we obtain enough samples, then ϵ_{stat} will also tend to 0, and we will be competitive with the comparison policy π^* . When $\epsilon_{\text{bias}} = 0$ and $\epsilon_{\text{stat}} = 0$, as in the tabular setting with exact gradients, the additional two terms become 0, consistent with Theorem 5.3 except that the convergence rate is $O(\sqrt{1/T})$ rather than the faster rate of $O(1/T)$. Obtaining a faster rate in the function approximation regime appears to require stronger conditions on how the approximation errors are controlled at each iteration.

⁷Technically, we only need the relative condition number $\sup_{w \in \mathbb{R}^d} \frac{w^\top \Sigma_{d^*} w}{w^\top \Sigma_{\pi^{(t)}} w}$ to be bounded for all t . We state this as a sufficient condition based on the initial distribution ν due to: this is more interpretable, and, as per Remark 6.3, this quantity can be bounded in a manner that is independent of the sequence of iterates produced by the algorithm.

The usual approximation-estimation error decomposition is that we can write our error as:

$$L(w^{(t)}; \theta^{(t)}, d^{(t)}) = \underbrace{L(w^{(t)}; \theta^{(t)}, d^{(t)}) - L(w_{\star}^{(t)}; \theta^{(t)}, d^{(t)})}_{\text{Excess risk}} + \underbrace{L(w_{\star}^{(t)}; \theta^{(t)}, d^{(t)})}_{\text{Approximation error}}$$

As we obtain more samples, we can drive the excess risk (the estimation error) to 0 (see Corollary 6.2). The approximation error above is due to modeling error. Importantly, for our Q -NPG performance bound, it is not this standard approximation error notion which is relevant, but it is this error under a different measure d^* , i.e. $L(w_{\star}^{(t)}; \theta^{(t)}, d^*)$. One appealing aspect about the transfer error is that this error is with respect to a fixed measure, namely d^* . Furthermore, in practice, modern machine learning methods often performs favorably with regards to transfer learning, substantially better than worst case theory might suggest.

The following corollary provides a performance bound in terms of the usual notion of approximation error, at the cost of also depending on the worst case distribution mismatch ratio. The corollary disentangles the estimation error from the approximation error.

Corollary 6.1. (*Estimation error/Approximation error bound for Q -NPG*) Consider the same setting as in Theorem 6.1. Rather than assuming the transfer error is bounded (part 2 in Assumption 6.1), suppose that, for all $t \leq T$,

$$\mathbb{E} \left[L(w_{\star}^{(t)}; \theta^{(t)}, d^{(t)}) \right] \leq \epsilon_{\text{approx}}.$$

We have that

$$\mathbb{E} \left[\min_{t < T} \{ V^{\pi^*}(\rho) - V^{(t)}(\rho) \} \right] \leq \frac{BW}{1-\gamma} \sqrt{\frac{2 \log |\mathcal{A}|}{T}} + \sqrt{\kappa \cdot \frac{4|\mathcal{A}|\epsilon_{\text{stat}}}{(1-\gamma)^3}} + \sqrt{\left\| \frac{d^*}{\nu} \right\|_{\infty} \cdot \frac{4|\mathcal{A}|\epsilon_{\text{approx}}}{(1-\gamma)^3}}.$$

Proof: We have the following crude upper bound on the transfer error:

$$L(w_{\star}^{(t)}; \theta^{(t)}, d^*) \leq \left\| \frac{d^*}{d^{(t)}} \right\|_{\infty} L(w_{\star}^{(t)}; \theta^{(t)}, d^{(t)}) \leq \frac{1}{1-\gamma} \left\| \frac{d^*}{\nu} \right\|_{\infty} L(w_{\star}^{(t)}; \theta^{(t)}, d^{(t)}),$$

where the last step uses the definition of $d^{(t)}$ (see (19)). This implies $\epsilon_{\text{bias}} \leq \frac{1}{1-\gamma} \left\| \frac{d^*}{\nu} \right\|_{\infty} \epsilon_{\text{approx}}$, and the corollary follows. ■

The above also shows the striking difference between the effects of estimation error and approximation error. The proof shows how the transfer error notion is weaker than previous conditions based on distribution mismatch coefficients or concentrability coefficients. Also, as discussed in Scherrer [2014], the (distribution mismatch) coefficient $\left\| \frac{d^*}{\nu} \right\|_{\infty}$ is already weaker than the more standard concentrability coefficients.

A few additional remarks are now in order. We now make a few observations with regards to κ .

Remark 6.3. (Dimension dependence in κ and the importance of ν) It is reasonable to think about κ as being dimension dependent (or worse), but it is not necessarily related to the size of the state

space. For example, if $\|\phi_{s,a}\|_2 \leq B$, then $\kappa \leq \frac{B^2}{\sigma_{\min}(\mathbb{E}_{s,a \sim \nu}[\phi_{s,a}\phi_{s,a}^\top])}$ though this bound may be pessimistic. Here, we also see the importance of choice of ν in having a small (relative) condition number; in particular, this is the motivation for considering the generalization which allows for a starting state-action distribution ν vs. just a starting state distribution μ (as we did in the tabular case). Roughly speaking, we desire a ν which provides good coverage over the features. As the following lemma shows, there always exists a universal distribution ν , which can be constructed only with knowledge of the feature set (without knowledge of d^*), such that $\kappa \leq d$.

Lemma 6.1. ($\kappa \leq d$ is always possible) *Let $\Phi = \{\phi(s, a) | (s, a) \in \mathcal{S} \times \mathcal{A}\} \subset \mathbb{R}^d$ and suppose Φ is a compact set. There always exists a state-action distribution ν , which is supported on at most d^2 state-action pairs and which can be constructed only with knowledge of Φ (without knowledge of the MDP or d^*), such that:*

$$\kappa \leq d.$$

Proof: The distribution can be found through constructing the minimal volume ellipsoid containing Φ , i.e. the Löwner-John ellipsoid [John, 1948]. In particular, this ν is supported on the contact points between this ellipsoid and Φ ; the lemma immediately follows from properties of this ellipsoid (e.g. see Ball [1997], Bubeck et al. [2012]). ■

It is also worth considering a more general example (beyond tabular MDPs) in which $\epsilon_{\text{bias}} = 0$ for the log-linear policy class.

Remark 6.4. ($\epsilon_{\text{bias}} = 0$ for “linear” MDPs) In the recent linear MDP model of Jin et al. [2019], Yang and Wang [2019], Jiang et al. [2017], where the transition dynamics are low rank, we have that $\epsilon_{\text{bias}} = 0$ provided we use the features of the linear MDP. Our guarantees also permit model misspecification of linear MDPs, with non worst-case approximation error where $\epsilon_{\text{bias}} \neq 0$.

Remark 6.5. (Comparison with POLITEX and EE-POLITEX) Compared with POLITEX [Abbasi-Yadkori et al., 2019a], Assumption 6.2 is substantially milder, in that it just assumes a good relative condition number for one policy rather than all possible policies (which cannot hold in general even for tabular MDPs). Changing this assumption to an analog of Assumption 6.2 is the main improvement in the analysis of the EE-POLITEX [Abbasi-Yadkori et al., 2019b] algorithm. They provide a regret bound for the average reward setting, which is qualitatively different from the suboptimality bound in the discounted setting that we study. They provide a specialized result for linear function approximation, similar to Theorem 6.1.

6.2.1 Q -NPG Sample Complexity

Assumption 6.3 (Episodic Sampling Oracle). *For a fixed state-action distribution ν , we assume the ability to: start at $s_0, a_0 \sim \nu$; continue to act thereafter in the MDP according to any policy π ; and terminate this “rollout” when desired. With this oracle, it is straightforward to obtain unbiased samples of $Q^\pi(s, a)$ (or $A^\pi(s, a)$) under $s, a \sim d_\nu^\pi$ for any π ; see Algorithms 1 and 3.*

Algorithm 2 provides a sample based version of the Q -NPG algorithm; it simply uses stochastic projected gradient ascent within each iteration. The following corollary shows this algorithm suffices to obtain an accurate sample based version of Q -NPG.

Algorithm 1 Sampler for: $s, a \sim d_\nu^\pi$ and unbiased estimate of $Q^\pi(s, a)$

Require: Starting state-action distribution ν .

- 1: Sample $s_0, a_0 \sim \nu$.
 - 2: Sample $s, a \sim d_\nu^\pi$ as follows: at every timestep h , with probability γ , act according to π ; else, accept (s_h, a_h) as the sample and proceed to Step 4. See (19).
 - 3: From s_h, a_h , continue to execute π , and use a termination probability of $1 - \gamma$. Upon termination, set $\widehat{Q}^\pi(s_h, a_h)$ as the *undiscounted* sum of rewards from time h onwards.
 - 4: **return** (s_h, a_h) and $\widehat{Q}^\pi(s_h, a_h)$.
-

Corollary 6.2. (Sample complexity of Q-NPG) Assume we are in the setting of Theorem 6.1 and that we have access to an episodic sampling oracle (i.e. Assumption 6.3). Suppose that the Sample Based Q-NPG Algorithm (Algorithm 2) is run for T iterations, with N gradient steps per iteration, with an appropriate setting of the learning rates η and α . We have that:

$$\begin{aligned} & \mathbb{E} \left[\min_{t < T} \{V^{\pi^*}(\rho) - V^{(t)}(\rho)\} \right] \\ & \leq \frac{BW}{1-\gamma} \sqrt{\frac{2 \log |\mathcal{A}|}{T}} + \sqrt{\frac{8\kappa |\mathcal{A}| BW (BW + 1)}{(1-\gamma)^4}} \frac{1}{N^{1/4}} + \frac{\sqrt{4|\mathcal{A}|\epsilon_{\text{bias}}}}{1-\gamma}. \end{aligned}$$

Furthermore, since each episode has expected length $2/(1-\gamma)$, the expected number of total samples used by Q-NPG is $2NT/(1-\gamma)$.

Proof: Note that our sampled gradients are bounded by $G := 2B(BW + \frac{1}{1-\gamma})$. Using $\alpha = \frac{W}{G\sqrt{N}}$, a standard analysis for stochastic projected gradient ascent (Theorem E.3) shows that:

$$\epsilon_{\text{stat}} \leq \frac{2BW(BW + \frac{1}{1-\gamma})}{\sqrt{N}}.$$

The proof is completed via substitution. ■

Remark 6.6. (Improving the scaling with N) Our current rate of convergence is $1/N^{1/4}$ due to our use of stochastic projected gradient ascent. Instead, for the least squares estimator, ϵ_{stat} would be $O(d/N)$ provided certain further regularity assumptions hold (a bound on the minimal eigenvalue of Σ_ν would be sufficient but not necessary. See Hsu et al. [2014] for such conditions). With such further assumptions, our rate of convergence would be $O(1/\sqrt{N})$.

6.3 NPG: Performance Bounds for Smooth Policy Classes

We now return to the analyzing the standard NPG update rule, which uses advantages rather than Q -values (see Section 6.1). It is helpful to define

$$L_A(w; \theta, \nu) := \mathbb{E}_{s, a \sim \nu} \left[\left(A^{\pi_\theta}(s, a) - w \cdot \nabla_\theta \log \pi_\theta(a|s) \right)^2 \right].$$

Algorithm 2 Sample-based Q -NPG for Log-linear Policies

Require: Learning rate η ; SGD learning rate α ; number of SGD iterations N

- 1: Initialize $\theta^{(0)} = 0$.
- 2: **for** $t = 0, 1, \dots, T - 1$ **do**
- 3: Initialize $w_0 = 0$
- 4: **for** $n = 0, 1, \dots, N - 1$ **do**
- 5: Call Algorithm 1 to obtain $s, a \sim d^{(t)}$ and an unbiased estimate $\widehat{Q}(s, a)$.
- 6: Update:

$$w_{n+1} = \text{Proj}_{\mathcal{W}} \left(w_n - 2\alpha \left(w_n \cdot \phi_{s,a} - \widehat{Q}(s, a) \right) \phi_{s,a} \right)$$

where $\mathcal{W} = \{w : \|w\|_2 \leq W\}$.

- 7: **end for**
 - 8: Set $\widehat{w}^{(t)} = \frac{1}{N} \sum_{n=1}^N w_n$.
 - 9: Update $\theta^{(t+1)} = \theta^{(t)} + \eta \widehat{w}^{(t)}$.
 - 10: **end for**
-

where ν is state-action distribution, and the subscript of A denotes the loss function uses advantages (rather than Q -values). The iterates of the NPG algorithm can be viewed as minimizing this loss under some appropriately chosen measure.

We now consider an approximate version of the NPG update rule:

$$\text{Approx. NPG: } \theta^{(t+1)} = \theta^{(t)} + \eta w^{(t)}, \quad w^{(t)} \approx \operatorname{argmin}_{\|w\|_2 \leq W} L_A(w; \theta^{(t)}, d^{(t)}), \quad (21)$$

where again we use the on-policy, fitting distribution $d^{(t)}$. As with Q -NPG, we also permit the use of a starting state-action distribution ν as opposed to just a starting state distribution (see Remark 6.3). Again, we let $w_\star^{(t)}$ denote the minimizer, i.e. $w_\star^{(t)} \in \operatorname{argmin}_{\|w\|_2 \leq W} L_A(w; \theta^{(t)}, d^{(t)})$.

For this section, our analysis will focus on more general policy classes, beyond log-linear policy classes. In particular, we make the following smoothness assumption on the policy class:

Assumption 6.4. (*Policy Smoothness*) Assume for all $s \in \mathcal{S}$ and $a \in \mathcal{A}$ that $\log \pi_\theta(a|s)$ is a β -smooth function of θ (to recall the definition of smoothness, see (24)).

It is not to difficult to verify that the tabular softmax policy parameterization is a 1-smooth policy class in the above sense. The more general class of log-linear policies is also smooth as we remark below.

Remark 6.7. (Smoothness of the log-linear policy class) For the log-linear policy class (see Section 6.1.1), smoothness is implied if the features ϕ have bounded Euclidean norm. Precisely, if the feature mapping ϕ satisfies $\|\phi_{s,a}\|_2 \leq B$, then it is not difficult to verify that $\log \pi_\theta(a|s)$ is a B^2 -smooth function.

For any state-action distribution ν , define:

$$\Sigma_\nu^\theta = \mathbb{E}_{s,a \sim \nu} \left[\nabla_\theta \log \pi_\theta(a|s) (\nabla_\theta \log \pi_\theta(a|s))^\top \right]$$

and, again, we use $\Sigma_v^{(t)}$ as shorthand for $\Sigma_v^{\theta^{(t)}}$.

Assumption 6.5. (*Estimation/Transfer/Conditioning*) Fix a state distribution ρ ; a state-action distribution ν ; an arbitrary comparator policy π^* (not necessarily an optimal policy). With respect to π^* , define the state-action measure d^* as

$$d^*(s, a) = d_\rho^{\pi^*}(s)\pi^*(a|s).$$

Note that, in comparison to Assumption 6.1, d^* is the state-action visitation measure of the comparator policy. Let us permit the sequence of iterates $w^{(0)}, w^{(1)}, \dots, w^{(T-1)}$ used by the NPG algorithm to be random, where the randomness could be due to sample-based, estimation error. Suppose the following holds for all $t < T$:

1. (*Excess risk*) Assume the estimation error is bounded as:

$$\mathbb{E} \left[L_A(w^{(t)}; \theta^{(t)}, d^{(t)}) - L_A(w_\star^{(t)}; \theta^{(t)}, d^{(t)}) \mid \theta^{(t)} \right] \leq \epsilon_{\text{stat}}$$

i.e. the above conditional expectation is bounded (with probability one).⁸ As we see in Corollary 6.2, we can guarantee ϵ_{stat} to drop as $\sqrt{1/N}$.

2. (*Transfer error*) Suppose that:

$$\mathbb{E} \left[L_A(w_\star^{(t)}; \theta^{(t)}, d^*) \right] \leq \epsilon_{\text{bias}}.$$

3. (*Relative condition number*) For all iterations t , assume the average relative condition number is bounded as follows:

$$\mathbb{E} \left[\sup_{w \in \mathbb{R}^d} \frac{w^\top \Sigma_{d^*}^{(t)} w}{w^\top \Sigma_\nu^{(t)} w} \right] \leq \kappa. \quad (22)$$

Note that term inside the expectation is a random quantity as $\theta^{(t)}$ is random.

In the above conditions, the expectation is with respect to the randomness in the sequence of iterates $w^{(0)}, w^{(1)}, \dots, w^{(T-1)}$.

Analogous to our Q-NPG theorem, our main theorem for NPG shows how the transfer error is relevant in addition the statistical error ϵ_{stat} .

Theorem 6.2. (*Agnostic learning with NPG*) Fix a state distribution ρ ; a state-action distribution ν ; an arbitrary comparator policy π^* (not necessarily an optimal policy). Suppose Assumption 6.4 holds. Suppose the NPG update rule (in (21)) starts with $\pi^{(0)}$ being the uniform distribution (at each state), $\eta = \sqrt{2 \log |\mathcal{A}| / (\beta W^2 T)}$, and the (random) sequence of iterates satisfies Assumption 6.5. We have that

$$\mathbb{E} \left[\min_{t < T} \{ V^{\pi^*}(\rho) - V^{(t)}(\rho) \} \right] \leq \frac{W}{1-\gamma} \sqrt{\frac{2\beta \log |\mathcal{A}|}{T}} + \sqrt{\frac{\kappa \epsilon_{\text{stat}}}{(1-\gamma)^3}} + \frac{\sqrt{\epsilon_{\text{bias}}}}{1-\gamma}.$$

⁸The use of a conditional expectation here (vs. the unconditional one in Assumption 6.1) permits the assumption to hold even in settings where we may reuse data in the sample-based approximation of L_A . Also, the expectation over the iterates allows a more natural assumption on the relative condition number, relevant for the more general case of smooth policies.

The proof is provided in Section 6.4.

Remark 6.8. (The $|\mathcal{A}|$ dependence: NPG vs. Q -NPG) Observe there is no polynomial dependence on $|\mathcal{A}|$ in the rate for NPG (in contrast to Theorem 6.1); also observe that here we define d^* as the state-action distribution of π^* in Assumption 6.5, as opposed to a uniform distribution over the actions, as in Assumption 6.1. The main difference arises in the analysis in that, even for Q -NPG, we need to bound the error in fitting the advantage estimates; this leads to the dependence on $|\mathcal{A}|$ (which can be removed with a path dependent bound, i.e. a bound which depends on the sequence of iterates produced by the algorithm)⁹. For NPG, the direct fitting of the advantage function sidesteps this conversion step. Note that the relative condition number assumption in Q -NPG (Assumption 6.2) is a weaker assumption, due to that it can be bounded independently of the path of the algorithm (see Remark 6.2), while NPG’s centering of the features makes the assumption on the relative condition number depend on the path of the algorithm.

Remark 6.9. (Generalizing Q -NPG for smooth policies) A similar reasoning as the analysis here can be also used to establish a convergence result for the Q -NPG algorithm in this more general setting of smooth policy classes. Concretely, we can analyze the Q -NPG update described for neural policy classes in Section 6.1.2, assuming that the function f_θ is Lipschitz-continuous in θ . Like for Theorem 6.2, the main modification is that Assumption 6.2 on relative condition numbers is now defined using the covariance matrix for the features $f_\theta(s, a)$, which depend on θ , as opposed to some a feature map $\phi(s, a)$ in the log-linear case. The rest of the analysis follows with an appropriate adaptation of the results above.

6.3.1 NPG Sample Complexity

Algorithm 4 provides a sample based version of the NPG algorithm, again using stochastic projected gradient ascent; it uses a slight modification of the Q -NPG algorithm to obtain unbiased gradient estimates. The following corollary shows that this algorithm provides an accurate sample based version of NPG.

Corollary 6.3. (*Sample complexity of NPG*) Assume we are in the setting of Theorem 6.2 and that we have access to an episodic sampling oracle (i.e. Assumption 6.3). Suppose that the Sample Based NPG Algorithm (Algorithm 4) is run for T iterations, with N gradient steps per iteration. Also, suppose that $\|\nabla_\theta \log \pi^{(t)}(a|s)\|_2 \leq B$ holds with probability one. There exists a setting of η and α such that:

$$\begin{aligned} & \mathbb{E} \left[\min_{t < T} \{V^{\pi^*}(\rho) - V^{(t)}(\rho)\} \right] \\ & \leq \frac{W}{1 - \gamma} \sqrt{\frac{2\beta \log |\mathcal{A}|}{T}} + \sqrt{\frac{8\kappa BW(BW + 1)}{(1 - \gamma)^4} \frac{1}{N^{1/4}}} + \frac{\sqrt{\epsilon_{\text{bias}}}}{1 - \gamma}. \end{aligned}$$

Furthermore, since each episode has expected length $2/(1 - \gamma)$, the expected number of total samples used by NPG is $2NT/(1 - \gamma)$.

⁹For Q -NPG, we have to bound two distribution shift terms to both π^* and $\pi^{(t)}$ at step t of the algorithm.

Algorithm 3 Sampler for: $s, a \sim d_\nu^\pi$ and unbiased estimate of $A^\pi(s, a)$

Require: Starting state-action distribution ν .

- 1: Set $\widehat{Q}^\pi = 0$ and $\widehat{V}^\pi = 0$.
 - 2: Start at state $s_0 \sim \nu$. Sample $a_0 \sim \nu(\cdot|s_0)$ (though do not necessarily execute a_0).
 - 3: (d_ν^π sampling) At every timestep $h \geq 0$,
 - With probability γ , execute a_h , transition to s_{h+1} , and sample $a_{h+1} \sim \pi(\cdot|s_{h+1})$.
 - Else accept (s_h, a_h) as the sample and proceed to Step 4.
 - 4: ($A^\pi(s, a)$ sampling) Set SampleQ = True with probability $1/2$.
 - If SampleQ = True, execute a_h at state s_h and then continue executing π with a termination probability of $1 - \gamma$. Upon termination, set \widehat{Q}^π as the *undiscounted* sum of rewards from time h onwards.
 - Else sample $a'_h \sim \pi(\cdot|s_h)$. Then execute a'_h at state s_h and then continue executing π with a termination probability of $1 - \gamma$. Upon termination, set \widehat{V}^π as the *undiscounted* sum of rewards from time h onwards.
 - 5: **return** (s_h, a_h) and $\widehat{A}^\pi(s_h, a_h) = 2(\widehat{Q}^\pi - \widehat{V}^\pi)$.
-

Proof: Let us see that the update direction in Step 6 of Algorithm 4 uses an unbiased estimate of the true gradient of the loss function L_A :

$$\begin{aligned}
& 2\mathbb{E}_{s,a \sim d^{(t)}} \left[\left(w_n \cdot \nabla_\theta \log \pi^{(t)}(a|s) - \widehat{A}(s, a) \right) \nabla_\theta \log \pi^{(t)}(a|s) \right] \\
&= 2\mathbb{E}_{s,a \sim d^{(t)}} \left[\left(w_n \cdot \nabla_\theta \log \pi^{(t)}(a|s) - \mathbb{E}[\widehat{A}(s, a)|s, a] \right) \nabla_\theta \log \pi^{(t)}(a|s) \right] \\
&= 2\mathbb{E}_{s,a \sim d^{(t)}} \left[\left(w_n \cdot \nabla_\theta \log \pi^{(t)}(a|s) - A^{(t)}(s, a) \right) \nabla_\theta \log \pi^{(t)}(a|s) \right] \\
&= \nabla_w L_A(w_n; \theta^{(t)}, d^{(t)})
\end{aligned}$$

where the last step follows due to that sampling procedure in Algorithm 3 produces a conditionally unbiased estimate.

Since $\|\nabla_\theta \log \pi^{(t)}(a|s)\|_2 \leq B$ and since $\widehat{A}(s, a) \leq 2/(1 - \gamma)$, our sampled gradients are bounded by $G := 8B(BW + \frac{1}{1-\gamma})$. The remainder of the proof follows that of Corollary 6.2 \blacksquare

6.4 Analysis

We first proceed by providing a general analysis of NPG, for arbitrary sequences. We then specialize it to complete the proof of our two main theorems in this section.

Algorithm 4 Sample-based NPG

Require: Learning rate η ; SGD learning rate α ; number of SGD iterations N

- 1: Initialize $\theta^{(0)} = 0$.
- 2: **for** $t = 0, 1, \dots, T - 1$ **do**
- 3: Initialize $w_0 = 0$
- 4: **for** $n = 0, 1, \dots, N - 1$ **do**
- 5: Call Algorithm 3 to obtain $s, a \sim d^{(t)}$, and an unbiased estimate $\widehat{A}(s, a)$ of $A^{(t)}(s, a)$.
- 6: Update:

$$w_{n+1} = \text{Proj}_{\mathcal{W}} \left(w_n - 2\alpha \left(w_n \cdot \nabla_{\theta} \log \pi^{(t)}(a|s) - \widehat{A}(s, a) \right) \nabla_{\theta} \log \pi^{(t)}(a|s) \right),$$

where $\mathcal{W} = \{w : \|w\|_2 \leq W\}$

- 7: **end for**
 - 8: Set $\widehat{w}^{(t)} = \frac{1}{N} \sum_{n=1}^N w_n$.
 - 9: Update $\theta^{(t+1)} = \theta^{(t)} + \eta \widehat{w}^{(t)}$.
 - 10: **end for**
-

6.4.1 The NPG “Regret Lemma”

It is helpful for us to consider NPG more abstractly, as an update rule of the form

$$\theta^{(t+1)} = \theta^{(t)} + \eta w^{(t)}. \quad (23)$$

We will now provide a lemma where $w^{(t)}$ is an *arbitrary* (bounded) sequence, which will be helpful when specialized.

Recall a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is said to be β -smooth if for all $x, x' \in \mathbb{R}^d$:

$$\|\nabla f(x) - \nabla f(x')\|_2 \leq \beta \|x - x'\|_2,$$

and, due to Taylor’s theorem, recall that this implies:

$$\left| f(x') - f(x) - \nabla f(x) \cdot (x' - x) \right| \leq \frac{\beta}{2} \|x' - x\|_2^2. \quad (24)$$

The following analysis of NPG is based on the mirror-descent approach developed in [Even-Dar et al., 2009], which motivates us to refer to it as a “regret lemma”.

Lemma 6.2. (*NPG Regret Lemma*) Fix a comparison policy $\tilde{\pi}$ and a state distribution ρ . Assume for all $s \in \mathcal{S}$ and $a \in \mathcal{A}$ that $\log \pi_{\theta}(a|s)$ is a β -smooth function of θ . Consider the update rule (23), where $\pi^{(0)}$ is the uniform distribution (for all states) and where the sequence of weights $w^{(0)}, \dots, w^{(T)}$, satisfies $\|w^{(t)}\|_2 \leq W$ (but is otherwise arbitrary). Define:

$$\text{err}_t = \mathbb{E}_{s \sim \tilde{d}} \mathbb{E}_{a \sim \tilde{\pi}(\cdot|s)} \left[A^{(t)}(s, a) - w^{(t)} \cdot \nabla_{\theta} \log \pi^{(t)}(a|s) \right].$$

We have that:

$$\min_{t < T} \{V^{\tilde{\pi}}(\rho) - V^{(t)}(\rho)\} \leq \frac{1}{1-\gamma} \left(\frac{\log |\mathcal{A}|}{\eta T} + \frac{\eta\beta W^2}{2} + \frac{1}{T} \sum_{t=0}^{T-1} \text{err}_t \right).$$

Proof: By smoothness (see (24)),

$$\begin{aligned} \log \frac{\pi^{(t+1)}(a|s)}{\pi^{(t)}(a|s)} &\geq \nabla_{\theta} \log \pi^{(t)}(a|s) \cdot (\theta^{(t+1)} - \theta^{(t)}) - \frac{\beta}{2} \|\theta^{(t+1)} - \theta^{(t)}\|_2^2 \\ &= \eta \nabla_{\theta} \log \pi^{(t)}(a|s) \cdot w^{(t)} - \eta^2 \frac{\beta}{2} \|w^{(t)}\|_2^2. \end{aligned}$$

We use \tilde{d} as shorthand for $d_{\rho}^{\tilde{\pi}}$ (note ρ and $\tilde{\pi}$ are fixed); for any policy π , we also use π_s as shorthand for the vector $\pi(\cdot|s)$. Using the performance difference lemma (Lemma 3.2),

$$\begin{aligned} &\mathbb{E}_{s \sim \tilde{d}} (\text{KL}(\tilde{\pi}_s || \pi_s^{(t)}) - \text{KL}(\tilde{\pi}_s || \pi_s^{(t+1)})) \\ &= \mathbb{E}_{s \sim \tilde{d}} \mathbb{E}_{a \sim \tilde{\pi}(\cdot|s)} \left[\log \frac{\pi^{(t+1)}(a|s)}{\pi^{(t)}(a|s)} \right] \\ &\geq \eta \mathbb{E}_{s \sim \tilde{d}} \mathbb{E}_{a \sim \tilde{\pi}(\cdot|s)} [\nabla_{\theta} \log \pi^{(t)}(a|s) \cdot w^{(t)}] - \eta^2 \frac{\beta}{2} \|w^{(t)}\|_2^2 \quad (\text{using previous display}) \\ &= \eta \mathbb{E}_{s \sim \tilde{d}} \mathbb{E}_{a \sim \tilde{\pi}(\cdot|s)} [A^{(t)}(s, a)] - \eta^2 \frac{\beta}{2} \|w^{(t)}\|_2^2 \\ &\quad + \eta \mathbb{E}_{s \sim \tilde{d}} \mathbb{E}_{a \sim \tilde{\pi}(\cdot|s)} [\nabla_{\theta} \log \pi^{(t)}(a|s) \cdot w^{(t)} - A^{(t)}(s, a)] \\ &= (1-\gamma)\eta \left(V^{\tilde{\pi}}(\rho) - V^{(t)}(\rho) \right) - \eta^2 \frac{\beta}{2} \|w^{(t)}\|_2^2 - \eta \text{err}_t \end{aligned}$$

Rearranging, we have:

$$V^{\tilde{\pi}}(\rho) - V^{(t)}(\rho) \leq \frac{1}{1-\gamma} \left(\frac{1}{\eta} \mathbb{E}_{s \sim \tilde{d}} (\text{KL}(\tilde{\pi}_s || \pi_s^{(t)}) - \text{KL}(\tilde{\pi}_s || \pi_s^{(t+1)})) + \frac{\eta\beta}{2} W^2 + \text{err}_t \right)$$

Proceeding,

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} (V^{\tilde{\pi}}(\rho) - V^{(t)}(\rho)) &\leq \frac{1}{\eta T(1-\gamma)} \sum_{t=0}^{T-1} \mathbb{E}_{s \sim \tilde{d}} (\text{KL}(\tilde{\pi}_s || \pi_s^{(t)}) - \text{KL}(\tilde{\pi}_s || \pi_s^{(t+1)})) \\ &\quad + \frac{1}{T(1-\gamma)} \sum_{t=0}^{T-1} \left(\frac{\eta\beta W^2}{2} + \text{err}_t \right) \\ &\leq \frac{\mathbb{E}_{s \sim \tilde{d}} \text{KL}(\tilde{\pi}_s || \pi_s^{(0)})}{\eta T(1-\gamma)} + \frac{\eta\beta W^2}{2(1-\gamma)} + \frac{1}{T(1-\gamma)} \sum_{t=0}^{T-1} \text{err}_t \\ &\leq \frac{\log |\mathcal{A}|}{\eta T(1-\gamma)} + \frac{\eta\beta W^2}{2(1-\gamma)} + \frac{1}{T(1-\gamma)} \sum_{t=0}^{T-1} \text{err}_t, \end{aligned}$$

which completes the proof. ■

6.4.2 Proofs of Theorem 6.1 and 6.2

Proof: (of Theorem 6.1) Using the NPG regret lemma (Lemma 6.2) and the smoothness of the log-linear policy class (see Example 6.7),

$$\mathbb{E} \left[\min_{t < T} \{V^{\pi^*}(\rho) - V^{(t)}(\rho)\} \right] \leq \frac{BW}{1-\gamma} \sqrt{\frac{2 \log |\mathcal{A}|}{T}} + \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \text{err}_t \right].$$

where we have used our setting of η .

We make the following decomposition of err_t :

$$\begin{aligned} \text{err}_t &= \mathbb{E}_{s \sim d_p^*, a \sim \pi^*(\cdot|s)} \left[A^{(t)}(s, a) - w_\star^{(t)} \cdot \nabla_\theta \log \pi^{(t)}(a|s) \right] \\ &\quad + \mathbb{E}_{s \sim d_p^*, a \sim \pi^*(\cdot|s)} \left[\left(w_\star^{(t)} - w^{(t)} \right) \cdot \nabla_\theta \log \pi^{(t)}(a|s) \right]. \end{aligned}$$

For the first term, using that $\nabla_\theta \log \pi_\theta(a|s) = \phi_{s,a} - \mathbb{E}_{a' \sim \pi_\theta(\cdot|s)}[\phi_{s,a'}]$ (see Section 6.1.1), we have:

$$\begin{aligned} &\mathbb{E}_{s \sim d_p^*, a \sim \pi^*(\cdot|s)} \left[A^{(t)}(s, a) - w_\star^{(t)} \cdot \nabla_\theta \log \pi^{(t)}(a|s) \right] \\ &= \mathbb{E}_{s \sim d_p^*, a \sim \pi^*(\cdot|s)} \left[Q^{(t)}(s, a) - w_\star^{(t)} \cdot \phi_{s,a} \right] - \mathbb{E}_{s \sim d_p^*, a' \sim \pi^{(t)}(\cdot|s)} \left[Q^{(t)}(s, a') - w_\star^{(t)} \cdot \phi_{s,a'} \right] \\ &\leq \sqrt{\mathbb{E}_{s \sim d_p^*, a \sim \pi^*(\cdot|s)} \left(Q^{(t)}(s, a) - w_\star^{(t)} \cdot \phi_{s,a} \right)^2} + \sqrt{\mathbb{E}_{s \sim d_p^*, a' \sim \pi^{(t)}(\cdot|s)} \left(Q^{(t)}(s, a') - w_\star^{(t)} \cdot \phi_{s,a'} \right)^2} \\ &\leq 2\sqrt{|\mathcal{A}| \mathbb{E}_{s \sim d_p^*, a \sim \text{Unif}_{\mathcal{A}}} \left[\left(Q^{(t)}(s, a) - w_\star^{(t)} \cdot \phi_{s,a} \right)^2 \right]} = 2\sqrt{|\mathcal{A}| L(w_\star^{(t)}; \theta^{(t)}, d^*)}. \end{aligned} \quad (25)$$

where we have used the definition of d^* and $L(w_\star^{(t)}; \theta^{(t)}, d^*)$ in the last step.

For the second term, let us now show that:

$$\begin{aligned} &\mathbb{E}_{s \sim d_p^*, a \sim \pi^*(\cdot|s)} \left[\left(w_\star^{(t)} - w^{(t)} \right) \cdot \nabla_\theta \log \pi^{(t)}(a|s) \right] \\ &\leq 2\sqrt{\frac{|\mathcal{A}| \kappa}{1-\gamma} \left(L(w^{(t)}; \theta^{(t)}, d^{(t)}) - L(w_\star^{(t)}; \theta^{(t)}, d^{(t)}) \right)} \end{aligned} \quad (26)$$

To see this, first observe that a similar argument to the above leads to:

$$\begin{aligned} &\mathbb{E}_{s \sim d_p^*, a \sim \pi^*(\cdot|s)} \left[\left(w_\star^{(t)} - w^{(t)} \right) \cdot \nabla_\theta \log \pi^{(t)}(a|s) \right] \\ &= \mathbb{E}_{s \sim d_p^*, a \sim \pi^*(\cdot|s)} \left[\left(w_\star^{(t)} - w^{(t)} \right) \cdot \phi_{s,a} \right] - \mathbb{E}_{s \sim d_p^*, a' \sim \pi^{(t)}(\cdot|s)} \left[\left(w_\star^{(t)} - w^{(t)} \right) \cdot \phi_{s,a'} \right] \\ &\leq 2\sqrt{|\mathcal{A}| \mathbb{E}_{s, a \sim d^*} \left[\left(\left(w_\star^{(t)} - w^{(t)} \right) \cdot \phi_{s,a} \right)^2 \right]} = 2\sqrt{|\mathcal{A}| \cdot \|w_\star^{(t)} - w^{(t)}\|_{\Sigma_{d^*}}^2}, \end{aligned}$$

where we use the notation $\|x\|_M^2 := x^\top M x$ for a matrix M and a vector x . From the definition of κ ,

$$\|w_\star^{(t)} - w^{(t)}\|_{\Sigma_{d^*}}^2 \leq \kappa \|w_\star^{(t)} - w^{(t)}\|_{\Sigma_\nu}^2 \leq \frac{\kappa}{1-\gamma} \|w_\star^{(t)} - w^{(t)}\|_{\Sigma_{d^{(t)}}}^2$$

using that $(1 - \gamma)\nu \leq d_\nu^{\pi^{(t)}}$ (see (19)). Due to that $w_\star^{(t)}$ minimizes $L(w; \theta^{(t)}, d^{(t)})$ over the set $\mathcal{W} := \{w : \|w\|_2 \leq W\}$, for any $w \in \mathcal{W}$ the first-order optimality conditions for $w_\star^{(t)}$ imply that:

$$(w - w_\star^{(t)}) \cdot \nabla L(w_\star^{(t)}; \theta^{(t)}, d^{(t)}) \geq 0.$$

Therefore, for any $w \in \mathcal{W}$,

$$\begin{aligned} & L(w; \theta^{(t)}, d^{(t)}) - L(w_\star^{(t)}; \theta^{(t)}, d^{(t)}) \\ &= \mathbb{E}_{s, a \sim d^{(t)}} \left[(w \cdot \phi(s, a) - w_\star \cdot \phi(s, a) + w_\star \cdot \phi(s, a) - Q^{(t)}(s, a))^2 \right] - L(w_\star^{(t)}; \theta^{(t)}, d^{(t)}) \\ &= \mathbb{E}_{s, a \sim d^{(t)}} \left[(w \cdot \phi(s, a) - w_\star \cdot \phi(s, a))^2 \right] + 2(w - w_\star) \mathbb{E}_{s, a \sim d^{(t)}} [\phi(s, a)(w_\star \cdot \phi(s, a) - Q^{(t)}(s, a))] \\ &= \|w - w_\star^{(t)}\|_{\Sigma_{d^{(t)}}}^2 + (w - w_\star^{(t)}) \cdot \nabla L(w_\star^{(t)}; \theta^{(t)}, d^{(t)}) \\ &\geq \|w - w_\star^{(t)}\|_{\Sigma_{d^{(t)}}}^2. \end{aligned}$$

Noting that $w^{(t)} \in \mathcal{W}$ by construction in Algorithm 4 yields the claimed bound on the second term in (26).

Using the bounds on the first and second terms in (25) and (26), along with concavity of the square root function, we have that:

$$\mathbb{E}[\text{err}_t] \leq 2\sqrt{|\mathcal{A}|\mathbb{E} \left[L(w_\star^{(t)}; \theta^{(t)}, d^\star) \right]} + 2\sqrt{\frac{|\mathcal{A}|\kappa}{1 - \gamma} \mathbb{E} \left[L(w^{(t)}; \theta^{(t)}, d^{(t)}) - L(w_\star^{(t)}; \theta^{(t)}, d^{(t)}) \right]}.$$

The proof is completed by substitution and using our assumptions on ϵ_{stat} and ϵ_{bias} . ■

The following proof for the NPG algorithm follows along similar lines.

Proof: (of Theorem 6.2) Using the NPG regret lemma and our setting of η ,

$$\mathbb{E} \left[\min_{t < T} \{V^{\pi^\star}(\rho) - V^{(t)}(\rho)\} \right] \leq \frac{W}{1 - \gamma} \sqrt{\frac{2\beta \log |\mathcal{A}|}{T}} + \mathbb{E} \left[\frac{1}{T} \sum_{t=0}^{T-1} \text{err}_t \right].$$

where the expectation is with respect to the sequence of iterates $w^{(0)}, w^{(1)}, \dots, w^{(T-1)}$.

Again, we make the following decomposition of err_t :

$$\begin{aligned} \text{err}_t &= \mathbb{E}_{s \sim d_\rho^\star, a \sim \pi^\star(\cdot|s)} \left[A^{(t)}(s, a) - w_\star^{(t)} \cdot \nabla_\theta \log \pi^{(t)}(a|s) \right] \\ &\quad + \mathbb{E}_{s \sim d_\rho^\star, a \sim \pi^\star(\cdot|s)} \left[(w_\star^{(t)} - w^{(t)}) \cdot \nabla_\theta \log \pi^{(t)}(a|s) \right]. \end{aligned}$$

For the first term,

$$\begin{aligned} & \mathbb{E}_{s \sim d_\rho^\star, a \sim \pi^\star(\cdot|s)} \left[A^{(t)}(s, a) - w_\star^{(t)} \cdot \nabla_\theta \log \pi^{(t)}(a|s) \right] \\ &\leq \sqrt{\mathbb{E}_{s \sim d_\rho^\star, a \sim \pi^\star(\cdot|s)} \left[\left(A^{(t)}(s, a) - w_\star^{(t)} \cdot \phi_{s,a} \right)^2 \right]} = \sqrt{L_A(w_\star^{(t)}; \theta^{(t)}, d^\star)}. \end{aligned}$$

where we have used the definition of $L_A(w_\star^{(t)}; \theta^{(t)}, d^\star)$ in the last step.

For the second term, a similar argument leads to:

$$\mathbb{E}_{s \sim d_\star^\star, a \sim \pi^\star(\cdot|s)} \left[(w_\star^{(t)} - w^{(t)}) \cdot \nabla_\theta \log \pi^{(t)}(a|s) \right] = \sqrt{\|w_\star^{(t)} - w^{(t)}\|_{\Sigma_{d^\star}^{(t)}}^2}.$$

Define $\kappa^{(t)} := \|(\Sigma_\nu^{(t)})^{-1/2} \Sigma_{d^\star} (\Sigma_\nu^{(t)})^{-1/2}\|_2$, which is the relative condition number at iteration t . We have

$$\begin{aligned} \|w_\star^{(t)} - w^{(t)}\|_{\Sigma_{d^\star}^{(t)}}^2 &\leq \|(\Sigma_\nu^{(t)})^{-1/2} \Sigma_{d^\star} (\Sigma_\nu^{(t)})^{-1/2}\|_2 \|w_\star^{(t)} - w^{(t)}\|_{\Sigma_\nu}^2 \\ &\leq \frac{\kappa^{(t)}}{1-\gamma} \|w_\star^{(t)} - w^{(t)}\|_{\Sigma_{d^{(t)}}}^2 \\ &\leq \frac{\kappa^{(t)}}{1-\gamma} (L_A(w^{(t)}; \theta^{(t)}, d^{(t)}) - L_A(w_\star^{(t)}; \theta^{(t)}, d^{(t)})) \end{aligned}$$

where the last step uses that $w_\star^{(t)}$ is a minimizer of L_A over \mathcal{W} and that $w^{(t)}$ is feasible as before (see the proof of Theorem 6.1). Now taking an expectation we have:

$$\begin{aligned} \mathbb{E} [\|w_\star^{(t)} - w^{(t)}\|_{\Sigma_{d^\star}^{(t)}}^2] &\leq \mathbb{E} \left[\frac{\kappa^{(t)}}{1-\gamma} (L_A(w^{(t)}; \theta^{(t)}, d^{(t)}) - L_A(w_\star^{(t)}; \theta^{(t)}, d^{(t)})) \right] \\ &= \mathbb{E} \left[\frac{\kappa^{(t)}}{1-\gamma} \mathbb{E} [L_A(w^{(t)}; \theta^{(t)}, d^{(t)}) - L_A(w_\star^{(t)}; \theta^{(t)}, d^{(t)}) \mid \theta^{(t)}] \right] \\ &\leq \mathbb{E} \left[\frac{\kappa^{(t)}}{1-\gamma} \right] \cdot \epsilon_{\text{stat}} \leq \frac{\kappa \epsilon_{\text{stat}}}{1-\gamma} \end{aligned}$$

where we have used our assumption on κ and ϵ_{stat} .

The proof is completed by substitution and using the concavity of the square root function. ■

7 Discussion

This work provides a systematic study of the convergence properties of policy optimization techniques, both in the tabular and the function approximation settings. At the core, our results imply that the non-convexity of the policy optimization problem is not the fundamental challenge for typical variants of the policy gradient approach. This is evidenced by the global convergence results which we establish and that demonstrate the relative niceness of the underlying optimization problem. At the same time, our results highlight that insufficient exploration can lead to the convergence to sub-optimal policies, as is also observed in practice; technically, we show how this is an issue of conditioning. Conversely, we can expect typical policy gradient algorithms to find the best policy from amongst those whose state-visitation distribution is adequately aligned with the policies we discover, provided a distribution-shifted notion of approximation error is small.

In the tabular case, our results show that the nature and severity of the exploration/distribution mismatch term differs in different policy optimization approaches. For instance, we find that doing policy gradient in its standard form for both the direct and softmax parameterizations can be

slow to converge, particularly in the face of distribution mismatch, even when policy gradients are computed exactly. Natural policy gradient, on the other hand, enjoys a fast dimension-free convergence when we are in tabular settings with exact gradients. On the other hand, for the function approximation setting, or when using finite samples, all algorithms suffer to some degree from the exploration issue captured through a conditioning effect.

With regards to function approximation, the guarantees herein are the first provable results that permit average case approximation errors, where the guarantees do not have explicit worst case dependencies over the state space. These worst case dependencies are avoided by precisely characterizing an approximation/estimation error decomposition, where the relevant approximation error is under distribution shift to an optimal policies measure. Here, we see that successful function approximation relies on two key aspects: good conditioning (related to exploration) and low distribution-shifted, approximation error. In particular, these results identify the relevant measure of the expressivity of a policy class, for the natural policy gradient.

With regards to sample size issues, we showed that simply using stochastic (projected) gradient ascent suffices for accurate policy optimization. However, in terms of improving sample efficiency and polynomial dependencies, there are number of important questions for future research, including variance reduction techniques along with data re-use.

There are number of compelling directions for further study. The first is in understanding how to remove the density ratio guarantees among prior algorithms; our results are suggestive that the incremental policy optimization approaches, including CPI [Kakade and Langford, 2002], PSDP [Bagnell et al., 2004], and MD-MPI Geist et al. [2019], may permit such an improved analysis. The question of understanding what representations are robust to distribution shift is well-motivated by the nature of our distribution-shifted, approximation error (the transfer error). Finally, we hope that policy optimization approaches can be combined with exploration approaches, so that, provably, these approaches can retain their robustness properties (in terms of their agnostic learning guarantees) while mitigating the need for a well conditioned initial starting distribution.

Acknowledgments

We thank the anonymous reviewers who provided detailed and constructive feedback that helped us significantly improve the presentation and exposition. Sham Kakade and Alekh Agarwal gratefully acknowledge numerous helpful discussions with Wen Sun with regards to the Q -NPG algorithm and our notion of transfer error. We also acknowledge numerous helpful comments from Ching-An Cheng and Andrea Zanette on an earlier draft of this work. We thank Nan Jiang, Bruno Scherrer, and Matthieu Geist for their comments with regards to the relationship between concentrability coefficients, the condition number, and the transfer error; this discussion ultimately lead to Corollary 6.1. Sham Kakade acknowledges funding from the Washington Research Foundation for Innovation in Data-intensive Discovery, the ONR award N00014-18-1-2247, and the DARPA award FA8650-18-2-7836. Jason D. Lee acknowledges support of the ARO under MURI Award W911NF-11-1-0303. This is part of the collaboration between US DOD, UK MOD and UK Engineering and Physical Research Council (EPSRC) under the Multidisciplinary University Research Initiative.

References

- Yasin Abbasi-Yadkori, Peter Bartlett, Kush Bhatia, Nevena Lazic, Csaba Szepesvari, and Gellért Weisz. POLITEX: Regret bounds for policy iteration using expert prediction. In *International Conference on Machine Learning*, pages 3692–3702, 2019a.
- Yasin Abbasi-Yadkori, Nevena Lazic, Csaba Szepesvari, and Gellert Weisz. Exploration-enhanced politex. *arXiv preprint arXiv:1908.10479*, 2019b.
- Abbas Abdolmaleki, Jost Tobias Springenberg, Yuval Tassa, Remi Munos, Nicolas Heess, and Martin Riedmiller. Maximum a posteriori policy optimisation. In *International Conference on Learning Representations*, 2018.
- Zafarali Ahmed, Nicolas Le Roux, Mohammad Norouzi, and Dale Schuurmans, editors. *Understanding the impact of entropy on policy optimization*, 2019. URL <https://arxiv.org/abs/1811.11214>.
- András Antos, Csaba Szepesvári, and Rémi Munos. Learning near-optimal policies with bellman-residual minimization based fitted policy iteration and a single sample path. *Machine Learning*, 71(1):89–129, 2008.
- Hédy Attouch, Jérôme Bolte, Patrick Redont, and Antoine Soubeyran. Proximal alternating minimization and projection methods for nonconvex problems: An approach based on the kurdyka-łojasiewicz inequality. *Mathematics of Operations Research*, 35(2):438–457, 2010.
- Mohammad Gheshlaghi Azar, Vicenç Gómez, and Hilbert J. Kappen. Dynamic policy programming. *J. Mach. Learn. Res.*, 13(1), November 2012. ISSN 1532-4435.
- J. A. Bagnell, Sham M Kakade, Jeff G. Schneider, and Andrew Y. Ng. Policy search by dynamic programming. In S. Thrun, L. K. Saul, and B. Schölkopf, editors, *Advances in Neural Information Processing Systems 16*, pages 831–838. MIT Press, 2004.
- J. Andrew Bagnell and Jeff Schneider. Covariant policy search. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence, IJCAI’03*, pages 1019–1024, San Francisco, CA, USA, 2003. Morgan Kaufmann Publishers Inc.
- Keith Ball. An elementary introduction to modern convex geometry. *Flavors of geometry*, 31: 1–58, 1997.
- A. Beck. *First-Order Methods in Optimization*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2017. doi: 10.1137/1.9781611974997.
- Richard Bellman and Stuart Dreyfus. Functional approximations and dynamic programming. *Mathematical Tables and Other Aids to Computation*, 13(68):247–251, 1959.
- Dimitri P Bertsekas and John N Tsitsiklis. *Neuro-Dynamic Programming*. Athena Scientific, Belmont, MA, 1996.

- Jalaj Bhandari and Daniel Russo. Global optimality guarantees for policy gradient methods. *CoRR*, abs/1906.01786, 2019. URL <http://arxiv.org/abs/1906.01786>.
- Shalabh Bhatnagar, Richard S Sutton, Mohammad Ghavamzadeh, and Mark Lee. Natural actor-critic algorithms. *Automatica*, 45(11):2471–2482, 2009.
- Jérôme Bolte, Aris Daniilidis, and Adrian Lewis. The Łojasiewicz inequality for nonsmooth sub-analytic functions with applications to subgradient dynamical systems. *SIAM Journal on Optimization*, 17(4):1205–1223, 2007.
- Justin A Boyan. Least-squares temporal difference learning. In *Proceedings of the Sixteenth International Conference on Machine Learning*, pages 49–56. Morgan Kaufmann Publishers Inc., 1999.
- Ronen I Brafman and Moshe Tennenholtz. R-max-a general polynomial time algorithm for near-optimal reinforcement learning. *The Journal of Machine Learning Research*, 3:213–231, 2003.
- Sébastien Bubeck, Nicolò Cesa-Bianchi, and Sham M. Kakade. Towards minimax policies for online linear optimization with bandit feedback. In *COLT 2012 - The 25th Annual Conference on Learning Theory, June 25-27, 2012, Edinburgh, Scotland*, volume 23 of *JMLR Proceedings*, pages 41.1–41.14, 2012.
- Qi Cai, Zhuoran Yang, Jason D. Lee, and Zhaoran Wang. Neural temporal-difference learning converges to global optima. *CoRR*, abs/1905.10027, 2019. URL <http://arxiv.org/abs/1905.10027>.
- Nicolo Cesa-Bianchi and Gabor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, New York, NY, USA, 2006. ISBN 0521841089.
- Jinglin Chen and Nan Jiang. Information-theoretic considerations in batch reinforcement learning. In *International Conference on Machine Learning*, pages 1042–1051, 2019.
- Eyal Even-Dar, Sham M Kakade, and Yishay Mansour. Online Markov decision processes. *Mathematics of Operations Research*, 34(3):726–736, 2009.
- Amir-massoud Farahmand, Csaba Szepesvári, and Rémi Munos. Error propagation for approximate policy and value iteration. In *Advances in Neural Information Processing Systems*, pages 568–576, 2010.
- Maryam Fazel, Rong Ge, Sham M Kakade, and Mehran Mesbahi. Global convergence of policy gradient methods for the linear quadratic regulator. *arXiv preprint arXiv:1801.05039*, 2018.
- Yoav Freund and Robert E Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences*, 55(1):119–139, 1997.
- Rong Ge, Furong Huang, Chi Jin, and Yang Yuan. Escaping from saddle points - online stochastic gradient for tensor decomposition. *Proceedings of The 28th Conference on Learning Theory, COLT 2015, Paris, France, July 3-6, 2015*, 2015.

- Matthieu Geist, Bruno Scherrer, and Olivier Pietquin. A theory of regularized markov decision processes. *arXiv preprint arXiv:1901.11275*, 2019.
- Saeed Ghadimi and Guanhui Lan. Accelerated gradient methods for nonconvex nonlinear and stochastic programming. *Mathematical Programming*, 156(1-2):59–99, 2016.
- Daniel Hsu, Sham M. Kakade, and Tong Zhang. Random design analysis of ridge regression. *Foundations of Computational Mathematics*, 14(3):569–600, 2014.
- Nan Jiang, Akshay Krishnamurthy, Alekh Agarwal, John Langford, and Robert E Schapire. Contextual decision processes with low bellman rank are PAC-learnable. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1704–1713. JMLR. org, 2017.
- Chi Jin, Rong Ge, Praneeth Netrapalli, Sham M. Kakade, and Michael I. Jordan. How to escape saddle points efficiently. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, pages 1724–1732, 2017.
- Chi Jin, Zhuoran Yang, Zhaoran Wang, and Michael I Jordan. Provably efficient reinforcement learning with linear function approximation. *arXiv preprint arXiv:1907.05388*, 2019.
- Fritz John. Extremum problems with inequalities as subsidiary conditions. Interscience Publishers, 1948.
- S. Kakade. A natural policy gradient. In *NIPS*, 2001.
- Sham Kakade and John Langford. Approximately Optimal Approximate Reinforcement Learning. In *Proceedings of the 19th International Conference on Machine Learning*, volume 2, pages 267–274, 2002.
- Hamed Karimi, Julie Nutini, and Mark Schmidt. Linear convergence of gradient and proximal-gradient methods under the polyak-łojasiewicz condition. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 795–811. Springer, 2016.
- Michael Kearns and Satinder Singh. Near-optimal reinforcement learning in polynomial time. *Machine Learning*, 49(2-3):209–232, 2002.
- Vijay R Konda and John N Tsitsiklis. Actor-critic algorithms. In *Advances in neural information processing systems*, pages 1008–1014, 2000.
- Alessandro Lazaric, Mohammad Ghavamzadeh, and Rémi Munos. Analysis of classification-based policy iteration algorithms. *The Journal of Machine Learning Research*, 17(1):583–612, 2016.
- Boyi Liu, Qi Cai, Zhuoran Yang, and Zhaoran Wang. Neural proximal/trust region policy optimization attains globally optimal policy. *CoRR*, abs/1906.10306, 2019. URL <http://arxiv.org/abs/1906.10306>.

- Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pages 1928–1937, 2016.
- Rémi Munos. Error bounds for approximate policy iteration. In *ICML*, volume 3, pages 560–567, 2003.
- Rémi Munos. Error bounds for approximate value iteration. In *AAAI*, 2005.
- Arkadii Semenovich Nemirovsky and David Borisovich Yudin. Problem complexity and method efficiency in optimization. 1983.
- Yurii Nesterov and Boris T. Polyak. Cubic regularization of newton method and its global performance. *Math. Program.*, pages 177–205, 2006.
- Gergely Neu, Andras Antos, András György, and Csaba Szepesvári. Online markov decision processes under bandit feedback. In *Advances in Neural Information Processing Systems 23*. Curran Associates, Inc., 2010.
- Gergely Neu, Anders Jonsson, and Vicenç Gómez. A unified view of entropy-regularized markov decision processes. *CoRR*, abs/1705.07798, 2017.
- Jan Peters and Stefan Schaal. Natural actor-critic. *Neurocomput.*, 71(7-9):1180–1190, 2008. ISSN 0925-2312.
- Jan Peters, Katharina Mülling, and Yasemin Altün. Relative entropy policy search. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI 2010)*, pages 1607–1612. AAAI Press, 2010.
- B. T. Polyak. Gradient methods for minimizing functionals. *USSR Computational Mathematics and Mathematical Physics*, 3(4):864–878, 1963.
- Aravind Rajeswaran, Kendall Lowrey, Emanuel V. Todorov, and Sham M Kakade. Towards generalization and simplicity in continuous control. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 6550–6561. Curran Associates, Inc., 2017.
- Bruno Scherrer. Approximate policy iteration schemes: A comparison. In *Proceedings of the 31st International Conference on International Conference on Machine Learning - Volume 32, ICML'14*. JMLR.org, 2014.
- Bruno Scherrer and Matthieu Geist. Local policy search in a convex space and conservative policy iteration as boosted policy search. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 35–50. Springer, 2014.
- Bruno Scherrer, Mohammad Ghavamzadeh, Victor Gabillon, Boris Lesner, and Matthieu Geist. Approximate modified policy iteration and its application to the game of tetris. *Journal of Machine Learning Research*, 16:1629–1676, 2015.

- John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International Conference on Machine Learning*, pages 1889–1897, 2015.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- S. Shalev-Shwartz and S. Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Understanding Machine Learning: From Theory to Algorithms. Cambridge University Press, 2014. ISBN 9781107057135. URL <https://books.google.com/books?id=ttJkAwAAQBAJ>.
- Shai Shalev-Shwartz et al. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2012.
- Lior Shani, Yonathan Efroni, and Shie Mannor. Adaptive trust region policy optimization: Global convergence and faster rates for regularized mdps, 2019.
- Richard S Sutton, David A McAllester, Satinder P Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. In *Advances in Neural Information Processing Systems*, volume 99, pages 1057–1063, 1999.
- Csaba Szepesvári and Rémi Munos. Finite time bounds for sampling based fitted value iteration. In *Proceedings of the 22nd international conference on Machine learning*, pages 880–887. ACM, 2005.
- Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.
- Ronald J Williams and Jing Peng. Function optimization using connectionist reinforcement learning algorithms. *Connection Science*, 3(3):241–268, 1991.
- Lin F. Yang and Mengdi Wang. Sample-optimal parametric q-learning using linearly additive features. In *International Conference on Machine Learning*, pages 6995–7004, 2019.

A Proofs for Section 3

Proof:[of Lemma 3.1] Recall the MDP in Figure 1. Note that since actions in terminal states s_3 , s_4 and s_5 do not change the expected reward, we only consider actions in states s_1 and s_2 . Let the "up/above" action as a_1 and "right" action as a_2 . Note that

$$V^\pi(s_1) = \pi(a_2|s_1)\pi(a_1|s_2) \cdot r$$

Consider

$$\theta^{(1)} = (\log 1, \log 3, \log 3, \log 1), \quad \theta^{(2)} = (-\log 1, -\log 3, -\log 3, -\log 1)$$

where θ is written as a tuple $(\theta_{a_1, s_1}, \theta_{a_2, s_1}, \theta_{a_1, s_2}, \theta_{a_2, s_2})$. Then, for the softmax parameterization, we have

$$\pi^{(1)}(a_2|s_1) = \frac{3}{4}; \quad \pi^{(1)}(a_1|s_2) = \frac{3}{4}; \quad V^{(1)}(s_1) = \frac{9}{16}r$$

and

$$\pi^{(2)}(a_2|s_1) = \frac{1}{4}; \quad \pi^{(2)}(a_1|s_2) = \frac{1}{4}; \quad V^{(2)}(s_1) = \frac{1}{16}r$$

Also, for $\theta^{(\text{mid})} = \frac{\theta^{(1)} + \theta^{(2)}}{2}$,

$$\pi^{(\text{mid})}(a_2|s_1) = \frac{1}{2}; \quad \pi^{(\text{mid})}(a_1|s_2) = \frac{1}{2}; \quad V^{(\text{mid})}(s_1) = \frac{1}{4}r$$

This gives

$$V^{(1)}(s_1) + V^{(2)}(s_1) > 2V^{(\text{mid})}(s_1)$$

which shows that V^π is non-concave. ■

Proof:[of Lemma 3.2] Let $\text{Pr}^\pi(\tau|s_0 = s)$ denote the probability of observing a trajectory τ when starting in state s and following the policy π . Using a telescoping argument, we have:

$$\begin{aligned} V^\pi(s) - V^{\pi'}(s) &= \mathbb{E}_{\tau \sim \text{Pr}^\pi(\tau|s_0=s)} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \right] - V^{\pi'}(s) \\ &= \mathbb{E}_{\tau \sim \text{Pr}^\pi(\tau|s_0=s)} \left[\sum_{t=0}^{\infty} \gamma^t \left(r(s_t, a_t) + V^{\pi'}(s_t) - V^{\pi'}(s_t) \right) \right] - V^{\pi'}(s) \\ &\stackrel{(a)}{=} \mathbb{E}_{\tau \sim \text{Pr}^\pi(\tau|s_0=s)} \left[\sum_{t=0}^{\infty} \gamma^t \left(r(s_t, a_t) + \gamma V^{\pi'}(s_{t+1}) - V^{\pi'}(s_t) \right) \right] \\ &\stackrel{(b)}{=} \mathbb{E}_{\tau \sim \text{Pr}^\pi(\tau|s_0=s)} \left[\sum_{t=0}^{\infty} \gamma^t \left(r(s_t, a_t) + \gamma \mathbb{E}[V^{\pi'}(s_{t+1})|s_t, a_t] - V^{\pi'}(s_t) \right) \right] \\ &= \mathbb{E}_{\tau \sim \text{Pr}^\pi(\tau|s_0=s)} \left[\sum_{t=0}^{\infty} \gamma^t A^{\pi'}(s_t, a_t) \right] = \frac{1}{1-\gamma} \mathbb{E}_{s' \sim d_s^\pi} \mathbb{E}_{a \sim \pi(\cdot|s)} \gamma^t A^{\pi'}(s', a), \end{aligned}$$

where (a) rearranges terms in the summation and cancels the $V^{\pi'}(s_0)$ term with the $-V^{\pi'}(s)$ outside the summation, and (b) uses the tower property of conditional expectations and the final equality follows from the definition of d_s^π . ■

B Proofs for Section 4

B.1 Proofs for Section 4.2

We first define first-order optimality for constrained optimization.

Definition B.1 (First-order Stationarity). A policy $\pi_\theta \in \Delta(\mathcal{A})^{|\mathcal{S}|}$ is ϵ -stationary with respect to the initial state distribution μ if

$$G(\pi_\theta) := \max_{\pi_\theta + \delta \in \Delta(\mathcal{A})^{|\mathcal{S}|}, \|\delta\|_2 \leq 1} \delta^\top \nabla_\pi V^{\pi_\theta}(\mu) \leq \epsilon.$$

where $\Delta(\mathcal{A})^{|\mathcal{S}|}$ is the set of all policies.

Due to that we are working with the direct parameterization (see (2)), we drop the θ subscript.

Remark B.1. If $\epsilon = 0$, then the definition simplifies to $\delta^\top \nabla_\pi V^\pi(\mu) \leq 0$. Geometrically, δ is a feasible direction of movement since the probability simplex $\Delta(\mathcal{A})^{|\mathcal{S}|}$ is convex. Thus the gradient is negatively correlated with any feasible direction of movement, and so π is first-order stationary.

Proposition B.1. Let $V^\pi(\mu)$ be β -smooth in π . Define the gradient mapping

$$G^\eta(\pi) = \frac{1}{\eta} \left(P_{\Delta(\mathcal{A})^{|\mathcal{S}|}}(\pi + \eta \nabla_\pi V^\pi(\mu)) - \pi \right),$$

and the update rule for the projected gradient is $\pi^+ = \pi + \eta G^\eta(\pi)$. If $\|G^\eta(\pi)\|_2 \leq \epsilon$, then

$$\max_{\pi + \delta \in \Delta(\mathcal{A})^{|\mathcal{S}|}, \|\delta\|_2 \leq 1} \delta^\top \nabla_\pi V^{\pi^+}(\mu) \leq \epsilon(\eta\beta + 1).$$

Proof: By Theorem E.2,

$$\nabla_\pi V^{\pi^+}(\mu) \in N_{\Delta(\mathcal{A})^{|\mathcal{S}|}}(\pi^+) + \epsilon(\eta\beta + 1)B_2,$$

where B_2 is the unit ℓ_2 ball, and $N_{\Delta(\mathcal{A})^{|\mathcal{S}|}}$ is the normal cone of the product simplex $\Delta(\mathcal{A})^{|\mathcal{S}|}$. Since $\nabla_\pi V^{\pi^+}(\mu)$ is $\epsilon(\eta\beta + 1)$ distance from the normal cone and δ is in the tangent cone, then $\delta^\top \nabla_\pi V^{\pi^+}(\mu) \leq \epsilon(\eta\beta + 1)$. ■

Proof:[of Theorem 4.1] Recall the definition of gradient mapping

$$G^\eta(\pi) = \frac{1}{\eta} \left(P_{\Delta(\mathcal{A})^{|\mathcal{S}|}}(\pi + \eta \nabla_\pi V^{(\pi)}(\mu)) - \pi \right)$$

From Lemma D.3, we have $V^\pi(s)$ is β -smooth for all states s (and also hence $V^\pi(\mu)$ is also β -smooth) with $\beta = \frac{2\gamma|\mathcal{A}|}{(1-\gamma)^3}$. Then, from standard result (Theorem E.1), we have that for $G^\eta(\pi)$ with step-size $\eta = \frac{1}{\beta}$,

$$\min_{t=0,1,\dots,T-1} \|G^\eta(\pi^{(t)})\|_2 \leq \frac{\sqrt{2\beta(V^*(\mu) - V^{(0)}(\mu))}}{\sqrt{T}}$$

Then, from Proposition B.1, we have

$$\min_{t=0,1,\dots,T} \max_{\pi^{(t)} + \delta \in \Delta(A)^{|\mathcal{S}|}, \|\delta\|_2 \leq 1} \delta^\top \nabla_\pi V^{\pi^{(t+1)}}(\mu) \leq (\eta\beta + 1) \frac{\sqrt{2\beta(V^*(\mu) - V^{(0)}(\mu))}}{\sqrt{T}}$$

Observe that

$$\begin{aligned} \max_{\bar{\pi} \in \Delta(A)^{|\mathcal{S}|}} (\bar{\pi} - \pi)^\top \nabla_\pi V^\pi(\mu) &= 2\sqrt{|\mathcal{S}|} \max_{\bar{\pi} \in \Delta(A)^{|\mathcal{S}|}} \frac{1}{2\sqrt{|\mathcal{S}|}} (\bar{\pi} - \pi)^\top \nabla_\pi V^\pi(\mu) \\ &\leq 2\sqrt{|\mathcal{S}|} \max_{\pi + \delta \in \Delta(A)^{|\mathcal{S}|}, \|\delta\|_2 \leq 1} \delta^\top \nabla_\pi V^\pi(\mu) \end{aligned}$$

where the last step follows as $\|\bar{\pi} - \pi\|_2 \leq 2\sqrt{|\mathcal{S}|}$. And then using Lemma 4.1 and $\eta\beta = 1$, we have

$$\min_{t=0,1,\dots,T} V^*(\rho) - V^{(t)}(\rho) \leq \frac{4\sqrt{|\mathcal{S}|}}{1-\gamma} \left\| \frac{d_\rho^{\pi^*}}{\mu} \right\|_\infty \frac{\sqrt{2\beta(V^*(\mu) - V^{(0)}(\mu))}}{\sqrt{T}}$$

We can get our required bound of ϵ , if we set T such that

$$\frac{4\sqrt{|\mathcal{S}|}}{1-\gamma} \left\| \frac{d_\rho^{\pi^*}}{\mu} \right\|_\infty \frac{\sqrt{2\beta(V^*(\mu) - V^{(0)}(\mu))}}{\sqrt{T}} \leq \epsilon$$

or, equivalently,

$$T \geq \frac{32|\mathcal{S}|\beta(V^*(\mu) - V^{(0)}(\mu))}{(1-\gamma)^2\epsilon^2} \left\| \frac{d_\rho^{\pi^*}}{\mu} \right\|_\infty^2.$$

Using $V^*(\mu) - V^{(0)}(\mu) \leq \frac{1}{1-\gamma}$ and $\beta = \frac{2\gamma|\mathcal{A}|}{(1-\gamma)^3}$ from Lemma D.3 leads to the desired result. \blacksquare

B.2 Proofs for Section 4.3

Recall the MDP in Figure 2. Each trajectory starts from the initial state s_0 , and we use the discount factor $\gamma = H/(H+1)$. Recall that we work with the direct parameterization, where $\pi_\theta(a|s) = \theta_{s,a}$ for $a = a_1, a_2, a_3$ and $\pi_\theta(a_4|s) = 1 - \theta_{s,a_1} - \theta_{s,a_2} - \theta_{s,a_3}$. Note that since states s_0 and s_{H+1} only have once action, therefore, we only consider the parameters for states s_1 to s_H . For this policy class and MDP, let P^θ be the state transition matrix under π_θ , i.e. $[P^\theta]_{s,s'}$ is the probability of going from state s to s' under policy π_θ :

$$[P^\theta]_{s,s'} = \sum_{a \in \mathcal{A}} \pi_\theta(a|s) P(s'|s, a).$$

For the MDP illustrated in Figure 2, the entries of this matrix are given as:

$$[P^\theta]_{s,s'} = \begin{cases} \theta_{s,a_1} & \text{if } s' = s_{i+1} \text{ and } s = s_i \text{ with } 1 \leq i \leq H \\ 1 - \theta_{s,a_1} & \text{if } s' = s_{i-1} \text{ and } s = s_i \text{ with } 1 \leq i \leq H \\ 1 & \text{if } s' = s_1 \text{ and } s = s_0 \\ 1 & \text{if } s' = s = s_{H+1} \\ 0 & \text{otherwise} \end{cases}. \quad (27)$$

With this definition, we recall that the value function in the initial state s_0 is given by

$$V^{\pi_\theta}(s_0) = \mathbb{E}_{\tau \sim \pi_\theta} \left[\sum_{t=0}^{\infty} \gamma^t r_t \right] = e_0^T (I - \gamma P^\theta)^{-1} r,$$

where e_0 is an indicator vector for the starting state s_0 . From the form of the transition probabilities (27), it is clear that the value function only depends on the parameters θ_{s,a_1} in any state s . While care is needed for derivatives as the parameters across actions are related by the simplex feasibility constraints, we have assumed each parameter is strictly positive, so that an infinitesimal change to any parameter other than θ_{s,a_1} does not affect the policy value and hence the policy gradients. With this understanding, we succinctly refer to θ_{s,a_1} as θ_s in any state s . We also refer to the state s_i simply as i to reduce subscripts.

For convenience, we also define \bar{p} (resp. \underline{p}) to be the largest (resp. smallest) of the probabilities θ_s across the states $s \in [1, H]$ in the MDP.

In this section, we prove Proposition 4.1, that is: for $0 < \theta < 1$ (componentwise across states and actions), $\bar{p} \leq 1/4$, and for all $k \leq \frac{H}{40 \log(2H)} - 1$, we have $\|\nabla_\theta^k V^{\pi_\theta}(s_0)\| \leq (1/3)^{H/4}$, where $\nabla_\theta^k V^{\pi_\theta}(s_0)$ is a tensor of the k_{th} order. Furthermore, we seek to show $V^*(s_0) - V^{\pi_\theta}(s_0) \geq (H+1)/8 - (H+1)^2/3^H$ (where θ^* are the optimal policy's parameters).

It is easily checked that $V^{\pi_\theta}(s_0) = M_{0,H+1}^\theta$, where

$$M^\theta := (I - \gamma P^\theta)^{-1},$$

since the only rewards are obtained in the state s_{H+1} . In order to bound the derivatives of the expected reward, we first establish some properties of the matrix M^θ .

Lemma B.1. *Suppose $\bar{p} \leq 1/4$. Fix any $\alpha \in \left[\frac{1 - \sqrt{1 - 4\gamma^2 \bar{p}(1-\underline{p})}}{2\gamma(1-\underline{p})}, \max \left\{ \frac{1 + \sqrt{1 - 4\gamma^2 \bar{p}(1-\underline{p})}}{2\gamma(1-\underline{p})}, 1 \right\} \right]$. Then*

1. $M_{a,b}^\theta \leq \frac{\alpha^{b-a-1}}{1-\gamma}$ for $0 \leq a \leq b \leq H$.
2. $M_{a,H+1}^\theta \leq \frac{\gamma \bar{p}}{1-\gamma} M_{a,H}^\theta \leq \frac{\gamma \bar{p}}{(1-\gamma)^2} \alpha^{H-a}$ for $0 \leq a \leq H$.

Proof: Let $\rho_{a,b}^k$ be the normalized discounted probability of reaching b , when the initial state is a , in k steps, that is

$$\rho_{a,b}^k := (1 - \gamma) \sum_{i=0}^k [(\gamma P^\theta)^i]_{a,b}, \quad (28)$$

where we recall the convention that U^0 is the identity matrix for any square matrix U . Observe that $0 \leq \rho_{a,b}^k \leq 1$, and, based on the form (27) of P^θ , we have the recursive relation for all $k > 0$:

$$\rho_{a,b}^k = \begin{cases} \gamma(1 - \theta_{b+1})\rho_{a,b+1}^{k-1} + \gamma\theta_{b-1}\rho_{a,b-1}^{k-1} & \text{if } 1 < b < H \\ \gamma\theta_{H-1}\rho_{a,H-1}^{k-1} & \text{if } b = H \\ \gamma\theta_H\rho_{a,H}^{k-1} + \gamma\rho_{a,H+1}^{k-1} & \text{if } b = H + 1 \text{ and } a < H + 1 \\ 1 - \gamma & \text{if } b = H + 1 \text{ and } a = H + 1 \\ \gamma(1 - \theta_2)\rho_{a,2}^{k-1} + \gamma\rho_{a,0}^{k-1} & \text{if } b = 0 \end{cases} \quad (29)$$

Note that $\rho_{a,b}^0 = 0$ for $a \neq b$ and $\rho_{a,b}^0 = 1 - \gamma$ for $a = b$. Now let us inductively prove that for all $k \geq 0$

$$\rho_{a,b}^k \leq \alpha^{b-a} \quad \text{for } 1 \leq a \leq b \leq H. \quad (30)$$

Clearly this holds for $k = 0$ since $\rho_{a,b}^0 = 0$ for $a \neq b$ and $\rho_{a,b}^0 = 1 - \gamma$ for $a = b$. Now, assuming the bound for all steps till $k - 1$, we now prove it for k case by case.

For $a = b$ the result follows since

$$\rho_{a,b}^k \leq 1 = \alpha^{b-a}.$$

For $1 < b < H$ and $a < b$, observe that the recursion (29) and the inductive hypothesis imply that

$$\begin{aligned} \rho_{a,b}^k &\leq \gamma(1 - \theta_{b+1})\alpha^{b+1-a} + \gamma\theta_{b-1}\alpha^{b-1-a} \\ &= \alpha^{b-a-1}(\gamma(1 - \theta_{b+1})\alpha^2 + \gamma\theta_{b-1}) \\ &\leq \alpha^{b-a-1}(\gamma(1 - \underline{p})\alpha^2 + \gamma\bar{p}) \\ &= \alpha^{b-a-1}(\alpha + \gamma(1 - \underline{p})\alpha^2 - \alpha + \gamma\bar{p}) \leq \alpha^{b-a}, \end{aligned}$$

where the last inequality follows since $\alpha^2\gamma(1 - \underline{p}) - \alpha + \gamma\bar{p} \leq 0$ due to that α is within the roots of this quadratic equation. Note the discriminant term in the square root is non-negative provided $\bar{p} < 1/4$, since the condition along with the knowledge that $\underline{p} \leq \bar{p}$ ensures that $4\gamma^2\bar{p}(1 - \underline{p}) \leq 1$.

For $b = H$ and $a < H$, we observe that

$$\begin{aligned} \rho_{a,b}^k &\leq \gamma\theta_{H-1}\alpha^{H-1-a} \\ &= \alpha^{H-a}\frac{\gamma\theta_{H-1}}{\alpha} \\ &\leq \alpha^{H-a}\left(\frac{\gamma\bar{p}}{\alpha}\right) \leq \alpha^{H-a}(\gamma(1 - \underline{p})\alpha + \frac{\gamma\bar{p}}{\alpha}) \leq \alpha^{H-a}. \end{aligned}$$

This proves the inductive claim (note that the cases of $b = a = 1$ and $b = a = H$ are already handled in the first part above). Next, we prove that for all $k \geq 0$

$$\rho_{0,b}^k \leq \alpha^{b-1}.$$

Clearly this holds for $k = 0$ and $b \neq 0$ since $\rho_{0,b}^0 = 0$. Furthermore, for all $k \geq 0$ and $b = 0$,

$$\rho_{0,b}^k \leq 1 \leq \alpha^{b-1},$$

since $\alpha \leq 1$ by construction and $b = 0$. Now, we consider the only remaining case when $k > 0$ and $b \in [1, H + 1]$. By (27), observe that for $k > 0$ and $b \in [1, H + 1]$,

$$[(P^\theta)^i]_{0,b} = [(P^\theta)^{i-1}]_{1,b}, \quad (31)$$

for all $i \geq 1$. Using the definition of $\rho_{a,b}^k$ (28) for $k > 0$ and $b \in [1, H + 1]$,

$$\begin{aligned}
\rho_{0,b}^k &= (1 - \gamma) \sum_{i=0}^k [(\gamma P^\theta)^i]_{0,b} = (1 - \gamma)[(\gamma P^\theta)^0]_{0,b} + (1 - \gamma) \sum_{i=1}^k [(\gamma P^\theta)^i]_{0,b} \\
&= 0 + (1 - \gamma) \sum_{i=1}^k \gamma^i [(P^\theta)^i]_{0,b} && \text{(since } b \geq 1) \\
&= (1 - \gamma) \sum_{i=1}^k \gamma^i [(P^\theta)^{i-1}]_{1,b} && \text{(using Equation (31))} \\
&= (1 - \gamma) \gamma \sum_{j=0}^{k-1} \gamma^j [(P^\theta)^j]_{1,b} && \text{(By substituting } j = i - 1) \\
&= \gamma \rho_{1,b}^{k-1} && \text{(using Equation (28))} \\
&\leq \alpha^{b-1} && \text{(using Equation (30) and } \gamma, \alpha \leq 1)
\end{aligned}$$

Hence, for all $k \geq 0$

$$\rho_{0,b}^k \leq \alpha^{b-1}$$

In conjunction with Equation (30), the above display gives for all $k \geq 0$,

$$\begin{aligned}
\rho_{a,b}^k &\leq \alpha^{b-a} \quad \text{for } 1 \leq a \leq b \leq H \\
\rho_{a,b}^k &\leq \alpha^{b-a-1} \quad \text{for } 0 = a \leq b \leq H
\end{aligned}$$

Also observe that

$$M_{a,b}^\theta = \lim_{k \rightarrow \infty} \frac{\rho_{a,b}^k}{1 - \gamma}.$$

Since the above bound holds for all $k \geq 0$, it also applies to the limiting value $M_{a,b}^\theta$, which shows that

$$\begin{aligned}
M_{a,b}^\theta &\leq \frac{\alpha^{b-a}}{1 - \gamma} \leq \frac{\alpha^{b-a-1}}{1 - \gamma} \quad \text{for } 1 \leq a \leq b \leq H \\
M_{a,b}^\theta &\leq \frac{\alpha^{b-a-1}}{1 - \gamma} \quad \text{for } 0 = a \leq b \leq H
\end{aligned}$$

which completes the proof of the first part of the lemma.

For the second claim, from recursion (29) and $b = H + 1$ and $a < H + 1$

$$\rho_{a,H+1}^k = \gamma \theta_H \rho_{a,H}^{k-1} + \gamma \rho_{a,H+1}^{k-1} \leq \gamma \bar{p} \rho_{a,H}^{k-1} + \gamma \rho_{a,H+1}^{k-1},$$

Taking the limit of $k \rightarrow \infty$, we see that

$$M_{a,H+1}^\theta \leq \gamma \bar{p} M_{a,H}^\theta + \gamma M_{a,H+1}^\theta.$$

Rearranging the terms in the above bound yields the second claim in the lemma. ■

Using the lemma above, we now bound the derivatives of M^θ .

Lemma B.2. *The k_{th} order partial derivatives of M satisfy:*

$$\left| \frac{\partial^k M_{0,H+1}^\theta}{\partial \theta_{\beta_1} \dots \partial \theta_{\beta_k}} \right| \leq \frac{\bar{p} 2^k \gamma^{k+1} k! \alpha^{H-2k}}{(1-\gamma)^{k+2}}.$$

where β denotes a k dimensional vector with entries in $\{1, 2, \dots, H\}$.

Proof: Since the parameter θ is fixed throughout, we drop the superscript in M^θ for brevity. Using $\nabla_\theta M = -M \nabla_\theta (I - \gamma P^\theta) M$, using the form of P^θ in (27), we get for any $h \in [1, H]$

$$-\frac{\partial M_{a,b}}{\partial \theta_h} = -\gamma \sum_{i,j=0}^{H+1} M_{a,i} \frac{\partial P_{i,j}}{\partial \theta_h} M_{j,b} = \gamma M_{a,h} (M_{h-1,b} - M_{h+1,b}) \quad (32)$$

where the second equality follows since $P_{h,h+1} = \theta_h$ and $P_{h,h-1} = 1 - \theta_h$ are the only two entries in the transition matrix which depend on θ_h for $h \in [1, H]$.

Next, let us consider a k_{th} order partial derivative of $M_{0,H+1}$, denoted as $\frac{\partial^k M_{0,H+1}}{\partial \theta_\beta}$. Note that β can have repeated entries to capture higher order derivative with respect to some parameter. We prove by induction for all $k \geq 1$, $-\frac{\partial^k M_{0,H+1}}{\partial \theta_\beta}$ can be written as $\sum_{n=1}^N c_n \zeta_n$ where

1. $|c_n| = \gamma^k$ and $N \leq 2^k k!$,
2. Each monomial ζ_n is of the form $M_{i_1, j_1} \dots M_{i_{k+1}, j_{k+1}}$, $i_1 = 0$, $j_{k+1} = H + 1$, $j_l \leq H$ and $i_{l+1} = j_l \pm 1$ for all $l \in [1, k]$.

The base case $k = 1$ follows from Equation (32), as we can write for any $h \in [H]$

$$-\frac{\partial M_{0,H+1}}{\partial \theta_h} = \gamma M_{0,h} M_{h-1,H+1} - \gamma M_{0,h} M_{h+1,H+1}$$

Clearly, the induction hypothesis is true with $|c_n| = \gamma$, $N = 2$, $i_1 = 0$, $j_2 = H + 1$, $j_1 \leq H$ and $i_2 = j_1 \pm 1$. Now, suppose the claim holds till $k - 1$. Then by the chain rule:

$$\frac{\partial^k M_{0,H+1}}{\partial \theta_{\beta_1} \dots \partial \theta_{\beta_k}} = \frac{\partial \frac{\partial^{k-1} M_{0,H+1}}{\partial \theta_{\beta'/1}}}{\partial \theta_{\beta_1}},$$

where β'/i is the vector β with the i_{th} entry removed. By inductive hypothesis,

$$-\frac{\partial^{k-1} M_{0,H+1}}{\partial \theta_{\beta'/1}} = \sum_{n=1}^N c_n \zeta_n$$

where

1. $|c_n| = \gamma^{k-1}$ and $N \leq 2^{k-1} (k - 1)!$,

2. Each monomial ζ_n is of the form $M_{i_1, j_1} \dots M_{i_k, j_k}$, $i_1 = 0$, $j_k = H + 1$, $j_l \leq H$ and $i_{l+1} = j_l \pm 1$ for all $l \in [1, k - 1]$.

In order to compute the $(k)_{th}$ derivative of $M_{0, H+1}$, we have to compute derivative of each monomial ζ_n with respect to θ_{β_1} . Consider one of the monomials in the $(k - 1)_{th}$ derivative, say, $\zeta = M_{i_1, j_1} \dots M_{i_k, j_k}$. We invoke the chain rule as before and replace one of the terms in ζ , say M_{i_m, j_m} , with $\gamma M_{i_m, \beta_1} M_{\beta_1 - 1, j_m} - \gamma M_{i_m, \beta_1} M_{\beta_1 + 1, j_m}$ using Equation 32. That is, the derivative of each entry gives rise to two monomials and therefore derivative of ζ leads to $2k$ monomials which can be written in the form $\zeta' = M_{i'_1, j'_1} \dots M_{i'_{k+1}, j'_{k+1}}$ where we have the following properties (by appropriately reordering terms)

1. $i'_l, j'_l = i_l, j_l$ for $l < m$
2. $i'_l, j'_l = i_{l-1}, j_{l-1}$ for $l > m + 1$
3. $i'_m, j'_m = i_m, \beta_1$ and $i'_{m+1}, j'_{m+1} = j_m \pm 1, j_m$

Using the induction hypothesis, we can write

$$-\frac{\partial^k M_{0, H+1}}{\partial \theta_{\beta_1} \dots \partial \theta_{\beta_k}} = \sum_{n=0}^{N'} c'_n \zeta'_n$$

where

1. $|c'_n| = \gamma |c_n| = \gamma^k$, since as shown above each coefficient gets multiplied by $\pm \gamma$.
2. $N' \leq 2k2^{k-1}(k-1)! = 2^k k!$, since as shown above each monomial ζ leads to $2k$ monomials ζ' .
3. Each monomial ζ'_n is of the form $M_{i_1, j_1} \dots M_{i_{k+1}, j_{k+1}}$, $i_1 = 0$, $j_{k+1} = H + 1$, $j_l \leq H$ and $i_{l+1} = j_l \pm 1$ for all $l \in [1, k]$.

This completes the induction.

Next we prove a bound on the magnitude of each of the monomials which arise in the derivatives of $M_{0, H+1}$. Specifically, we show that for each monomial $\zeta = M_{i_1, j_1} \dots M_{i_{k+1}, j_{k+1}}$, we have

$$\left| M_{i_1, j_1} \dots M_{i_{k+1}, j_{k+1}} \right| \leq \frac{\gamma \bar{\rho} \alpha^{H-2k}}{(1-\gamma)^{k+2}} \quad (33)$$

We observe that it suffices to only consider pairs of indices i_l, j_l where $i_l < j_l$. Since $|M_{i, j}| \leq \frac{1}{1-\gamma}$ for all i, j ,

$$\begin{aligned}
\left| \prod_{l=1}^{k+1} M_{i'_l, j'_l} \right| &\leq \left| \prod_{1 \leq l \leq k : i'_l < j'_l} M_{i'_l, j'_l} \right| \left| \prod_{1 \leq l \leq k : i'_l \geq j'_l} \frac{1}{1-\gamma} \right| \left| M_{i'_{k+1}, j'_{k+1}} \right| \\
&= \left| \prod_{1 \leq l \leq k : i'_l < j'_l} M_{i'_l, j'_l} \right| \left| \prod_{1 \leq l \leq k : i'_l \geq j'_l} \frac{1}{1-\gamma} \right| \left| M_{i'_{k+1}, H+1} \right| \\
&\hspace{15em} \text{(by the inductive claim shown above)} \\
&\leq \frac{\alpha^{\sum_{\{1 \leq l \leq k : i'_l < j'_l\}} j'_l - i'_l - 1} \gamma \bar{p} \alpha^{H - i'_{k+1}}}{(1-\gamma)^k (1-\gamma)^2} \\
&\hspace{10em} \text{(using Lemma B.1, parts 1 and 2 on the first and last terms resp.)} \\
&= \frac{\gamma \bar{p} \alpha^{\sum_{\{1 \leq l \leq k+1 : i'_l < j'_l\}} j'_l - i'_l - 1}}{(1-\gamma)^{k+2}} \tag{34}
\end{aligned}$$

The last step follows from $H + 1 = j'_{k+1} \geq i'_{k+1}$. Note that

$$\sum_{\{1 \leq l \leq k+1 : i'_l < j'_l\}} j'_l - i'_l \geq \sum_{l=1}^{k+1} j'_l - i'_l = j'_{k+1} - i'_1 + \sum_{l=1}^k (j'_{l+1} - i'_l) \geq H + 1 - k \geq 0$$

where the first inequality follows from adding only non-positive terms to the sum, the second equality follows from rearranging terms and the third inequality follows from $i'_1 = 0$, $j'_{k+1} = H + 1$ and $i'_{l+1} = j'_l \pm 1$ for all $l \in [1, k]$. Therefore,

$$\sum_{\{1 \leq l \leq k+1 : i'_l < j'_l\}} j'_l - i'_l - 1 \geq H - 2k$$

Using Equation (34) and $\alpha \leq 1$ with above display gives

$$\left| \prod_{l=1}^{k+1} M_{i'_l, j'_l} \right| \leq \frac{\gamma \bar{p} \alpha^{H-2k}}{(1-\gamma)^{k+2}}$$

This proves the bound. Now using the claim that

$$\frac{\partial^k M_{0, H+1}}{\partial \theta_\beta} = \sum_{n=1}^N c_n \zeta_n$$

where $|c_n| = \gamma^k$ and $N \leq 2^k k!$, we have shown that

$$\left| \frac{\partial^k M_{0, H+1}}{\partial \theta_\beta} \right| \leq \frac{\bar{p} 2^k \gamma^{k+1} k! \alpha^{H-2k}}{(1-\gamma)^{k+2}},$$

which completes the proof. ■

We are now ready to prove Proposition 4.1.

Proof:[Proof of Proposition 4.1] The k_{th} order partial derivative of $V^{\pi_\theta}(s_0)$ is equal to

$$\frac{\partial^k V^{\pi_\theta}(s_0)}{\partial \theta_{\beta_1} \dots \partial \theta_{\beta_k}} = \frac{\partial^k M_{0,H+1}^\theta}{\partial \theta_{\beta_1} \dots \partial \theta_{\beta_k}}.$$

Given vectors u^1, \dots, u^k which are unit vectors in \mathbb{R}^{H^k} (we denote the unit sphere by \mathbb{S}^{H^k}), the norm of this gradient tensor is given by:

$$\begin{aligned} \|\nabla_\theta^k V^{\pi_\theta}(s_0)\| &= \max_{u^1, \dots, u^k \in \mathbb{S}^{H^k}} \left| \sum_{\beta \in [H]^k} \frac{\partial^k V^{\pi_\theta}(s_0)}{\partial \theta_{\beta_1} \dots \partial \theta_{\beta_k}} u_{\beta_1}^1 \dots u_{\beta_k}^k \right| \\ &\leq \max_{u^1, \dots, u^k \in \mathbb{S}^{H^k}} \sqrt{\sum_{\beta \in [H]^k} \left(\frac{\partial^k V^{\pi_\theta}(s_0)}{\partial \theta_{\beta_1} \dots \partial \theta_{\beta_k}} \right)^2} \sqrt{\sum_{\beta \in [H]^k} (u_{\beta_1}^1 \dots u_{\beta_k}^k)^2} \\ &= \max_{u^1, \dots, u^k \in \mathbb{S}^{H^k}} \sqrt{\sum_{\beta \in [H]^k} \left(\frac{\partial^k V^{\pi_\theta}(s_0)}{\partial \theta_{\beta_1} \dots \partial \theta_{\beta_k}} \right)^2} \sqrt{\prod_{i=1}^k \|u^i\|_2^2} \\ &= \sqrt{\sum_{\beta \in [H]^k} \left(\frac{\partial^k V^{\pi_\theta}(s_0)}{\partial \theta_{\beta_1} \dots \partial \theta_{\beta_k}} \right)^2} = \sqrt{\sum_{\beta \in [H]^k} \left(\frac{\partial^k M_{0,H+1}^\theta}{\partial \theta_{\beta_1} \dots \partial \theta_{\beta_k}} \right)^2} \\ &\leq \sqrt{\frac{H^k \bar{p}^2 2^{2k} \gamma^{2k+2} (k!)^2 \alpha^{2H-4k}}{(1-\gamma)^{2k+4}}}, \end{aligned}$$

where the last inequality follows from Lemma B.2. In order to proceed further, we need an upper bound on the smallest admissible value of α . To do so, let us consider all possible parameters θ such that $\bar{p} \leq 1/4$ in accordance with the theorem statement. In order to bound α , it suffices to place an upper bound on the lower end of the range for α in Lemma B.1 (note Lemma B.1 holds for any choice of α in the range). Doing so, we see that

$$\begin{aligned} \frac{1 - \sqrt{1 - 4\gamma^2 \bar{p}(1 - \underline{p})}}{2\gamma(1 - \underline{p})} &\leq \frac{1 - 1 + 2\gamma \sqrt{\bar{p}(1 - \underline{p})}}{2\gamma(1 - \underline{p})} \\ &= \sqrt{\frac{\bar{p}}{1 - \underline{p}}} \leq \sqrt{\frac{4\bar{p}}{3}}, \end{aligned}$$

where the first inequality uses $\sqrt{x-y} \geq \sqrt{x} - \sqrt{y}$, by triangle inequality while the last inequality uses $\underline{p} \leq \bar{p} \leq 1/4$.

Hence, we have the bound

$$\begin{aligned}
\max_{u^1, \dots, u^k \in \mathbb{S}^{H^k}} \left| \sum_{\beta \in [H]^h} \frac{\partial^k V^{\pi_\theta}(s_0)}{\partial \theta_{\beta_1} \dots \partial \theta_{\beta_k}} u_{\beta_1}^1 \dots u_{\beta_k}^k \right| &\leq \sqrt{\frac{H^k \bar{p}^2 2^{2k} \gamma^{2k+2} (k!)^2 \left(\frac{4\bar{p}}{3}\right)^{H-2k}}{(1-\gamma)^{2k+4}}} \\
&\stackrel{(a)}{\leq} \sqrt{(H+1)^{2k+4} H^k \bar{p}^2 2^{2k} \gamma^{2k+2} (k!)^2 \left(\frac{4\bar{p}}{3}\right)^{H-2k}} \\
&\stackrel{(b)}{\leq} \sqrt{(2H)^{2k+4} H^k 2^{2k} (H)^{2k} \left(\frac{4\bar{p}}{3}\right)^{H-2k}} \\
&= \sqrt{(2)^{4k+4} (H)^{5k+4} \left(\frac{4\bar{p}}{3}\right)^{H-2k}}
\end{aligned}$$

where (a) uses $\gamma = H/(H+1)$, (b) follows since $\bar{p} \leq 1$, $H, k \geq 1$, $\gamma \leq 1$ and $k \leq H$.

Requiring that the gradient norm be no larger than $\left(\frac{4\bar{p}}{3}\right)^{H/4}$, we would like to satisfy

$$(2)^{4k+4} (H)^{5k+4} \left(\frac{4\bar{p}}{3}\right)^{H-2k} \leq \left(\frac{4\bar{p}}{3}\right)^{H/2},$$

for which it suffices to have

$$k \leq k_0 := \frac{\frac{H}{2} \log(3/4\bar{p}) - \log(2^4 H^4)}{\log(2^4 H^5) + 2 \log(3/4\bar{p})}.$$

Since,

$$\begin{aligned}
&\frac{\frac{H}{2} \log(3/4\bar{p}) - \log(2^4 H^4)}{\log(2^4 H^5) + 2 \log(3/4\bar{p})} \\
&\stackrel{(a)}{\geq} \frac{\frac{H}{2} \log(3/4\bar{p}) - \log(2^4 H^4)}{2 \log(2^4 H^5) 2 \log(3/4\bar{p})} \\
&\geq \frac{H}{8 \log(2^4 H^5)} - \frac{\log(2^4 H^4)}{4 \log(2^4 H^5) \log(3/4\bar{p})} \\
&\stackrel{(b)}{\geq} \frac{H}{8 \log(2^4 H^5)} - \frac{\log(2^4 H^4)}{4 \log(2^4 H^4) \log(3)} \\
&\geq \frac{H}{40 \log(2H)} - 1
\end{aligned}$$

where (a) follows from $a+b \leq 2ab$ when $a, b \geq 1$, (b) follows from $H \geq 1$ and $\bar{p} \leq 1/4$. Therefore, in order to obtain the smallest value of k_0 for all choices of $0 \leq \bar{p} < 1/4$, we further lower bound k_0 as

$$k_0 \geq \frac{H}{40 \log(2H)} - 1,$$

Thus, the norm of the gradient is bounded by $\left(\frac{4\bar{p}}{3}\right)^{H/4} \leq (1/3)^{H/4}$ for all $k \leq \frac{H}{40 \log(2H)} - 1$ as long as $\bar{p} \leq 1/4$, which gives the first part of the lemma.

For the second part, note that the optimal policy always chooses the action a_1 , and gets a discounted reward of

$$\gamma^{H+2}/(1-\gamma) = (H+1) \left(1 - \frac{1}{H+1}\right)^{H+2} \geq \frac{H+1}{8},$$

where the final inequality uses $(1 - 1/x)^x \geq 1/8$ for $x \geq 1$. On the other hand, the value of π_θ is upper bounded by

$$\begin{aligned} M_{0,H+1} &\leq \frac{\gamma \bar{p} \alpha^H}{(1-\gamma)^2} \leq \frac{\gamma \bar{p}}{(1-\gamma)^2} \left(\frac{4\bar{p}}{3}\right)^H \\ &\leq \frac{(H+1)^2}{3^H}. \end{aligned}$$

This gives the second part of the lemma. ■

C Proofs for Section 5

We first give a useful lemma about the structure of policy gradients for the softmax parameterization. We use the notation $\Pr^\pi(\tau|s_0 = s)$ to denote the probability of observing a trajectory τ when starting in state s and following the policy π and $\Pr_\mu^\pi(\tau)$ be $\mathbb{E}_{s \sim \mu}[\Pr^\pi(\tau|s_0 = s)]$ for a distribution μ over states.

Lemma C.1. *For the softmax policy class, we have:*

$$\frac{\partial V^{\pi_\theta}(\mu)}{\partial \theta_{s,a}} = \frac{1}{1-\gamma} d_\mu^{\pi_\theta}(s) \pi_\theta(a|s) A^{\pi_\theta}(s, a)$$

Proof: First note that

$$\frac{\partial \log \pi_\theta(a|s)}{\partial \theta_{s',a'}} = \mathbb{1}[s = s'] \left(\mathbb{1}[a = a'] - \pi_\theta(a'|s) \right) \quad (35)$$

where $\mathbb{1}[\mathcal{E}]$ is the indicator of \mathcal{E} being true.

Using this along with the policy gradient expression (6), we have:

$$\begin{aligned} \frac{\partial V^{\pi_\theta}(\mu)}{\partial \theta_{s,a}} &= \mathbb{E}_{\tau \sim \Pr_\mu^{\pi_\theta}} \left[\sum_{t=0}^{\infty} \gamma^t \mathbb{1}[s_t = s] \left(\mathbb{1}[a_t = a] A^{\pi_\theta}(s, a) - \pi_\theta(a|s) A^{\pi_\theta}(s_t, a_t) \right) \right] \\ &= \mathbb{E}_{\tau \sim \Pr_\mu^{\pi_\theta}} \left[\sum_{t=0}^{\infty} \gamma^t \mathbb{1}[(s_t, a_t) = (s, a)] A^{\pi_\theta}(s, a) \right] \\ &\quad - \pi_\theta(a|s) \sum_{t=0}^{\infty} \gamma^t \mathbb{E}_{\tau \sim \Pr_\mu^{\pi_\theta}} [\mathbb{1}[s_t = s] A^{\pi_\theta}(s_t, a_t)] \\ &= \frac{1}{1-\gamma} \mathbb{E}_{s' \sim d_\mu^{\pi_\theta}} \mathbb{E}_{a' \sim \pi_\theta(\cdot|s)} \left[\mathbb{1}[(s', a') = (s, a)] A^{\pi_\theta}(s, a) \right] - 0 \\ &= \frac{1}{1-\gamma} d_\mu^{\pi_\theta}(s) \pi_\theta(a|s) A^{\pi_\theta}(s, a), \end{aligned}$$

where the second to last step uses that for any policy $\sum_a \pi(a|s)A^\pi(s, a) = 0$. ■

C.1 Proofs for Section 5.1

We now prove Theorem 5.1, i.e. we show that for the updates given by

$$\theta^{(t+1)} = \theta^{(t)} + \eta \nabla V^{(t)}(\mu), \quad (36)$$

policy gradient converges to optimal policy for the softmax parameterization.

We prove this theorem by first proving a series of supporting lemmas. First, we show in Lemma C.2, that $V^{(t)}(s)$ is monotonically increasing for all states s using the fact that for appropriately chosen stepsizes GD makes monotonic improvement for smooth objectives.

Lemma C.2 (Monotonic Improvement in $V^{(t)}(s)$). *For all states s and actions a , for updates (36) with learning rate $\eta \leq \frac{(1-\gamma)^2}{5}$, we have*

$$V^{(t+1)}(s) \geq V^{(t)}(s); \quad Q^{(t+1)}(s, a) \geq Q^{(t)}(s, a).$$

Proof: The proof will consist of showing that:

$$\sum_{a \in \mathcal{A}} \pi^{(t+1)}(a|s)A^{(t)}(s, a) \geq \sum_{a \in \mathcal{A}} \pi^{(t)}(a|s)A^{(t)}(s, a) = 0. \quad (37)$$

holds for all states s . To see this, observe that since the above holds for all states s' , the performance difference lemma (Lemma 3.2) implies

$$V^{(t+1)}(s) - V^{(t)}(s) = \frac{1}{1-\gamma} \mathbb{E}_{s' \sim d_s^{(t+1)}} \mathbb{E}_{a \sim \pi^{(t+1)}(\cdot|s')} [A^{(t)}(s', a)] \geq 0,$$

which would complete the proof.

Let us use the notation $\theta_s \in \mathbb{R}^{|\mathcal{A}|}$ to refer to the vector of θ_s , for some fixed state s . Define the function

$$F_s(\theta_s) := \sum_{a \in \mathcal{A}} \pi_{\theta_s}(a|s)c(s, a) \quad (38)$$

where $c(s, a)$ is constant, which we later set to be $A^{(t)}(s, a)$; note we do not treat $c(s, a)$ as a function of θ . Thus,

$$\begin{aligned} \frac{\partial F_s(\theta_s)}{\partial \theta_{s,a}} \Big|_{\theta_s^{(t)}} &= \sum_{a' \in \mathcal{A}} \frac{\partial \pi_{\theta_s}(a'|s)}{\partial \theta_{s,a}} \Big|_{\theta_s^{(t)}} c(s, a') \\ &= \pi^{(t)}(a|s)(1 - \pi^{(t)}(a|s))c(s, a) - \sum_{a' \neq a} \pi^{(t)}(a|s)\pi^{(t)}(a'|s)c(s, a') \\ &= \pi^{(t)}(a|s) \left(c(s, a) - \sum_{a' \in \mathcal{A}} \pi^{(t)}(a'|s)c(s, a') \right) \end{aligned}$$

Taking $c(s, a)$ to be $A^{(t)}(s, a)$ implies $\sum_{a' \in \mathcal{A}} \pi^{(t)}(a'|s)c(s, a') = \sum_{a' \in \mathcal{A}} \pi^{(t)}(a'|s)A^{(t)}(s, a') = 0$,

$$\left. \frac{\partial F_s(\theta_s)}{\partial \theta_{s,a}} \right|_{\theta_s^{(t)}} = \pi^{(t)}(a|s)A^{(t)}(s, a) \quad (39)$$

Observe that for the softmax parameterization,

$$\theta_s^{(t+1)} = \theta_s^{(t)} + \eta \nabla_s V^{(t)}(\mu)$$

where ∇_s is gradient w.r.t. θ_s and from Lemma C.1 that:

$$\frac{\partial V^{(t)}(\mu)}{\partial \theta_{s,a}} = \frac{1}{1-\gamma} d_{\mu}^{\pi^{(t)}}(s) \pi^{(t)}(a|s) A^{(t)}(s, a)$$

This gives using Equation (39)

$$\theta_s^{(t+1)} = \theta_s^{(t)} + \eta \frac{1}{1-\gamma} d_{\mu}^{\pi^{(t)}}(s) \nabla_s F_s(\theta_s) \Big|_{\theta_s^{(t)}}$$

Recall that for a β smooth function, gradient ascent will decrease the function value provided that $\eta \leq 1/\beta$ (Theorem E.1). Because $F_s(\theta_s)$ is β -smooth for $\beta = \frac{5}{1-\gamma}$ (Lemma D.1 and $|A^{(t)}(s, a)| \leq \frac{1}{1-\gamma}$), then our assumption that

$$\eta \leq \frac{(1-\gamma)^2}{5} = (1-\gamma)\beta^{-1}$$

implies that $\eta \frac{1}{1-\gamma} d_{\mu}^{\pi^{(t)}}(s) \leq 1/\beta$, and so we have

$$F_s(\theta_s^{(t+1)}) \geq F_s(\theta_s^{(t)})$$

which implies (37). ■

Next, we show the limit for iterates $V^{(t)}(s)$ and $Q^{(t)}(s, a)$ exists for all states s and actions a .

Lemma C.3. *For all states s and actions a , there exists values $V^{(\infty)}(s)$ and $Q^{(\infty)}(s, a)$ such that as $t \rightarrow \infty$, $V^{(t)}(s) \rightarrow V^{(\infty)}(s)$ and $Q^{(t)}(s, a) \rightarrow Q^{(\infty)}(s, a)$. Define*

$$\Delta = \min_{\{s,a|A^{(\infty)}(s,a) \neq 0\}} |A^{(\infty)}(s, a)|$$

where $A^{(\infty)}(s, a) = Q^{(\infty)}(s, a) - V^{(\infty)}(s)$. Furthermore, there exists a T_0 such that for all $t > T_0$, $s \in \mathcal{S}$, and $a \in \mathcal{A}$, we have

$$Q^{(t)}(s, a) \geq Q^{(\infty)}(s, a) - \Delta/4 \quad (40)$$

Proof: Observe that $Q^{(t+1)}(s, a) \geq Q^{(t)}(s, a)$ (by Lemma C.2) and $Q^{(t)}(s, a) \leq \frac{1}{1-\gamma}$, therefore by monotone convergence theorem, $Q^{(t)}(s, a) \rightarrow Q^{(\infty)}(s, a)$ for some constant $Q^{(\infty)}(s, a)$. Similarly it follows that $V^{(t)}(s) \rightarrow V^{(\infty)}(s)$ for some constant $V^{(\infty)}(s)$. Due to the limits existing, this implies we can choose T_0 , such that the result (40) follows. ■

Based on the limits $V^{(\infty)}(s)$ and $Q^{(\infty)}(s, a)$, define following sets:

$$\begin{aligned} I_0^s &:= \{a | Q^{(\infty)}(s, a) = V^{(\infty)}(s)\} \\ I_+^s &:= \{a | Q^{(\infty)}(s, a) > V^{(\infty)}(s)\} \\ I_-^s &:= \{a | Q^{(\infty)}(s, a) < V^{(\infty)}(s)\}. \end{aligned}$$

In the following lemmas C.5- C.11, we first show that probabilities $\pi^{(t)}(a|s) \rightarrow 0$ for actions $a \in I_+^s \cup I_-^s$ as $t \rightarrow \infty$. We then show that for actions $a \in I_-^s$, $\lim_{t \rightarrow \infty} \theta_{s,a}^{(t)} = -\infty$ and for all actions $a \in I_+^s$, $\theta^{(t)}(a|s)$ is bounded from below as $t \rightarrow \infty$.

Lemma C.4. *We have that there exists a T_1 such that for all $t > T_1$, $s \in \mathcal{S}$, and $a \in \mathcal{A}$, we have*

$$A^{(t)}(s, a) < -\frac{\Delta}{4} \text{ for } a \in I_-^s; \quad A^{(t)}(s, a) > \frac{\Delta}{4} \text{ for } a \in I_+^s \quad (41)$$

Proof: Since, $V^{(t)}(s) \rightarrow V^{(\infty)}(s)$, we have that there exists $T_1 > T_0$ such that for all $t > T_1$,

$$V^{(t)}(s) > V^{(\infty)}(s) - \frac{\Delta}{4}. \quad (42)$$

Using Equation (40), it follows that for $t > T_1 > T_0$, for $a \in I_-^s$

$$\begin{aligned} A^{(t)}(s, a) &= Q^{(t)}(s, a) - V^{(t)}(s) \\ &\leq Q^{(\infty)}(s, a) - V^{(t)}(s) \\ &\leq Q^{(\infty)}(s, a) - V^{(\infty)}(s) + \Delta/4 && \text{(Equation (42))} \\ &\leq -\Delta + \Delta/4 && \text{(definition of } I_-^s \text{ and Lemma C.3)} \\ &< -\Delta/4 \end{aligned}$$

Similarly $A^{(t)}(s, a) = Q^{(t)}(s, a) - V^{(t)}(s) > \Delta/4$ for $a \in I_+^s$ as

$$\begin{aligned} A^{(t)}(s, a) &= Q^{(t)}(s, a) - V^{(t)}(s) \\ &\geq Q^{(\infty)}(s, a) - \Delta/4 - V^{(t)}(s) && \text{(Lemma C.3)} \\ &\geq Q^{(\infty)}(s, a) - V^{(\infty)}(s) - \Delta/4 && (V^{(\infty)}(s) \geq V^{(t)}(s) \text{ from Lemma C.2)} \\ &\geq \Delta - \Delta/4 \\ &> \Delta/4 \end{aligned}$$

which completes the proof. ■

Lemma C.5. $\frac{\partial V^{(t)}(\mu)}{\partial \theta_{s,a}} \rightarrow 0$ as $t \rightarrow \infty$ for all states s and actions a . This implies that for $a \in I_+^s \cup I_-^s$, $\pi^{(t)}(a|s) \rightarrow 0$ and that $\sum_{a \in I_0^s} \pi^{(t)}(a|s) \rightarrow 1$.

Proof: Because $V^{\pi_\theta}(\mu)$ is smooth (Lemma D.4) as a function of θ , it follows from standard optimization results (Theorem E.1) that $\frac{\partial V^{(t)}(\mu)}{\partial \theta_{s,a}} \rightarrow 0$ for all states s and actions a . We have from Lemma C.1

$$\frac{\partial V^{(t)}(\mu)}{\partial \theta_{s,a}} = \frac{1}{1-\gamma} d_\mu^{\pi^{(t)}}(s) \pi^{(t)}(a|s) A^{(t)}(s, a).$$

Since, $|A^{(t)}(s, a)| > \frac{\Delta}{4}$ for all $t > T_1$ (from Lemma C.4) for all $a \in I_-^s \cup I_+^s$ and $d_\mu^{\pi^{(t)}}(s) \geq \frac{\mu(s)}{1-\gamma} > 0$ (using the strict positivity of μ in our assumption in Theorem 5.1), we have $\pi^{(t)}(a|s) \rightarrow 0$. ■

Lemma C.6. (Monotonicity in $\theta_{s,a}^{(t)}$). For all $a \in I_+^s$, $\theta_{s,a}^{(t)}$ is strictly increasing for $t \geq T_1$. For all $a \in I_-^s$, $\theta_{s,a}^{(t)}$ is strictly decreasing for $t \geq T_1$.

Proof: We have from Lemma C.1

$$\frac{\partial V^{(t)}(\mu)}{\partial \theta_{s,a}} = \frac{1}{1-\gamma} d_\mu^{\pi^{(t)}}(s) \pi^{(t)}(a|s) A^{(t)}(s, a)$$

From Lemma C.4, we have for all $t > T_1$

$$A^{(t)}(s, a) > 0 \text{ for } a \in I_+^s; \quad A^{(t)}(s, a) < 0 \text{ for } a \in I_-^s$$

Since $d_\mu^{\pi^{(t)}}(s) > 0$ and $\pi^{(t)}(a|s) > 0$ for the softmax parameterization, we have for all $t > T_1$

$$\frac{\partial V^{(t)}(\mu)}{\partial \theta_{s,a}} > 0 \text{ for } a \in I_+^s; \quad \frac{\partial V^{(t)}(\mu)}{\partial \theta_{s,a}} < 0 \text{ for } a \in I_-^s$$

This implies for all $a \in I_+^s$, $\theta_{s,a}^{(t+1)} - \theta_{s,a}^{(t)} = \frac{\partial V^{(t)}(\mu)}{\partial \theta_{s,a}} > 0$ i.e. $\theta_{s,a}^{(t)}$ is strictly increasing for $t \geq T_1$. The second claim follows similarly. ■

Lemma C.7. For all s where $I_+^s \neq \emptyset$, we have that:

$$\max_{a \in I_0^s} \theta_{s,a}^{(t)} \rightarrow \infty, \quad \min_{a \in \mathcal{A}} \theta_{s,a}^{(t)} \rightarrow -\infty$$

Proof: Since $I_+^s \neq \emptyset$, there exists some action $a_+ \in I_+^s$. From Lemma C.5,

$$\pi^{(t)}(a_+|s) \rightarrow 0, \text{ as } t \rightarrow \infty$$

or equivalently by softmax parameterization,

$$\frac{\exp(\theta_{s,a_+}^{(t)})}{\sum_a \exp(\theta_{s,a}^{(t)})} \rightarrow 0, \text{ as } t \rightarrow \infty$$

From Lemma C.6, for any action $a \in I_+^s$ and in particular for a_+ , $\theta_{s,a_+}^{(t)}$ is monotonically increasing for $t > T_1$. That is the numerator in previous display is monotonically increasing. Therefore, the denominator should go to infinity i.e.

$$\sum_a \exp(\theta_{s,a}^{(t)}) \rightarrow \infty, \text{ as } t \rightarrow \infty.$$

From Lemma C.5,

$$\sum_{a \in I_0^s} \pi^{(t)}(a|s) \rightarrow 1, \text{ as } t \rightarrow \infty$$

or equivalently

$$\frac{\sum_{a \in I_0^s} \exp(\theta_{s,a}^{(t)})}{\sum_a \exp(\theta_{s,a}^{(t)})} \rightarrow 1, \text{ as } t \rightarrow \infty$$

Since, denominator goes to ∞ ,

$$\sum_{a \in I_0^s} \exp(\theta_{s,a}^{(t)}) \rightarrow \infty, \text{ as } t \rightarrow \infty$$

which implies

$$\max_{a \in I_0^s} \theta_{s,a}^{(t)} \rightarrow \infty, \text{ as } t \rightarrow \infty$$

Note this also implies $\max_{a \in \mathcal{A}} \theta_{s,a}^{(t)} \rightarrow \infty$. The last part of the proof is completed using that the gradients sum to 0, i.e. $\sum_a \frac{\partial V^{(t)}(\mu)}{\partial \theta_{s,a}} = 0$. From gradient sum to 0, we get that $\sum_{a \in \mathcal{A}} \theta_{s,a}^{(t)} = \sum_{a \in \mathcal{A}} \theta_{s,a}^{(0)} := c$ for all $t > 0$ where c is defined as the sum (over \mathcal{A}) of initial parameters. That is $\min_{a \in \mathcal{A}} \theta_{s,a}^{(t)} < -\frac{1}{|\mathcal{A}|} \max_{a \in \mathcal{A}} \theta_{s,a}^{(t)} + c$. Since, $\max_{a \in \mathcal{A}} \theta_{s,a}^{(t)} \rightarrow \infty$, the result follows. ■

Lemma C.8. *Suppose $a_+ \in I_+^s$. For any $a \in I_0^s$, if there exists a $t \geq T_0$ such that $\pi^{(t)}(a|s) \leq \pi^{(t)}(a_+|s)$, then for all $\tau \geq t$, $\pi^{(\tau)}(a|s) \leq \pi^{(\tau)}(a_+|s)$.*

Proof: The proof is inductive. Suppose $\pi^{(t)}(a|s) \leq \pi^{(t)}(a_+|s)$, this implies from Lemma C.1

$$\begin{aligned} \frac{\partial V^{(t)}(\mu)}{\partial \theta_{s,a}} &= \frac{1}{1-\gamma} d_\mu^{(t)}(s) \pi^{(t)}(a|s) \left(Q^{(t)}(s, a) - V^{(t)}(s) \right) \\ &\leq \frac{1}{1-\gamma} d_\mu^{(t)}(s) \pi^{(t)}(a_+|s) \left(Q^{(t)}(s, a_+) - V^{(t)}(s) \right) = \frac{\partial V^{(t)}(\mu)}{\partial \theta_{s,a_+}}. \end{aligned}$$

where the second to last step follows from $Q^{(t)}(s, a_+) \geq Q^{(\infty)}(s, a_+) - \Delta/4 \geq Q^{(\infty)}(s, a) + \Delta - \Delta/4 > Q^{(t)}(s, a)$ for $t > T_0$. This implies that $\pi^{(t+1)}(a|s) \leq \pi^{(t+1)}(a_+|s)$ which completes the proof. ■

Consider an arbitrary $a_+ \in I_+^s$. Let us partition the set I_0^s into $B_0^s(a_+)$ and $\bar{B}_0^s(a_+)$ as follows: $B_0^s(a_+)$ is the set of all $a \in I_0^s$ such that for all $t \geq T_0$, $\pi^{(t)}(a_+|s) < \pi^{(t)}(a|s)$, and $\bar{B}_0^s(a_+)$ contains the remainder of the actions from I_0^s . We drop the argument (a_+) when clear from the context.

Lemma C.9. *Suppose $I_+^s \neq \emptyset$. For all $a_+ \in I_+^s$, we have that $B_0^s(a_+) \neq \emptyset$ and that*

$$\sum_{a \in B_0^s(a_+)} \pi^{(t)}(a|s) \rightarrow 1, \text{ as } t \rightarrow \infty.$$

This implies that:

$$\max_{a \in B_0^s(a_+)} \theta_{s,a}^{(t)} \rightarrow \infty.$$

Proof: Let $a_+ \in I_+^s$. Consider any $a \in \bar{B}_0^s$. Then, by definition of \bar{B}_0^s , there exists $t' > T_0$ such that $\pi^{(t)}(a_+|s) \geq \pi^{(t)}(a|s)$. From Lemma C.8, for all $\tau > t$ $\pi^{(\tau)}(a_+|s) \geq \pi^{(\tau)}(a|s)$. Also, since $\pi^{(t)}(a_+|s) \rightarrow 0$, this implies

$$\pi^{(t)}(a|s) \rightarrow 0 \text{ for all } a \in \bar{B}_0^s$$

Since, $B_0^s \cup \bar{B}_0^s = I_0^s$ and $\sum_{a \in I_0^s} \pi^{(t)}(a|s) \rightarrow 1$ (from Lemma C.5), this implies that $B_0^s \neq \emptyset$ and that means

$$\sum_{a \in B_0^s} \pi^{(t)}(a|s) \rightarrow 1, \text{ as } t \rightarrow \infty,$$

which completes the proof of the first claim. The proof of the second claim is identical to the proof in Lemma C.7 where instead of $\sum_{a \in I_0^s} \pi^{(t)}(a|s) \rightarrow 1$, we use $\sum_{a \in B_0^s} \pi^{(t)}(a|s) \rightarrow 1$. ■

Lemma C.10. Consider any s where $I_+^s \neq \emptyset$. Then, for any $a_+ \in I_+^s$, there exists an iteration T_{a_+} such that for all $t > T_{a_+}$,

$$\pi^{(t)}(a_+|s) > \pi^{(t)}(a|s)$$

for all $a \in \bar{B}_0^s(a_+)$.

Proof: The proof follows from definition of $\bar{B}_0^s(a_+)$. That is if $a \in \bar{B}_0^s(a_+)$, then there exists a iteration $t_a > T_0$ such that $\pi^{(t_a)}(a_+|s) > \pi^{(t_a)}(a|s)$. Then using Lemma C.8, for all $\tau > t_a$, $\pi^{(\tau)}(a_+|s) > \pi^{(\tau)}(a|s)$. Choosing

$$T_{a_+} = \max_{a \in \bar{B}_0^s(a_+)} t_a$$

completes the proof. ■

Lemma C.11. For all actions $a \in I_+^s$, we have that $\theta_{s,a}^{(t)}$ is bounded from below as $t \rightarrow \infty$. For all actions $a \in I_-^s$, we have that $\theta_{s,a}^{(t)} \rightarrow -\infty$ as $t \rightarrow \infty$.

Proof: For the first claim, from Lemma C.6, we know that after T_1 , $\theta_{s,a}^{(t)}$ is strictly increasing for $a \in I_+^s$, i.e. for all $t > T_1$

$$\theta_{s,a}^{(t)} \geq \theta_{s,a}^{(T_1)}.$$

For the second claim, we know that after T_1 , $\theta_{s,a}^{(t)}$ is strictly decreasing for $a \in I_-^s$ (Lemma C.6). Therefore, by monotone convergence theorem, $\lim_{t \rightarrow \infty} \theta_{s,a}^{(t)}$ exists and is either $-\infty$ or some constant θ_0 . We now prove the second claim by contradiction. Suppose $a \in I_-^s$ and that there exists a θ_0 , such that $\theta_{s,a}^{(t)} > \theta_0$, for $t \geq T_1$. By Lemma C.7, there must exist an action where $a' \in \mathcal{A}$ such that

$$\liminf_{t \rightarrow \infty} \theta_{s,a'}^{(t)} = -\infty. \quad (43)$$

Let us consider some $\delta > 0$ such that $\theta_{s,a'}^{(T_1)} \geq \theta_0 - \delta$. Now for $t \geq T_1$ define $\tau(t)$ as follows: $\tau(t) = k$ if k is the largest iteration in the interval $[T_1, t]$ such that $\theta_{s,a'}^{(k)} \geq \theta_0 - \delta$ (i.e. $\tau(t)$ is the latest iteration before $\theta_{s,a'}$ crosses below $\theta_0 - \delta$). Define $\mathcal{T}^{(t)}$ as the subsequence of iterations $\tau(t) < t' < t$ such that $\theta_{s,a'}^{(t')}$ decreases, i.e.

$$\frac{\partial V^{(t')}(\mu)}{\partial \theta_{s,a'}} \leq 0, \text{ for } \tau(t) < t' < t.$$

Define Z_t as the sum (if $\mathcal{T}^{(t)} = \emptyset$, we define $Z_t = 0$):

$$Z_t = \sum_{t' \in \mathcal{T}^{(t)}} \frac{\partial V^{(t')}(\mu)}{\partial \theta_{s,a'}}.$$

For non-empty $\mathcal{T}^{(t)}$, this gives:

$$\begin{aligned} Z_t &= \sum_{t' \in \mathcal{T}^{(t)}} \frac{\partial V^{(t')}(\mu)}{\partial \theta_{s,a'}} \leq \sum_{t'=\tau(t)-1}^{t-1} \frac{\partial V^{(t')}(\mu)}{\partial \theta_{s,a'}} \leq \sum_{t'=\tau(t)}^{t-1} \frac{\partial V^{(t')}(\mu)}{\partial \theta_{s,a'}} + \frac{1}{1-\gamma^2} \\ &= \frac{1}{\eta}(\theta_{s,a'}^{(t)} - \theta_{s,a'}^{(\tau(t))}) + \frac{1}{1-\gamma^2} \leq \frac{1}{\eta} \left(\theta_{s,a'}^{(t)} - (\theta_0 - \delta) \right) + \frac{1}{1-\gamma^2}, \end{aligned}$$

where we have used that $|\frac{\partial V^{(t')}(\mu)}{\partial \theta_{s,a'}}| \leq 1/(1-\gamma)$. By (43), this implies that:

$$\liminf_{t \rightarrow \infty} Z_t = -\infty.$$

For any $\mathcal{T}^{(t)} \neq \emptyset$, this implies that for all $t' \in \mathcal{T}^{(t)}$, from Lemma C.1

$$\begin{aligned} \left| \frac{\partial V^{(t')}(\mu)/\partial \theta_{s,a}}{\partial V^{(t')}(\mu)/\partial \theta_{s,a'}} \right| &= \left| \frac{\pi^{(t')}(a|s)A^{(t')}(s,a)}{\pi^{(t')}(a'|s)A^{(t')}(s,a')} \right| \geq \exp(\theta_0 - \theta_{s,a'}^{(t)}) \frac{(1-\gamma)\Delta}{4} \\ &\geq \exp(\delta) \frac{(1-\gamma)\Delta}{4} \end{aligned}$$

where we have used that $|A^{(t')}(s,a')| \leq 1/(1-\gamma)$ and $|A^{(t')}(s,a)| \geq \frac{\Delta}{4}$ for all $t' > T_1$ (from Lemma C.4). Note that since $\frac{\partial V^{(t')}(\mu)}{\partial \theta_{s,a}} < 0$ and $\frac{\partial V^{(t')}(\mu)}{\partial \theta_{s,a'}} < 0$ over the subsequence $\mathcal{T}^{(t)}$, the sign of the inequality reverses. In particular, for any $\mathcal{T}^{(t)} \neq \emptyset$

$$\begin{aligned} \frac{1}{\eta}(\theta_{s,a}^{(T_1)} - \theta_{s,a}^{(t)}) &= \sum_{t'=T_1}^{t-1} \frac{\partial V^{(t')}(\mu)}{\partial \theta_{s,a}} \leq \sum_{t' \in \mathcal{T}^{(t)}} \frac{\partial V^{(t')}(\mu)}{\partial \theta_{s,a}} \\ &\leq \exp(\delta) \frac{(1-\gamma)\Delta}{4} \sum_{t' \in \mathcal{T}^{(t)}} \frac{\partial V^{(t')}(\mu)}{\partial \theta_{s,a'}} \\ &= \exp(\delta) \frac{(1-\gamma)\Delta}{4} Z_t \end{aligned}$$

where the first step follows from that $\theta_{s,a}^{(t)}$ is monotonically decreasing, i.e. $\frac{\partial V^{(t)}(\mu)}{\partial \theta_{s,a}} < 0$ for $t \notin \mathcal{T}$ (Lemma C.6). Since,

$$\liminf_{t \rightarrow \infty} Z_t = -\infty,$$

this contradicts that $\theta_{s,a}^{(t)}$ is lower bounded from below, which completes the proof. \blacksquare

Lemma C.12. Consider any s where $I_+^s \neq \emptyset$. Then, for any $a_+ \in I_+^s$,

$$\sum_{a \in B_0^s(a_+)} \theta_{s,a}^{(t)} \rightarrow \infty, \text{ as } t \rightarrow \infty$$

Proof: Consider any $a \in B_0^s$. We have by definition of B_0^s that $\pi^{(t)}(a_+|s) < \pi^{(t)}(a|s)$ for all $t > T_0$. This implies by the softmax parameterization that $\theta_{s,a_+}^{(t)} < \theta_{s,a}^{(t)}$. Since, $\theta_{s,a_+}^{(t)}$ is lower bounded as $t \rightarrow \infty$ (using Lemma C.11), this implies $\theta_{s,a}^{(t)}$ is lower bounded as $t \rightarrow \infty$ for all $a \in B_0^s$. This in conjunction with $\max_{a \in B_0^s(a_+)} \theta_{s,a}^{(t)} \rightarrow \infty$ implies

$$\sum_{a \in B_0^s} \theta_{s,a}^{(t)} \rightarrow \infty, \quad (44)$$

which proves this claim. ■

We are now ready to complete the proof for Theorem 5.1. We prove it by showing that I_+^s is empty for all states s or equivalently $V^{(t)}(s_0) \rightarrow V^*(s_0)$ as $t \rightarrow \infty$.

Proof:[Proof for Theorem 5.1] Suppose the set I_+^s is non-empty for some s , else the proof is complete. Let $a_+ \in I_+^s$. Then, from Lemma C.12,

$$\sum_{a \in B_0^s} \theta_{s,a}^{(t)} \rightarrow \infty, \quad (45)$$

Now we proceed by showing a contradiction. For $a \in I_-^s$, we have that since $\frac{\pi^{(t)}(a|s)}{\pi^{(t)}(a_+|s)} = \exp(\theta_{s,a}^{(t)} - \theta_{s,a_+}^{(t)}) \rightarrow 0$ (as $\theta_{s,a_+}^{(t)}$ is lower bounded and $\theta_{s,a}^{(t)} \rightarrow -\infty$ by Lemma C.11), there exists $T_2 > T_0$ such that

$$\frac{\pi^{(t)}(a|s)}{\pi^{(t)}(a_+|s)} < \frac{(1-\gamma)\Delta}{16|\mathcal{A}|}$$

or, equivalently,

$$-\sum_{a \in I_-^s} \frac{\pi^{(t)}(a|s)}{1-\gamma} > -\pi^{(t)}(a_+|s) \frac{\Delta}{16}. \quad (46)$$

For $a \in \bar{B}_0^s$, we have $A^{(t)}(s, a) \rightarrow 0$ (by definition of set I_0^s and $\bar{B}_0^s \subset I_0^s$) and $1 < \frac{\pi^{(t)}(a_+|s)}{\pi^{(t)}(a|s)}$ for all $t > T_{a_+}$ from Lemma C.10. Thus, there exists $T_3 > T_2, T_{a_+}$ such that

$$|A^{(t)}(s, a)| < \frac{\pi^{(t)}(a_+|s)}{\pi^{(t)}(a|s)} \frac{\Delta}{16|\mathcal{A}|}$$

which implies

$$\sum_{a \in \bar{B}_0^s} \pi^{(t)}(a|s) |A^{(t)}(s, a)| < \pi^{(t)}(a_+|s) \frac{\Delta}{16}$$

$$-\pi^{(t)}(a_+|s)\frac{\Delta}{16} < \sum_{a \in \bar{B}_0^s} \pi^{(t)}(a|s)A^{(t)}(s, a) < \pi^{(t)}(a_+|s)\frac{\Delta}{16} \quad (47)$$

We have for $t > T_3$, from $\sum_{a \in \mathcal{A}} \pi^{(t)}(a|s)A^{(t)}(s, a) = 0$,

$$\begin{aligned} 0 &= \sum_{a \in I_0^s} \pi^{(t)}(a|s)A^{(t)}(s, a) + \sum_{a \in I_+^s} \pi^{(t)}(a|s)A^{(t)}(s, a) + \sum_{a \in I_-^s} \pi^{(t)}(a|s)A^{(t)}(s, a) \\ &\stackrel{(a)}{\geq} \sum_{a \in B_0^s} \pi^{(t)}(a|s)A^{(t)}(s, a) + \sum_{a \in \bar{B}_0^s} \pi^{(t)}(a|s)A^{(t)}(s, a) + \pi^{(t)}(a_+|s)A^{(t)}(s, a_+) \\ &\quad + \sum_{a \in I_-^s} \pi^{(t)}(a|s)A^{(t)}(s, a) \\ &\stackrel{(b)}{\geq} \sum_{a \in B_0^s} \pi^{(t)}(a|s)A^{(t)}(s, a) + \sum_{a \in \bar{B}_0^s} \pi^{(t)}(a|s)A^{(t)}(s, a) + \pi^{(t)}(a_+|s)\frac{\Delta}{4} - \sum_{a \in I_-^s} \frac{\pi^{(t)}(a|s)}{1-\gamma} \\ &\stackrel{(c)}{>} \sum_{a \in B_0^s} \pi^{(t)}(a|s)A^{(t)}(s, a) - \pi^{(t)}(a_+|s)\frac{\Delta}{16} + \pi^{(t)}(a_+|s)\frac{\Delta}{4} - \pi^{(t)}(a_+|s)\frac{\Delta}{16} \\ &> \sum_{a \in B_0^s} \pi^{(t)}(a|s)A^{(t)}(s, a) \end{aligned}$$

where in the step (a), we used $A^{(t)}(s, a) > 0$ for all actions $a \in I_+^s$ for $t > T_3 > T_1$ from Lemma C.4. In the step (b), we used $A^{(t)}(s, a_+) \geq \frac{\Delta}{4}$ for $t > T_3 > T_1$ from Lemma C.4 and $A^{(t)}(s, a) \geq -\frac{1}{1-\gamma}$. In the step (c), we used Equation (46) and left inequality in (47). This implies that for all $t > T_3$

$$\sum_{a \in B_0^s} \frac{\partial V^{(t)}(\mu)}{\partial \theta_{s,a}} < 0$$

This contradicts Equation (45) which requires

$$\lim_{t \rightarrow \infty} \sum_{a \in B_0^s} (\theta_{s,a}^{(t)} - \theta_{s,a}^{(T_3)}) = \eta \sum_{t=T_3}^{\infty} \sum_{a \in B_0^s} \frac{\partial V^{(t)}(\mu)}{\partial \theta_{s,a}} \rightarrow \infty.$$

Therefore, the set I_+^s must be empty, which completes the proof. \blacksquare

C.2 Proofs for Section 5.2

Proof:[of Corollary 5.1] Using Theorem 5.2, the desired optimality gap ϵ will follow if we set

$$\lambda = \frac{\epsilon(1-\gamma)}{2 \left\| \frac{d_{\rho}^{\pi^*}}{\mu} \right\|_{\infty}} \quad (48)$$

and if $\|\nabla_{\theta} L_{\lambda}(\theta)\|_2 \leq \lambda/(2|\mathcal{S}||\mathcal{A}|)$. In order to complete the proof, we need to bound the iteration complexity of making the gradient sufficiently small.

Since the optimization is deterministic and unconstrained, we can appeal to standard results (Theorem E.1) which give that after T iterations of gradient ascent with stepsize of $1/\beta_{\lambda}$, we have

$$\min_{t \leq T} \|\nabla_{\theta} L_{\lambda}(\theta^{(t)})\|_2^2 \leq \frac{2\beta_{\lambda}(L_{\lambda}(\theta^*) - L_{\lambda}(\theta^{(0)}))}{T} \leq \frac{2\beta_{\lambda}}{(1-\gamma)T}, \quad (49)$$

where β_{λ} is an upper bound on the smoothness of $L_{\lambda}(\theta)$. We seek to ensure

$$\epsilon_{\text{opt}} \leq \sqrt{\frac{2\beta_{\lambda}}{(1-\gamma)T}} \leq \frac{\lambda}{2|\mathcal{S}||\mathcal{A}|}$$

Choosing $T \geq \frac{8\beta_{\lambda}|\mathcal{S}||\mathcal{A}|^2}{(1-\gamma)\lambda^2}$ satisfies the above inequality. By Lemma D.4, we can take $\beta_{\lambda} = \frac{8\gamma}{(1-\gamma)^3} + \frac{2\lambda}{|\mathcal{S}|}$, and so

$$\begin{aligned} \frac{8\beta_{\lambda}|\mathcal{S}||\mathcal{A}|^2}{(1-\gamma)\lambda^2} &\leq \frac{64|\mathcal{S}||\mathcal{A}|^2}{(1-\gamma)^4\lambda^2} + \frac{16|\mathcal{S}||\mathcal{A}|^2}{(1-\gamma)\lambda} \\ &\leq \frac{80|\mathcal{S}||\mathcal{A}|^2}{(1-\gamma)^4\lambda^2} \\ &= \frac{320|\mathcal{S}||\mathcal{A}|^2}{(1-\gamma)^6\epsilon^2} \left\| \frac{d_{\rho}^{\pi^*}}{\mu} \right\|_{\infty}^2 \end{aligned}$$

where we have used that $\lambda < 1$. This completes the proof. \blacksquare

C.3 Proofs for Section 5.3

Proof:[of Lemma 5.1] Following the definition of compatible function approximation in Sutton et al. [1999], which was also invoked in Kakade [2001], for a vector $w \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|}$, we define the error function

$$L^{\theta}(w) = \mathbb{E}_{s \sim d_{\rho}^{\pi_{\theta}}, a \sim \pi_{\theta}(\cdot|s)} (w^{\top} \nabla_{\theta} \log \pi_{\theta}(\cdot|s) - A^{\pi_{\theta}}(s, a))^2.$$

Let w_{θ}^* be the minimizer of $L^{\theta}(w)$ with the smallest ℓ_2 norm. Then by definition of Moore-Penrose pseudoinverse, it is easily seen that

$$w_{\theta}^* = F_{\rho}(\theta)^{\dagger} \mathbb{E}_{s \sim d_{\rho}^{\pi_{\theta}}, a \sim \pi_{\theta}(a|s)} [\nabla_{\theta} \log \pi_{\theta}(a|s) A^{\pi_{\theta}}(s, a)] = (1-\gamma) F_{\rho}(\theta)^{\dagger} \nabla_{\theta} V^{\pi_{\theta}}(\rho). \quad (50)$$

In other words, w_{θ}^* is precisely proportional to the NPG update direction. Note further that for the Softmax policy parameterization, we have by (35),

$$w^{\top} \nabla_{\theta} \log \pi_{\theta}(a|s) = w_{s,a} - \sum_{a' \in \mathcal{A}} w_{s,a'} \pi_{\theta}(a'|s).$$

Since $\sum_{a \in \mathcal{A}} \pi(a|s)A^\pi(s, a) = 0$, this immediately yields that $L^\theta(A^{\pi_\theta}) = 0$. However, this might not be the unique minimizer of L^θ , which is problematic since $w^*(\theta)$ as defined in terms of the Moore-Penrose pseudoinverse is formally the smallest norm solution to the least-squares problem, which A^{π_θ} may not be. However, given any vector $v \in \mathbb{R}^{|\mathcal{S}||\mathcal{A}|}$, let us consider solutions of the form $A^{\pi_\theta} + v$. Due to the form of the derivatives of the policy for the softmax parameterization (recall Equation 35), we have for any state s, a such that s is reachable under ρ ,

$$v^\top \nabla_\theta \log \pi_\theta(a|s) = \sum_{a' \in \mathcal{A}} (v_{s,a'} \mathbb{1}[a = a'] - v_{s,a'} \pi_\theta(a'|s)) = v_{s,a} - \sum_{a' \in \mathcal{A}} v_{s,a'} \pi_\theta(a'|s).$$

Note that here we have used that π_θ is a stochastic policy with $\pi_\theta(a|s) > 0$ for all actions a in each state s , so that if a state is reachable under ρ , it will also be reachable using π_θ , and hence the zero derivative conditions apply at each reachable state. For $A^{\pi_\theta} + v$ to minimize L^θ , we would like $v^\top \nabla_\theta \log \pi_\theta(a|s) = 0$ for all s, a so that $v_{s,a}$ is independent of the action and can be written as a constant c_s for each s by the above equality. Hence, the minimizer of $L^\theta(w)$ is determined up to a state-dependent offset, and

$$F_\rho(\theta)^\dagger \nabla_\theta V^{\pi_\theta}(\rho) = \frac{A^{\pi_\theta}}{1 - \gamma} + v,$$

where $v_{s,a} = c_s$ for some $c_s \in \mathbb{R}$ for each state s and action a . Finally, we observe that this yields the updates

$$\theta^{(t+1)} = \theta^{(t)} + \frac{\eta}{1 - \gamma} A^{(t)} + \eta v \quad \text{and} \quad \pi^{(t+1)}(a|s) = \pi^{(t)}(a|s) \frac{\exp(\eta A^{(t)}(s, a)/(1 - \gamma) + \eta c_s)}{Z_t(s)}.$$

Owing to the normalization factor $Z_t(s)$, the state dependent offset c_s cancels in the updates for π , so that resulting policy is invariant to the specific choice of c_s . Hence, we pick $c_s \equiv 0$, which yields the statement of the lemma. \blacksquare

D Smoothness Proofs

Various convergence guarantees we show leverage results from smooth, non-convex optimization. In this section, we collect the various results on smoothness of policies and value functions in the different parameterizations which are needed in our analysis.

Define the Hadamard product of two vectors:

$$[x \odot y]_i = x_i y_i$$

Define $\text{diag}(x)$ for a column vector x as the diagonal matrix with diagonal as x .

Lemma D.1 (Smoothness of F (see Equation 38)). *Fix a state s . Let $\theta_s \in \mathbb{R}^{|\mathcal{A}|}$ be the column vector of parameters for state s . Let $\pi_\theta(\cdot|s)$ be the corresponding vector of action probabilities given by the softmax parameterization. For some fixed vector $c \in \mathbb{R}^{|\mathcal{A}|}$, define:*

$$F(\theta) := \pi_\theta(\cdot|s) \cdot c = \sum_a \pi_\theta(a|s) c_a.$$

Then

$$\|\nabla_{\theta_s} F(\theta_s) - \nabla_{\theta_s} F(\theta'_s)\|_2 \leq \beta \|\theta_s - \theta'_s\|_2$$

where

$$\beta = 5\|c\|_\infty.$$

Proof: For notational convenience, we do not explicitly state the s dependence. For the softmax parameterization, we have that

$$\nabla_\theta \pi_\theta = \text{diag}(\pi_\theta) - \pi_\theta \pi_\theta^\top.$$

We can then write (as $\nabla_\theta \pi_\theta$ is symmetric),

$$\nabla_\theta(\pi_\theta \cdot c) = (\text{diag}(\pi_\theta) - \pi_\theta \pi_\theta^\top)c = \pi_\theta \odot c - (\pi_\theta \cdot c)\pi_\theta \quad (51)$$

and therefore

$$\nabla_\theta^2(\pi_\theta \cdot c) = \nabla_\theta(\pi_\theta \odot c - (\pi_\theta \cdot c)\pi_\theta).$$

For the first term, we get

$$\nabla_\theta(\pi_\theta \odot c) = \text{diag}(\pi_\theta \odot c) - \pi_\theta(\pi_\theta \odot c)^\top,$$

and the second term, we can decompose by chain rule

$$\nabla_\theta((\pi_\theta \cdot c)\pi_\theta) = (\pi_\theta \cdot c)\nabla_\theta \pi_\theta + (\nabla_\theta(\pi_\theta \cdot c))\pi_\theta^\top$$

Substituting these back, we get

$$\nabla_\theta^2(\pi_\theta \cdot c) = \text{diag}(\pi_\theta \odot c) - \pi_\theta(\pi_\theta \odot c)^\top - (\pi_\theta \cdot c)\nabla_\theta \pi_\theta - (\nabla_\theta(\pi_\theta \cdot c))\pi_\theta^\top. \quad (52)$$

Note that

$$\begin{aligned} \max(\|\text{diag}(\pi_\theta \odot c)\|_2, \|\pi_\theta \odot c\|_2, |\pi_\theta \cdot c|) &\leq \|c\|_\infty \\ \|\nabla_\theta \pi_\theta\|_2 = \|\text{diag}(\pi_\theta) - \pi_\theta \pi_\theta^\top\|_2 &\leq 1 \\ \|\nabla_\theta(\pi_\theta \cdot c)\|_2 &\leq \|\pi_\theta \odot c\|_2 + \|(\pi_\theta \cdot c)\pi_\theta\|_2 \leq 2\|c\|_\infty, \end{aligned}$$

which gives

$$\|\nabla_\theta^2(\pi_\theta \cdot c)\|_2 \leq 5\|c\|_\infty. \quad \blacksquare$$

Before we prove the smoothness results for $\nabla_\pi V^\pi(s_0)$ and $\nabla_\theta V^{\pi_\theta}(s_0)$, we prove the following helpful lemma. This lemma is general and not specific to the direct or softmax policy parameterizations.

Lemma D.2. Let $\pi_\alpha := \pi_{\theta+\alpha u}$ and let $\tilde{V}(\alpha)$ be the corresponding value at a fixed state s_0 , i.e.

$$\tilde{V}(\alpha) := V^{\pi_\alpha}(s_0).$$

Assume that

$$\sum_{a \in \mathcal{A}} \left| \frac{d\pi_\alpha(a|s_0)}{d\alpha} \Big|_{\alpha=0} \right| \leq C_1, \quad \sum_{a \in \mathcal{A}} \left| \frac{d^2\pi_\alpha(a|s_0)}{(d\alpha)^2} \Big|_{\alpha=0} \right| \leq C_2$$

Then

$$\max_{\|u\|_2=1} \left| \frac{d^2\tilde{V}(\alpha)}{(d\alpha)^2} \Big|_{\alpha=0} \right| \leq \frac{C_2}{(1-\gamma)^2} + \frac{2\gamma C_1^2}{(1-\gamma)^3}.$$

Proof: Consider a unit vector u and let $\tilde{P}(\alpha)$ be the state-action transition matrix under π , i.e.

$$[\tilde{P}(\alpha)]_{(s,a) \rightarrow (s',a')} = \pi_\alpha(a'|s')P(s'|s,a).$$

We can differentiate $\tilde{P}(\alpha)$ w.r.t α to get

$$\left[\frac{d\tilde{P}(\alpha)}{d\alpha} \Big|_{\alpha=0} \right]_{(s,a) \rightarrow (s',a')} = \frac{d\pi_\alpha(a'|s')}{d\alpha} \Big|_{\alpha=0} P(s'|s,a).$$

For an arbitrary vector x ,

$$\left[\frac{d\tilde{P}(\alpha)}{d\alpha} \Big|_{\alpha=0} x \right]_{s,a} = \sum_{a',s'} \frac{d\pi_\alpha(a'|s')}{d\alpha} \Big|_{\alpha=0} P(s'|s,a)x_{a',s'}$$

and therefore

$$\begin{aligned} \max_{\|u\|_2=1} \left| \left[\frac{d\tilde{P}(\alpha)}{d\alpha} \Big|_{\alpha=0} x \right]_{s,a} \right| &= \max_{\|u\|_2=1} \left| \sum_{a',s'} \frac{d\pi_\alpha(a'|s')}{d\alpha} \Big|_{\alpha=0} P(s'|s,a)x_{a',s'} \right| \\ &\leq \sum_{a',s'} \left| \frac{d\pi_\alpha(a'|s')}{d\alpha} \Big|_{\alpha=0} \right| P(s'|s,a)|x_{a',s'}| \\ &\leq \sum_{s'} P(s'|s,a)\|x\|_\infty \sum_{a'} \left| \frac{d\pi_\alpha(a'|s')}{d\alpha} \Big|_{\alpha=0} \right| \\ &\leq \sum_{s'} P(s'|s,a)\|x\|_\infty C_1 \\ &\leq C_1\|x\|_\infty. \end{aligned}$$

By definition of ℓ_∞ norm,

$$\max_{\|u\|_2=1} \left\| \frac{d\tilde{P}(\alpha)}{d\alpha} x \right\|_\infty \leq C_1 \|x\|_\infty$$

Similarly, differentiating $\tilde{P}(\alpha)$ twice w.r.t. α , we get

$$\left[\frac{d^2\tilde{P}(\alpha)}{(d\alpha)^2} \Big|_{\alpha=0} \right]_{(s,a) \rightarrow (s',a')} = \frac{d^2\pi_\alpha(a'|s')}{(d\alpha)^2} \Big|_{\alpha=0} P(s'|s,a).$$

An identical argument leads to that, for arbitrary x ,

$$\max_{\|u\|_2=1} \left\| \frac{d^2 \tilde{P}(\alpha)}{(d\alpha)^2} \Big|_{\alpha=0} x \right\|_{\infty} \leq C_2 \|x\|_{\infty}$$

Let $Q^\alpha(s_0, a_0)$ be the corresponding Q -function for policy π_α at state s_0 and action a_0 . Observe that $Q^\alpha(s_0, a_0)$ can be written as:

$$Q^\alpha(s_0, a_0) = e_{(s_0, a_0)}^\top (\mathbf{I} - \gamma \tilde{P}(\alpha))^{-1} r = e_{(s_0, a_0)}^\top M(\alpha) r$$

where $M(\alpha) := (\mathbf{I} - \gamma \tilde{P}(\alpha))^{-1}$ and differentiating twice w.r.t α gives:

$$\begin{aligned} \frac{dQ^\alpha(s_0, a)}{d\alpha} &= \gamma e_{(s_0, a)}^\top M(\alpha) \frac{d\tilde{P}(\alpha)}{d\alpha} M(\alpha) r, \\ \frac{d^2 Q^\alpha(s_0, a_0)}{(d\alpha)^2} &= 2\gamma^2 e_{(s_0, a_0)}^\top M(\alpha) \frac{d\tilde{P}(\alpha)}{d\alpha} M(\alpha) \frac{d\tilde{P}(\alpha)}{d\alpha} M(\alpha) r \\ &\quad + \gamma e_{(s_0, a_0)}^\top M(\alpha) \frac{d^2 \tilde{P}(\alpha)}{(d\alpha)^2} M(\alpha) r. \end{aligned}$$

By using power series expansion of matrix inverse, we can write $M(\alpha)$ as:

$$M(\alpha) = (\mathbf{I} - \gamma \tilde{P}(\alpha))^{-1} = \sum_{n=0}^{\infty} \gamma^n \tilde{P}(\alpha)^n$$

which implies that $M(\alpha) \geq 0$ (componentwise) and $M(\alpha)\mathbf{1} = \frac{1}{1-\gamma}\mathbf{1}$, i.e. each row of $M(\alpha)$ is positive and sums to $1/(1-\gamma)$. This implies:

$$\max_{\|u\|_2=1} \|M(\alpha)x\|_{\infty} \leq \frac{1}{1-\gamma} \|x\|_{\infty}$$

This gives using expression for $\frac{d^2 Q^\alpha(s_0, a_0)}{(d\alpha)^2}$ and $\frac{dQ^\alpha(s_0, a)}{d\alpha}$,

$$\begin{aligned} \max_{\|u\|_2=1} \left| \frac{d^2 Q^\alpha(s_0, a_0)}{(d\alpha)^2} \Big|_{\alpha=0} \right| &\leq 2\gamma^2 \left\| M(\alpha) \frac{d\tilde{P}(\alpha)}{d\alpha} M(\alpha) \frac{d\tilde{P}(\alpha)}{d\alpha} M(\alpha) r \right\|_{\infty} \\ &\quad + \gamma \left\| M(\alpha) \frac{d^2 \tilde{P}(\alpha)}{(d\alpha)^2} M(\alpha) r \right\|_{\infty} \\ &\leq \frac{2\gamma^2 C_1^2}{(1-\gamma)^3} + \frac{\gamma C_2}{(1-\gamma)^2} \\ \max_{\|u\|_2=1} \left| \frac{dQ^\alpha(s_0, a)}{d\alpha} \Big|_{\alpha=0} \right| &\leq \left\| \gamma M(\alpha) \frac{d\tilde{P}(\alpha)}{d\alpha} M(\alpha) r \right\|_{\infty} \\ &\leq \frac{\gamma C_1}{(1-\gamma)^2} \end{aligned}$$

Consider the identity:

$$\tilde{V}(\alpha) = \sum_a \pi_\alpha(a|s_0) Q^\alpha(s_0, a),$$

By differentiating $\tilde{V}(\alpha)$ twice w.r.t α , we get

$$\frac{d^2 \tilde{V}(\alpha)}{(d\alpha)^2} = \sum_a \frac{d^2 \pi_\alpha(a|s_0)}{(d\alpha)^2} Q^\alpha(s_0, a) + 2 \sum_a \frac{d\pi_\alpha(a|s_0)}{d\alpha} \frac{dQ^\alpha(s_0, a)}{d\alpha} + \sum_a \pi_\alpha(a|s_0) \frac{d^2 Q^\alpha(s_0, a)}{(d\alpha)^2}.$$

Hence,

$$\begin{aligned} \max_{\|u\|_2=1} \left| \frac{d^2 \tilde{V}(\alpha)}{(d\alpha)^2} \right| &\leq \frac{C_2}{1-\gamma} + \frac{2\gamma C_1^2}{(1-\gamma)^2} + \frac{2\gamma^2 C_1^2}{(1-\gamma)^3} + \frac{\gamma C_2}{(1-\gamma)^2} \\ &= \frac{C_2}{(1-\gamma)^2} + \frac{2\gamma C_1^2}{(1-\gamma)^3}, \end{aligned}$$

which completes the proof. \blacksquare

Using this lemma, we now establish smoothness for: the value functions under the direct policy parameterization and the log barrier regularized objective 12 for the softmax parameterization.

Lemma D.3 (Smoothness for direct parameterization). *For all starting states s_0 ,*

$$\left\| \nabla_\pi V^\pi(s_0) - \nabla_\pi V^{\pi'}(s_0) \right\|_2 \leq \frac{2\gamma|\mathcal{A}|}{(1-\gamma)^3} \|\pi - \pi'\|_2$$

Proof: By differentiating π_α w.r.t α gives

$$\sum_{a \in \mathcal{A}} \left| \frac{d\pi_\alpha(a|s_0)}{d\alpha} \right| \leq \sum_{a \in \mathcal{A}} |u_{a,s}| \leq \sqrt{|\mathcal{A}|}$$

and differentiating again w.r.t α gives

$$\sum_{a \in \mathcal{A}} \left| \frac{d^2 \pi_\alpha(a|s_0)}{(d\alpha)^2} \right| = 0$$

Using this with Lemma D.2 with $C_1 = \sqrt{|\mathcal{A}|}$ and $C_2 = 0$, we get

$$\max_{\|u\|_2=1} \left| \frac{d^2 \tilde{V}(\alpha)}{(d\alpha)^2} \right|_{\alpha=0} \leq \frac{C_2}{(1-\gamma)^2} + \frac{2\gamma C_1^2}{(1-\gamma)^3} \leq \frac{2\gamma|\mathcal{A}|}{(1-\gamma)^3}$$

which completes the proof. \blacksquare

We now present a smoothness result for the entropy regularized policy optimization problem which we study for the softmax parameterization.

Lemma D.4 (Smoothness for log barrier regularized softmax). *For the softmax parameterization and*

$$L_\lambda(\theta) = V^{\pi_\theta}(\mu) + \frac{\lambda}{|\mathcal{S}||\mathcal{A}|} \sum_{s,a} \log \pi_\theta(a|s),$$

we have that

$$\|\nabla_\theta L_\lambda(\theta) - \nabla_\theta L_\lambda(\theta')\|_2 \leq \beta_\lambda \|\theta - \theta'\|_2$$

where

$$\beta_\lambda = \frac{8}{(1-\gamma)^3} + \frac{2\lambda}{|\mathcal{S}|}$$

Proof: Let us first bound the smoothness of $V^{\pi_\theta}(\mu)$. Consider a unit vector u . Let $\theta_s \in \mathbb{R}^{|\mathcal{A}|}$ denote the parameters associated with a given state s . We have:

$$\nabla_{\theta_s} \pi_\theta(a|s) = \pi_\theta(a|s) \left(e_a - \pi(\cdot|s) \right)$$

and

$$\nabla_{\theta_s}^2 \pi_\theta(a|s) = \pi_\theta(a|s) \left(e_a e_a^\top - e_a \pi(\cdot|s)^\top - \pi(\cdot|s) e_a^\top + 2\pi(\cdot|s) \pi(\cdot|s)^\top - \text{diag}(\pi(\cdot|s)) \right),$$

where e_a is a standard basis vector and $\pi(\cdot|s)$ is a vector of probabilities. We also have by differentiating $\pi_\alpha(a|s)$ once w.r.t α ,

$$\begin{aligned} \sum_{a \in \mathcal{A}} \left| \frac{d\pi_\alpha(a|s)}{d\alpha} \Big|_{\alpha=0} \right| &\leq \sum_{a \in \mathcal{A}} \left| u^\top \nabla_{\theta+\alpha u} \pi_\alpha(a|s) \Big|_{\alpha=0} \right| \\ &\leq \sum_{a \in \mathcal{A}} \pi_\theta(a|s) |u_s^\top e_a - u_s^\top \pi(\cdot|s)| \\ &\leq \max_{a \in \mathcal{A}} \left(|u_s^\top e_a| + |u_s^\top \pi(\cdot|s)| \right) \leq 2 \end{aligned}$$

Similarly, differentiating once again w.r.t. α , we get

$$\begin{aligned} \sum_{a \in \mathcal{A}} \left| \frac{d^2 \pi_\alpha(a|s)}{(d\alpha)^2} \Big|_{\alpha=0} \right| &\leq \sum_{a \in \mathcal{A}} \left| u^\top \nabla_{\theta+\alpha u}^2 \pi_\alpha(a|s) \Big|_{\alpha=0} \right| u \\ &\leq \max_{a \in \mathcal{A}} \left(|u_s^\top e_a e_a^\top u_s| + |u_s^\top e_a \pi(\cdot|s)^\top u_s| + |u_s^\top \pi(\cdot|s) e_a^\top u_s| \right. \\ &\quad \left. + 2 |u_s^\top \pi(\cdot|s) \pi(\cdot|s)^\top u_s| + |u_s^\top \text{diag}(\pi(\cdot|s)) u_s| \right) \\ &\leq 6 \end{aligned}$$

Using this with Lemma D.2 for $C_1 = 2$ and $C_2 = 6$, we get

$$\max_{\|u\|_2=1} \left| \frac{d^2 \tilde{V}(\alpha)}{(d\alpha)^2} \Big|_{\alpha=0} \right| \leq \frac{C_2}{(1-\gamma)^2} + \frac{2\gamma C_1^2}{(1-\gamma)^3} \leq \frac{6}{(1-\gamma)^2} + \frac{8\gamma}{(1-\gamma)^3} \leq \frac{8}{(1-\gamma)^3}$$

or equivalently for all starting states s and hence for all starting state distributions μ ,

$$\|\nabla_{\theta} V^{\pi_{\theta}}(\mu) - \nabla_{\theta} V^{\pi_{\theta'}}(\mu)\|_2 \leq \beta \|\theta - \theta'\|_2 \quad (53)$$

where $\beta = \frac{8}{(1-\gamma)^3}$.

Now let us bound the smoothness of the regularizer $\frac{\lambda}{|S|}R(\theta)$, where

$$R(\theta) := \frac{1}{|\mathcal{A}|} \sum_{s,a} \log \pi_{\theta}(a|s)$$

We have

$$\frac{\partial R(\theta)}{\partial \theta_{s,a}} = \frac{1}{|\mathcal{A}|} - \pi_{\theta}(a|s).$$

Equivalently,

$$\nabla_{\theta_s} R(\theta) = \frac{1}{|\mathcal{A}|} \mathbf{1} - \pi_{\theta}(\cdot|s).$$

Hence,

$$\nabla_{\theta_s}^2 R(\theta) = -\text{diag}(\pi_{\theta}(\cdot|s)) + \pi_{\theta}(\cdot|s)\pi_{\theta}(\cdot|s)^{\top}.$$

For any vector u_s ,

$$|u_s^{\top} \nabla_{\theta_s}^2 R(\theta) u_s| = |u_s^{\top} \text{diag}(\pi_{\theta}(\cdot|s)) u_s - (u_s \cdot \pi_{\theta}(\cdot|s))^2| \leq 2\|u_s\|_{\infty}^2.$$

Since $\nabla_{\theta_s} \nabla_{\theta_{s'}} R(\theta) = 0$ for $s \neq s'$,

$$|u^{\top} \nabla_{\theta}^2 R(\theta) u| = \left| \sum_s u_s^{\top} \nabla_{\theta_s}^2 R(\theta) u_s \right| \leq 2 \sum_s \|u_s\|_{\infty}^2 \leq 2\|u\|_2^2.$$

Thus R is 2-smooth and $\frac{\lambda}{|S|}R$ is $\frac{2\lambda}{|S|}$ -smooth, which completes the proof. ■

E Standard Optimization Results

In this section, we present the standard optimization results from Ghadimi and Lan [2016], Beck [2017] used in our proofs. We consider solving the following problem

$$\min_{x \in C} \{f(x)\} \quad (54)$$

with C being a nonempty closed and convex set. We assume the following

Assumption E.1. $f : \mathbb{R}^d \rightarrow (-\infty, \infty)$ is proper and closed, $\text{dom}(f)$ is convex and f is β smooth over $\text{int}(\text{dom}(f))$.

Throughout the section, we will denote the optimal f value by $f(x^*)$.

Definition E.1 (Gradient Mapping). We define the gradient mapping $G^\eta(x)$ as

$$G^\eta(x) := \frac{1}{\eta} (x - P_C(x - \eta \nabla f(x))) \quad (55)$$

where P_C is the projection onto C .

Note that when $C = \mathbb{R}^d$, the gradient mapping $G^\eta(x) = \nabla f(x)$.

Theorem E.1 (Theorem 10.15 Beck [2017]). *Suppose that Assumption E.1 holds and let $\{x_k\}_{k \geq 0}$ be the sequence generated by the gradient descent algorithm for solving the problem (54) with the stepsize $\eta = 1/\beta$. Then,*

1. *The sequence $\{F(x_t)\}_{t \geq 0}$ is non-increasing.*
2. *$G^\eta(x_t) \rightarrow 0$ as $t \rightarrow \infty$*
3. *$\min_{t=0,1,\dots,T-1} \|G^\eta(x_t)\| \leq \frac{\sqrt{2\beta f(x_0) - f(x^*)}}{\sqrt{T}}$*

Theorem E.2 (Lemma 3 Ghadimi and Lan [2016]). *Suppose that Assumption E.1 holds. Let $x^+ = x - \eta G^\eta(x)$. Then,*

$$\nabla f(x^+) \in N_C(x^+) + \epsilon(\eta\beta + 1)B_2,$$

where B_2 is the unit ℓ_2 ball, and N_C is the normal cone of the set C .

We now consider the stochastic projected gradient descent algorithm where at each time step t , we update x_t by sampling a random v_t such that

$$x_{t+1} = P_C(x_t - \eta v_t), \quad \text{where } \mathbb{E}[v_t | x_t] = \nabla f(x_t) \quad (56)$$

Theorem E.3 (Theorem 14.8 and Lemma 14.9 Shalev-Shwartz and Ben-David [2014]). *Assume $C = \{x : \|x\| \leq B\}$, for some $B > 0$. Let f be a convex function and let $x^* \in \text{argmin}_{x: \|x\| \leq B} f(x)$. Assume also that for all t , $\|v_t\| \leq \rho$, and that stochastic projected gradient descent is run for N iterations with $\eta = \sqrt{\frac{B^2}{\rho^2 N}}$. Then,*

$$\mathbb{E} \left[f \left(\frac{1}{N} \sum_{t=1}^N x_t \right) \right] - f(x^*) \leq \frac{B\rho}{\sqrt{N}} \quad (57)$$