

Efficient Planning in Large MDPs with Weak Linear Function Approximation

Roshan Shariff

University of Alberta & Amii
roshan.shariff@ualberta.ca

Csaba Szepesvári

DeepMind & University of Alberta & Amii
szepesva@ualberta.ca

Abstract

Large-scale Markov decision processes (MDPs) require planning algorithms with runtime independent of the number of states of the MDP. We consider the planning problem in MDPs using linear value function approximation with only weak requirements: low approximation error for the optimal value function, and a small set of “core” states whose features span those of other states. In particular, we make no assumptions about the representability of policies or value functions of non-optimal policies. Our algorithm produces almost-optimal actions for any state using a generative oracle (simulator) for the MDP, while its computation time scales polynomially with the number of features, core states, and actions and the effective horizon.

1 Introduction

Markov decision processes (MDPs) are commonly used to model sequential decision making under uncertainty and have a wide range of applications [see 43, 33, 7, for example]. We consider planning in large-scale, expected discounted total reward MDPs. Computing an optimal policy in the discounted setting is known to require “reading” all states at least once [6]. As the state space for most interesting applications is intractably large if not infinite (“Bellman’s curse of dimensionality”), it is common to consider restrictions to the problem that can allow efficient calculation of near-optimal actions. One such relaxation is *online planning* — the MDP can be accessed through a simulator and we ask only for a good action at a given state [25, 29]. While in this problem the complexity of computing a “good action” can be independent of the number of states, the complexity is exponential in the planning horizon [25]. An alternative idea, which can be traced back to at least the work of Schweitzer and Seidmann [35], is to assume that one has access to a *feature representation* (that is, a vector of features for each state) and the planner needs to work well for those MDPs where the *optimal value function* of the MDP can be uniformly well approximated over all states by some appropriate weighted combination of the features. Since an accurate approximation of the optimal value function is known to be sufficient to generate near-optimal behavior,¹ the problem simplifies to producing a good estimate of the unknown feature weights with a computation cost that is independent of the number of states.

In this paper we consider the intersection of these two problem formulations. More precisely, our goal is to construct planning algorithms that produce an action for any given input state, using black-box access to the MDP through a simulator which takes a state and an action as input and produces a random next state and immediate reward. The planner can also access the feature representation of any state as a d -dimensional feature vector. Assume that the optimal value function of the MDP can be uniformly well approximated — to an accuracy of $\varepsilon_{\text{approx}}$ — as a linear combination of the features with fixed, but unknown, coefficients. Our goal is a randomized planning algorithm that interacts with

¹See Proposition 1; or, for example, Szepesvári [39, Lemma 5.17], Kearns et al. [25, Lemma 5], Kallenberg [24, Theorem 3.7].

the simulator $\text{poly}(1/\varepsilon, d, H, A, \dots)$ times to produce an action, such that following this action in every state results in an $O(\varepsilon + c\varepsilon_{\text{approx}})$ -optimal policy. Here $H = 1/(1 - \gamma)$ is the effective planning horizon for the discount factor $0 \leq \gamma < 1$, which is used in the definition of the values of policies; A is the number of actions; and $c > 0$ is an error inflation factor that may depend on γ , d , and A .

We call the features “weak” as we only require the optimal value function to be accurately representable by their linear combinations, in contrast to “strong” features that can accurately represent the value functions of *all* policies; in this latter case the problem of efficient planning in the presence of a simulator is known to have a solution, both in the episodic and discounted settings [44, 28]. As pointed out by Du et al. [12], with only weak features, the problem of efficient planning has not yet been solved.

Our Contributions We design a randomized algorithm that positively answers the challenge posed above under one extra assumption — that the feature vectors of all states lie within the convex hull of the feature vectors of a few selected “core states” that the algorithm is given. In particular, we show that the query-complexity and runtime of our algorithm is polynomial in the relevant quantities and the number of core states, providing a partial positive answer to the previously open problem of efficient planning in the presence of weak features.

To achieve our result, we start from the *approximate linear programming* (ALP) approach where the value function is approximated using the feature vectors. Following Lakshminarayanan et al. [27], we construct a *relaxed* ALP that drops all constraints except at the core states. In their work, Lakshminarayanan et al. gave bounds on the error of the value function that is obtained from solving this relaxed ALP. The authors also suggested a way to turn this error bound into an efficient planning method, though without a detailed analysis. The main contribution of the present work is to fill this gap, in addition to simplifying, strengthening and streamlining the earlier results. In particular, we propose using a randomized saddle-point solver that substantially reduces the computational requirements compared to the procedure hinted at by [27].

Paper Organization The rest of the paper is organized as follows: [Sections 1.1](#) and [1.2](#) give background on MDPs and introduce the linear programming (LP) approach to planning. [Section 2](#) formally defines the problem. Then, in [Section 3](#), we present the linear program that we start with and give our first results, bounding the value loss of the policy that can be read out from optimal solutions of the linear program. [Section 4](#) gives the efficient algorithm to solve the linear program and our main result. [Section 5](#) discusses related work. The paper is concluded in [Section 6](#). The proofs of the results are moved to [Appendix A](#).

Notation The set of real numbers is denoted by $\mathbb{R}$, whereas $\mathbb{R}_+ = [0, \infty)$. $\mathbb{R}^d$ denotes the vectors with d dimensions, while the $m \times n$ matrices are $\mathbb{R}^{m \times n}$. We use bold letters for vectors ($\mathbf{r}$) and bold capitals for matrices ($\mathbf{P}$); their elements are written as r_i and $P_{i,j}$ and matrix rows are $\mathbf{P}_i$. For vectors of identical dimension, $\mathbf{x} \leq \mathbf{y}$ means element-wise comparison: $x_i \leq y_i$ for each index i . The standard basis vector $\mathbf{e}_i$ has $e_{i,i} = 1$ and $e_{i,j} = 0$ for $i \neq j$, and the constant 0 or 1 vector is denoted by $\mathbf{0}, \mathbf{1} \in \mathbb{R}^d$; their dimension depends on the context. All vectors are considered column vectors by default. The probability simplex over any finite set $\mathcal{A}$ is denoted $\Delta_{\mathcal{A}} := \{\mathbf{p} \in \mathbb{R}_+^{|\mathcal{A}|} \mid \|\mathbf{p}\|_1 = 1\}$. For a finite set $\mathcal{S}$ with cardinality $S = |\mathcal{S}|$, we will think of functions $v : \mathcal{S} \rightarrow \mathbb{R}$ or $\varphi : \mathcal{S} \rightarrow \mathbb{R}^d$ as vectors or matrices, respectively, and use both notations: for example, $\mathbf{v} \in \mathbb{R}^S$, v_s , or $v(s)$; and $\mathbf{\Phi} \in \mathbb{R}^{S \times d}$, φ_s , or $\varphi(s)$ where $s \in \mathcal{S}$. When the domain takes the form $\mathcal{S} \times \mathcal{A}$ with respective cardinalities S and A , we use intuitive double indices of the form sa , e.g., with $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ we index components of $\mathbf{r}$ using the notation r_{sa} (i.e., $r_{sa} = r(s, a)$). In this case we also write $\mathbf{r} \in \mathbb{R}^{SA}$.

For convenience, an [Index of Notation](#) section is included in the [Appendix](#).

1.1 Background

A (finite, discounted) Markov Decision Process (MDP) is defined by the entities $(\mathcal{S}, \mathcal{A}, \mathbf{P}, \mathbf{r}, \gamma)$ where $\mathcal{S}$ and $\mathcal{A}$ are finite sets of states and actions, respectively. Without loss of generality we let $\mathcal{S} = [S]$ and $\mathcal{A} = [A]$, using the notation $[n] = \{1, \dots, n\}$ for integers $n > 0$. When the process is in state $s \in \mathcal{S}$ and action $a \in \mathcal{A}$ is chosen, a random reward is received with expectation $r_{sa} \in \mathbb{R}$ and the process transitions to a new state $s' \in \mathcal{S}$ with probability $P_{sa,s'}$. For convenience, we arrange the transition probabilities into a matrix $\mathbf{P} \in \mathbb{R}^{SA \times S}$ and the rewards into a vector $\mathbf{r} \in \mathbb{R}^{SA}$. Thus, $\mathbf{P}$ is a

row-stochastic matrix — each row $\mathbf{P}_{sa}$ for a state s and action a is a valid probability distribution for the next state.

For our purposes it will be sufficient to consider stationary policies, which we will just call *policies*. A policy $\pi : \mathcal{S} \rightarrow \Delta_{\mathcal{A}}$ is a function from states to probability distributions over actions — we use $\pi(a|s)$ to denote the probability assigned by π to action a in state s . Following a policy means that upon visiting state s , an action $a \sim \pi(s)$ is chosen at random. This gives rise to an infinite sequence of states, actions, and corresponding rewards. The *value* $v_\pi(s)$ of a policy for a process starting at state $s \in \mathcal{S}$ is defined as the total expected γ -discounted sum of the rewards incurred:

$$\mathbf{v}_\pi = \sum_{t=0}^{\infty} (\gamma \mathbf{P}_\pi)^t \mathbf{r}_\pi, \quad \text{where } [\mathbf{r}_\pi]_s = \sum_{a \in \mathcal{A}} \pi(a|s) r_{sa}, \quad \text{and } [\mathbf{P}_\pi]_{s,s'} = \sum_{a \in \mathcal{A}} \pi(a|s) P_{sa,s'}; \quad (1)$$

$\mathbf{r}_\pi \in \mathbb{R}^S$ is the expected reward and $\mathbf{P}_\pi \in \mathbb{R}^{S \times S}$ is the state transition matrix. A policy π^* is called *optimal* if $\mathbf{v}_{\pi^*} \geq \mathbf{v}_\pi$ for every policy π , where the inequality is component-wise. Every MDP has an optimal policy, and all optimal policies have the same value function $\mathbf{v}^*$, the *optimal value function*. Further, there always exist *deterministic* optimal policies, which concentrate all their probability on a single action for each state. We will also need $\mathbf{q}^* \in \mathbb{R}^{SA}$, which is defined via $\mathbf{q}^* = \mathbf{r} + \gamma \mathbf{P} \mathbf{v}^*$.

1.2 Linear Programming

Our approach to the MDP planning problem is based on the standard linear programming (LP) formulation; for details, see Puterman [32, §6.9], who swaps the primal and dual problems:

$$v^*(s_0) = \min \left\{ \mathbf{e}_{s_0}^\top \mathbf{v} \mid \mathbf{v} \in \mathbb{R}^S, \quad \mathbf{r} + (\gamma \mathbf{P} - \mathbf{E}) \mathbf{v} \leq \mathbf{0} \right\} \quad (\text{Primal LP})$$

$$= \max \left\{ \boldsymbol{\mu}^\top \mathbf{r} \mid \boldsymbol{\mu} \in \mathbb{R}_+^{SA}, \quad \mathbf{e}_{s_0} + \boldsymbol{\mu}^\top (\gamma \mathbf{P} - \mathbf{E}) = \mathbf{0} \right\}. \quad (\text{Dual LP})$$

The matrix $\mathbf{E} : \mathbb{R}^{SA \times S}$ has elements $E_{sa,s} = 1$ and $E_{sa,s'} = 0$ for $s \neq s'$. It maps vectors from $\mathbb{R}^S$ to $\mathbb{R}^{SA}$ by duplicating their elements over all actions: $[\mathbf{E} \mathbf{v}]_{sa} = v_s$ for all s, a . Both the primal and dual optimization problems above have the same optimal value: the optimal value of state s_0 . The dual variables $\boldsymbol{\mu}$ are *discounted state-action occupancy measures* of a policy π starting at state s_0 :

$$\boldsymbol{\mu}^\top = \boldsymbol{\rho}_\pi^\top \sum_{t=0}^{\infty} (\gamma \tilde{\mathbf{P}}_\pi)^t, \quad \text{where } [\boldsymbol{\rho}_\pi]_{sa} = \pi(a|s_0) e_{s,s_0}, \quad \text{and } [\tilde{\mathbf{P}}_\pi]_{sa,s'a'} = \pi(a'|s') [\mathbf{P}_\pi]_{sa,s'}; \quad (2)$$

$\boldsymbol{\rho}_\pi \in \Delta_{S \times \mathcal{A}}$ is the initial distribution over state-action pairs and $\tilde{\mathbf{P}}_\pi \in \mathbb{R}^{SA \times SA}$ is the state-action transition probability matrix; the state transition matrix $\mathbf{P}_\pi$ is defined in (1). The constraint $\mathbf{e}_{s_0} + \boldsymbol{\mu}^\top (\gamma \mathbf{P} - \mathbf{E}) = \mathbf{0}$ enforces that $\boldsymbol{\mu}$ has this form, and is therefore generated by some stochastic policy. Thus the dual problem can be seen as a linear formulation of policy optimization, where policies are represented by their occupancy measures and the objective is to maximize expected discounted reward — in particular, the policy corresponding to any feasible $\boldsymbol{\mu}$ can be obtained by conditioning on the state: $\pi_\mu(a|s) = \mu_{sa} / \sum_{a'} \mu_{sa'}$ for any state with non-zero occupancy measure.

The *approximate linear program* (ALP) of Schweitzer and Seidmann [35] is obtained from (Primal LP) by restricting $\mathbf{v}$ to lie in the span of a *feature matrix* $\boldsymbol{\Phi} \in \mathbb{R}^{S \times d}$ — $\mathbf{v} = \boldsymbol{\Phi} \boldsymbol{\theta}$ for $\boldsymbol{\theta} \in \mathbb{R}^d$. This reduces the number of variables from S to the feature dimension d , but is still intractable to solve because of the many constraints — one for each state-action pair. The *relaxed ALP* of Lakshminarayanan et al. [27] addresses this issue by keeping only a small number of constraints that are positive linear combinations of the original constraints; this is the foundation of our approach.

2 Problem Definition

In the *online MDP planning problem*, a randomized planner is given a state of the MDP $s_0 \in \mathcal{S}$ as input, and needs to return an action [e.g., 25, 16]. Letting $\pi(a|s_0)$ denote the probability that action a is returned for input s_0 , the planner’s *value loss* at state s is defined as $v^*(s) - v_\pi(s)$. The goal is to design planning algorithms with a small value loss for *every* state s regardless of the MDP.

For large MDPs, we want the computation time to be independent of the number of states S and depend polynomially on the number of actions and the “planning horizon” $H = 1/(1 - \gamma)$. To make this even remotely possible, we assume that the planner has access to a suitable feature map $\boldsymbol{\varphi} : \mathcal{S} \rightarrow \mathbb{R}^d$ (Assumption 1 below), is given a suitable set of “core states” (Assumption 2), and can

access a simulator of the MDP (Assumption 3).² Further, the planner is only required to perform well if the optimal value function lies uniformly close to the span of the features, which means that

$$\varepsilon_{\text{approx}} := \inf_{\theta \in \mathbb{R}^d} \max_{s \in \mathcal{S}} |v^*(s) - \varphi_s^\top \theta| \quad (3)$$

is small. To be precise, the planner’s value loss is allowed to degrade with $\varepsilon_{\text{approx}}$. Note that the class of MDPs where $\varepsilon_{\text{approx}}$ is small for a given feature map is a strict superset of those which are nearly linear up to the error $\varepsilon_{\text{approx}}$ in the sense of Jin et al. [19]. Hence, we call the features *weak* because we only require that $\varepsilon_{\text{approx}}$ as defined in (3) be small. Du et al. [12] posed the open problem of designing efficient online planning algorithms under this condition.

Assumptions. For the convenience of the reader, we now restate our assumptions in a concise form:

1. *Features:* The planner can access $\varphi(s) \in \mathbb{R}^d$ for any state $s \in \mathcal{S}$. Further, there is some $\eta \in \mathbb{R}^d$ such that $\varphi_s^\top \eta = 1$ for all s — this can be ensured easily by adding a “bias” feature that is always 1.
2. *Core States:* There is a set of core states $\mathcal{S}_* \subset \mathcal{S}$ (with $|\mathcal{S}_*| = m$) that are available to the algorithm, and the feature vector of every other state can be written as a positive linear combination of the core state features: $\Phi = \mathbf{Z}\Phi_*$ for some non-negative matrix $\mathbf{Z} \in \mathbb{R}_+^{S \times m}$, where $\Phi \in \mathbb{R}^{S \times d}$ and $\Phi_* \in \mathbb{R}^{m \times d}$ consist of the stacked feature vectors for all states and the core states, respectively.
3. *Simulator:* The planner can call a randomized function $\text{SIMULATE}(s, a)$ that returns $s' \sim \mathbf{P}_{sa}$ and a reward $\hat{r}$ with $\mathbb{E}[\hat{r}] = r_{sa}$. For simplicity, we assume $|\hat{r}| \leq 1$.

We do not need to explicitly assume that the feature vectors of all states lie within the convex hull of the core state features — that is a consequence of Assumptions 1 and 2 (specifically, $\mathbf{1} \in \text{span}(\Phi)$).³ Under Assumption 2, each of the core states is a “soft state aggregation” [36] that respects the feature representation. Zanette et al. [46] impose a similar requirement for core states, observing that “anchoring” the value function at states with “extreme” feature representations (i.e., on the boundary of the convex hull of feature vectors) results in the values of other states being *interpolated* (not *extrapolated*) from the values of the core states using their respective feature representations — this is sufficient to accurately deduce the values of all states when $\varepsilon_{\text{approx}}$ is small [44].

Without assuming extra structure, finding a set of core states requires checking the feature vectors of all states, which is intractable in large MDPs. It remains an interesting question what extra structure would make it possible to discover near-minimal core sets with an effort independent of the size of the MDP. We also note that for some feature maps the size of the minimal core set can be as large as the number of states S . Since the run time of our algorithm depends on the size of the core set, one should obviously avoid such feature maps. It remains an intriguing question whether requiring a small core set is necessary for efficient planning.

The algorithm we design can work with weaker assumptions — it only accesses $\varphi(s)$ for the planning state s_0 , the core states, and for the next state s' produced by the simulator: $(s', \hat{r}) \leftarrow \text{SIMULATE}(s, a)$, which it only queries with $s \in \{s_0\} \cup \mathcal{S}_*$. We also note in passing that the $\mathbf{Z}$ matrix of Assumption 2 is not used by the algorithm and need not be known; it need only exist. Finally, our results also hold if the core states are so-called “meta-states” — probability distributions over underlying MDP states; the simulator must then be able to sample states according to these probability distributions.

3 CoreLP — A Linear Program for Planning with Core States

Throughout this section, we will use $s_0 \in \mathcal{S}$ to refer to the current planning state — our goal is to output a random close-to-optimal action $a \sim \pi(s_0)$. Consider the following result, which follows immediately from the well-known “performance difference lemma” [23, Lemma 6.1]:

Proposition 1. Let π be an arbitrary policy. Then,

$$\max_{s \in \mathcal{S}} v^*(s) - v_\pi(s) \leq \frac{1}{1 - \gamma} \max_{s_0 \in \mathcal{S}} \mathbb{E}_{a \sim \pi(s_0)} [v^*(s_0) - q^*(s_0, a)].$$

²With no additional assumptions on the MDP, any online planning algorithm implementing an ε -suboptimal policy may need up to $\Omega((1/\varepsilon)^{H-1})$ simulator queries to find each action; this is exponential in the planning horizon H for any constant $\varepsilon > 0$ — see Kearns et al. [25, Theorem 2], noting that their H is different from ours.

³We have $\mathbf{Z}\mathbf{1} = \mathbf{Z}\Phi_*\eta = \Phi\eta = \mathbf{1}$ for some $\eta \in \mathbb{R}^d$, so the rows of $\mathbf{Z}$ must sum to one.

In light of this, we will design a randomized planning algorithm that guarantees $\mathbb{E}_{a \sim \pi(s_0)}[q^*(s_0, a)] \approx v^*(s_0)$ for any input state s_0 . Our approach stems from the relaxed linear program of Lakshminarayanan et al. [27] — more precisely, we start with the ALP but keep only the constraints corresponding to the actions of the core states and the current planning state (see Section 1.2). We then construct the corresponding dual LP, to which we add another constraint that allows us to “read out” an action distribution from the values of the dual variables.

Recall that $\mathcal{S}_* = \{s_1, \dots, s_m\}$ is the set of core states and define $\mathcal{S}_+$ as the sequence $(s_0, s_1, \dots, s_m)$. Note that s_0 is always the first state in $\mathcal{S}_+$ but may appear again if it is also a core state. For each of the $1 + m$ states in $\mathcal{S}_+$, we will select the A constraints in the ALP corresponding to the state-action pairs $(s_i, a) \in \mathcal{S}_+ \times \mathcal{A}$ — a total of $(1 + m)A$ constraints. Unlike the ALP, our linear program CORELP is based on the (Dual LP) of Section 1.2; the name refers to the CORE states and features that, respectively, CONSTRAIN and RELAX it.

Theorem 2 (CoreLP). *Suppose Assumptions 1 and 2 hold, $s_0 \in \mathcal{S}$, $\varphi_0 := \varphi_{s_0}$, $\mathbf{W} \in \{0, 1\}^{(1+m)A \times SA}$ has rows $[\mathbf{W}_{s_i a}]_{s_i \in \mathcal{S}_+, a \in \mathcal{A}} = \mathbf{e}_{s_i a}$, and $\Lambda := \{\lambda \in \mathbb{R}_+^{(1+m)A} \mid \sum_{a \in \mathcal{A}} \lambda_{s_0 a} = 1\}$. Define*

$$V^\dagger = \max \{ \lambda^\top \mathbf{W} \mathbf{r} \mid \lambda \in \Lambda, \varphi_0^\top + \lambda^\top \mathbf{W}(\gamma \mathbf{P} - \mathbf{E}) \Phi = \mathbf{0} \}. \quad (\text{CoreLP})$$

Let $\lambda^\dagger \in \Lambda$ be a maximizer of (CoreLP) and let $\pi^\dagger \in \Delta_{\mathcal{A}}$ be given by $\pi^\dagger(a) = \lambda_{s_0 a}^\dagger$. Then

$$|V^\dagger - v^*(s_0)| \leq \frac{10\gamma\epsilon_{\text{approx}}}{1 - \gamma}, \quad v^*(s_0) - \sum_{a \in \mathcal{A}} \pi^\dagger(a) q^*(s_0, a) \leq \frac{20\gamma\epsilon_{\text{approx}}}{1 - \gamma}.$$

This bound matches up to constant factors (and improves by a γ factor) the landmark result of de Farias and Van Roy [14] for the approximation error of the ALP (defined in Section 1.2). In other words, core states satisfying Assumption 2 lead to essentially no additional error in the solution of (CoreLP) compared to the ALP — this was pointed out by Lakshminarayanan et al. [27], whose result we improve upon in Theorem 5 (Appendix A.1). The theorem also implies that the linear program is both bounded and feasible, meaning its value is not $\pm\infty$; this is an important consideration when relaxing the ALP [4]. We present the detailed proof in Appendix A.2.

The feature matrix Φ , which in the ALP constrains the value functions of (Primal LP), instead *relaxes* the constraint in (Dual LP) to be $[\mathbf{e}_{s_0}^\top + \mu^\top(\gamma \mathbf{P} - \mathbf{E})] \Phi = \mathbf{0}$. As a result, μ may no longer be a discounted state-action occupancy distribution, although it behaves like one with respect to expectations of functions in the span of Φ — using the notation of Sections 1.1 and 1.2, any feasible μ satisfies $\mu^\top \mathbf{E} \mathbf{f} = \mathbf{e}_{s_0}^\top \sum_{t=0}^\infty (\gamma \mathbf{P} \pi)^t \mathbf{f}$ for some policy π and any $\mathbf{f} = \Phi \theta$; compare this with (2).

Conversely, the $\mathbf{W}$ matrix *constrains* μ to be non-zero only on the actions of core states and the current planning state: $\mu = \lambda^\top \mathbf{W}$. The discounted visits to all other states are “soft-aggregated” as positive linear combinations of the core states, as discussed in Section 2. Such aggregation is acceptable because (i) the above relaxation means μ only needs to be accurate for functions in the span of Φ ; and (ii) Assumption 2 ensures that the features of all states lie in the convex hull of the core state features. Thus the simultaneous constraint and relaxation complement each other, incurring the same $O(\epsilon_{\text{approx}}/(1 - \gamma))$ error as the ALP which also restricts value functions to the span of Φ .

Significantly, this theorem also specifies how to select an action that achieves the promised value for the planning state s_0 . This is made possible by (i) adding s_0 to the set of core states; and (ii) requiring (via the definition of Λ) that $\mu_{s_0 a} \equiv \lambda_{s_0 a}$ sum to one. This last constraint forces μ to directly represent the action probabilities at the planning state, not indirectly by being aggregated as linear combinations of the core state actions. As a result, a solution $\mu \equiv \lambda^\top \mathbf{W}$ of (CoreLP) yields an almost-optimal action distribution $\pi(a|s_0) = \mu_{s_0 a}$ for the planning state s_0 .

Unfortunately, the soft state aggregation which makes (CoreLP) tractable to solve (as in Section 4) comes at a price — μ directly encodes only an action distribution for the current planning state, not a policy for other states (unlike the original (Dual LP) of Section 1.2). We believe that solving a separate optimization problem for each planning state is unavoidable without restricting ourselves to a compactly representable class of policies; see Section 5 for a discussion of such approaches.

As a final remark, the value loss of the policy resulting from Theorem 2 is $O(\gamma\epsilon_{\text{approx}}/(1 - \gamma)^2)$. Here, an extra $1/(1 - \gamma)$ factor is incurred in Proposition 1, while the other $1/(1 - \gamma)$ factor is incurred in Theorem 2. This is similar to the bounds obtained in previous works [e.g., 14, 44, 28].

4 CoreStoMP — A Stochastic Saddle-Point Algorithm

Having formulated the planning problem as a linear program with few variables and constraints, the remaining issue is that the constraints still involve quantities of the form $\mathbf{W}\mathbf{P}\Phi$, which cannot be calculated exactly in time independent of S . However, since these are actually expectations, the simulator can be used to estimate them. One possibility would be to estimate $\mathbf{P}_{sa}\Phi \in \mathbb{R}^d$ for the initial and core states and use a plug-in estimator — often called *sample average approximation*. Instead, we pursue the *stochastic approximation* approach — using well-known first-order optimization methods to directly solve (CoreLP) by using the simulator to produce stochastic estimates of gradients that are intractable to compute exactly [20, 21]. This optimization-based approach is attractive to us because the resulting algorithm, by design, is *incremental* and *anytime* — the quality of the solution steadily improves if the algorithm is given more time.

We first rewrite (CoreLP) as an unconstrained “saddle point” problem, retaining only the constraint introduced by the definition of Λ :

$$V^\dagger = \max_{\lambda \in \Lambda} \min_{\theta \in \mathbb{R}^d} [f(\lambda, \theta) := \lambda^\top \mathbf{W}\mathbf{r} + \varphi_0^\top \theta + \lambda^\top \mathbf{W}(\gamma \mathbf{P} - \mathbf{E})\Phi \theta]. \quad (\text{Saddle CoreLP})$$

To be able to use first-order methods, we calculate the gradients of f :

$$f_\lambda(\theta) := \nabla_\lambda f(\lambda, \theta) = \mathbf{W}(\mathbf{r} + (\gamma \mathbf{P} - \mathbf{E})\Phi \theta), \quad (4)$$

$$f_\theta(\lambda) := \nabla_\theta f(\lambda, \theta) = \varphi_0^\top + \lambda^\top \mathbf{W}(\gamma \mathbf{P} - \mathbf{E})\Phi. \quad (5)$$

Note that f is *bilinear*: its gradient with respect to θ only depends linearly on λ , and vice versa. The transition probabilities, which present the major computational challenge, appear only through the matrix $\mathbf{B} := \mathbf{W}(\gamma \mathbf{P} - \mathbf{E})\Phi$, whose rows correspond to state-action pairs $(s, a) \in \mathcal{S}_+ \times \mathcal{A}$. Each row $\mathbf{B}_{sa} = \gamma \mathbf{P}_{sa}\Phi - \varphi_s^\top$ is the (discounted) expected change in the feature vector when taking action a in state s , which suggests how to estimate it using the simulator — define $\Delta\varphi(s, s') := \gamma\varphi_{s'} - \varphi_s$ and sample $s' \sim \mathbf{P}_{sa}$; then $\Delta\tilde{\varphi} := \Delta\varphi(s, s')$ is an unbiased estimator of $\mathbf{B}_{sa}$. The construction of the matrix $\mathbf{W}$ ensures that s is either the current state or one of the core states. Further, we only use s' through its feature representation $\varphi_{s'}$. Putting all this together, our gradient estimates are:

$$[\hat{f}_\lambda(\theta)]_{sa} := \hat{r} + \Delta\varphi(s, s')^\top \theta, \quad \forall s \in \mathcal{S}_+, a \in \mathcal{A}, \quad \text{where } (\hat{r}, s') \sim \text{SIMULATE}(s, a), \quad (6)$$

$$\hat{f}_\theta(\lambda) := \varphi_0^\top + \|\lambda\|_1 \Delta\varphi(s, s'), \quad \text{where } (s, a) \sim \lambda / \|\lambda\|_1 \text{ and } s' \sim \text{SIMULATE}(s, a). \quad (7)$$

Sampling both gradients requires a total of $1 + (1 + m)A$ queries of the simulator and an additional $O(dmA)$ computation time. By slightly abusing notation, we will use $\xi \sim \hat{f}_\theta(\lambda)$ (and $\rho \sim \hat{f}_\lambda(\theta)$) to denote a random d -dimensional (resp., $(1 + m)A$ -dimensional) vector taken from the distribution of $\hat{f}_\theta(\lambda)$ (resp., that of $\hat{f}_\lambda(\theta)$) as defined above. Finally, we remark in passing that the gradient estimate $[\hat{f}_\lambda(\theta)]_{sa}$ is the “temporal difference error” [38] of the value function $\Phi\theta$ at state s with action a .

We use these gradient estimates with the Stochastic Mirror-Prox algorithm of Juditsky et al. [22]. Instantiating the algorithm requires several choices — for the dual variables $\lambda \in \mathbb{R}_+^{(1+m)A}$, we use the 1-norm and the “unnormalized negentropy” regularizer; for the primal variables $\theta \in \mathbb{R}^d$ we use the norm $\|\theta\| = \|\Phi_*\theta\|_2$ and the regularizer $\|\theta\|^2/2$. The result is [Algorithm 1](#) (COREStoMP).

Theorem 3 (CoreStoMP). *Suppose [Assumptions 1, 2, and 3](#) hold, and define*

$$B := \frac{(9/8)\sqrt{m}}{1 - \gamma}, \quad C := \frac{(9/4)\sqrt{m(1 + 2\log A + 2\gamma\log m)}}{(1 - \gamma)^2}.$$

Let $\hat{\lambda}$ be the result of running [Algorithm 1](#) for T iterations with the parameter B and the step size $\eta = C^{-1}\sqrt{2/7T}$, which requires $2T(1 + (1 + m)A)$ simulator queries. Define $\hat{\pi} \in \Delta_{\mathcal{A}}$ by $\hat{\pi}(a) = \hat{\lambda}_{s_0a}$ (as in [Theorem 2](#)) and $a \sim \hat{\pi}$. Then

$$v^*(s_0) - \mathbb{E}[q^*(s_0, a)] \leq \frac{32\varepsilon_{\text{approx}}}{1 - \gamma} + \frac{21}{2(1 - \gamma)^2} \sqrt{\frac{3m(1 + 2\log A + 2\gamma\log m)}{T}}.$$

Note that the expectation on the left-hand side is both for the randomness of the algorithm and the action a . While the bound does not have a direct dependence on the dimension of the features, the number of core states, m , must exceed the rank of Φ . It is notable that the approximation error does

Algorithm 1 CORESTOMP: Stochastic Mirror-Prox for Planning with Core States

Parameters: T, B, η
Initialization: $\theta_0 \leftarrow \mathbf{0} \in \mathbb{R}^d$, $[\lambda_0]_{s_0a} \leftarrow 1/A$, $[\lambda_0]_{sa} \leftarrow \gamma/((1-\gamma)mA) \quad \forall s \in \mathcal{S}_*, a \in \mathcal{A}$
for $\tau = 1, 2, \dots, T$ **do**
 $(\theta'_\tau, \lambda'_\tau) \leftarrow \text{PROXUPDATE}(B, \eta, (\theta_{\tau-1}, \lambda_{\tau-1}), (\xi, \rho)) \quad \text{where } \xi \sim \hat{f}_\theta(\lambda_{\tau-1}), \rho \sim \hat{f}_\lambda(\theta_{\tau-1})$
 $(\theta_\tau, \lambda_\tau) \leftarrow \text{PROXUPDATE}(B, \eta, (\theta_{\tau-1}, \lambda_{\tau-1}), (\xi', \rho')) \quad \text{where } \xi' \sim \hat{f}_\theta(\lambda'_\tau), \rho' \sim \hat{f}_\lambda(\theta'_\tau)$
end for
return $(\sum_{\tau=1}^T \lambda_\tau)/T$

function PROXUPDATE($B, \eta, (\theta, \lambda), (\xi, \rho)$)
 $\tilde{\theta} \leftarrow \theta - \eta\xi$
 $\theta' \leftarrow \tilde{\theta} / \max\{1, \|\Phi_*\theta\|_2/B\}$
 $\tilde{\lambda} \leftarrow \exp(\log \lambda + \eta\rho)$
 $\lambda'_{s_0} \leftarrow \tilde{\lambda}_{s_0} / \|\tilde{\lambda}_{s_0}\|_1 \quad \text{where } \tilde{\lambda}_{s_0} := [\tilde{\lambda}_{s_0a}]_{a \in \mathcal{A}} \text{ and similarly for } \lambda'.$
 $\lambda'_* \leftarrow (\gamma/(1-\gamma))\tilde{\lambda}_* / \|\tilde{\lambda}_*\|_1 \quad \text{where } \tilde{\lambda}_* := [\tilde{\lambda}_{sa}]_{s \in \mathcal{S}_*, a \in \mathcal{A}} \text{ and similarly for } \lambda'.$
return (θ', λ')
end function

not get inflated by a rank-related quantity, as one would expect in the worst-case [28]; this is due to [Assumption 2](#). The increase in the leading term of the approximation error compared to [Theorem 2](#) is because of the need to bound the domain of θ by B ; it remains for future work to avoid this necessity. Altogether, [Algorithm 1](#) gives the following positive result for the online planning problem for MDPs.

Corollary 4. *Under [Assumptions 1, 2, and 3](#), [Algorithm 1](#) is a randomized planning algorithm that, for any $\varepsilon > 0$, uses $O(m^2 A(1 + \log A + \gamma \log m)/\varepsilon^2)$ simulator queries and $\text{poly}(d, A, m, 1/\varepsilon)$ computation to output an action. Following this action in every state gives a stochastic policy with value loss at most $O(\varepsilon_{\text{approx}}/(1-\gamma)^2 + \varepsilon/(1-\gamma)^3)$.*

5 Related Work

The online MDP planning problem formulation we adopt — where the planner is given an input state and asked to produce a close-to-optimal action using a generative model of the MDP as a subroutine — was proposed by Kearns et al. [25] as an alternative to requiring a compact, structured representation of the MDP. Their approach, also adopted by Kocsis and Szepesvári [26] for their UCT algorithm, is to build a (sparse) look-ahead tree. Generally, the problem is that the tree needs to be sufficiently deep and the branching factor can be as large as the number of actions, which leads to an exponential blow-up as a function of the planning horizon (see [footnote 2](#) on page 4). The focus is thus to characterize those MDPs where the planning time can be kept polynomial in the effective horizon [30, 16].

Planning with Feature Representations The broader context of this work is the problem posed by the recent paper of Du et al. [12], which asks whether “good features” (or representation) are sufficient in various RL contexts — including efficient online planning in large MDPs with a generative model. Their main (negative) result states that even when the features are good enough to represent the action-value functions of all policies up to a uniform error of $\varepsilon_{\text{approx}}$, a planning algorithm that is required to produce an $O(\varepsilon_{\text{approx}})$ -optimal policy needs to check at least 2^H states in some H -horizon episodic problems. Lattimore et al. [28] along with Van Roy and Dong [40] point out that if the feature space is d -dimensional, the exponential blowup with the planning horizon can be avoided if the policy only needs to be $O(\varepsilon_{\text{approx}} \sqrt{d}H^2)$ -optimal (where the horizon is $H = 1/(1-\gamma)$, as their results are for discounted problems). They also describe an instance of approximate policy iteration that achieves this bound with $\tilde{O}(d/(\varepsilon_{\text{approx}}^2(1-\gamma)^4))$ queries, where $\tilde{O}$ hides logarithmic factors.

For the finite-horizon setting, Du et al. also present a positive result [12, Theorem C.1] for the case when a simulator of the environment is available and the optimal action-value function can be represented with no error (i.e., $\varepsilon_{\text{approx}} = 0$). The proposed method is a randomized algorithm — an instance of fitted value iteration. In addition to the usual inputs, the algorithm also takes as input δ , a target failure probability. The algorithm returns an optimal policy with probability $1 - \delta$, while issuing at most $\text{poly}(d, H, \log(1/\delta), 1/\rho)$ queries to the simulator, where ρ is the minimum action-value gap that also needs to be known to the algorithm. The algorithm also relies on an oracle to construct a “core set” of d state-action pairs for each stage of the H -horizon problem whose feature vectors form a *barycentric spanner* of the set of all feature vectors at that stage. The idea of the algorithm is to construct a policy backwards by estimating the action value functions via interpolation: In each stage, the action-value of each member of the core set is estimated by using sufficiently many rollouts using the policy constructed for the further stages. The estimated values are used with barycentric interpolation to produce values for all the other state-action pairs.

For the same finite-horizon setting but allowing for an $\varepsilon_{\text{approx}}$ error in approximating the optimal action-value function, Zanette et al. [46] describe a similar algorithm. The main difference is that their algorithm uses the estimated values in a Monte Carlo procedure in place of policy roll-outs. They also propose using a core set (which they call the anchor points) and a similar barycentric extrapolation procedure. Unfortunately, the errors propagate multiplicatively between the stages and thus, in the worst case, the error can be as large as C^H where $C > 1$ depends on the choice of the features. Lattimore et al. [28] show that $1 \leq C \leq \sqrt{d}$; we note in passing that “state aggregation” gives rise to $C = 1$.

A number of authors have studied the problem of learning and planning with exact linear optimal action-value function under various extra conditions. Positive results have been shown for deterministic MDPs [42], the so-called “low Bellman rank” MDPs [18], and under a specific low variance and large gap condition [13]. Yang and Wang [44] assume the transition matrix has a linear structure and also use least-squares regression with data from a pre-selected collection of anchor state/action pairs. Their assumption — the same as ours — is that the features of all state-action pairs can be written as convex combinations of the anchoring features. They show that their algorithm needs at most $\text{poly}(d, 1/(1 - \gamma), \log(1/\delta), m)$ queries, where m is the number of anchor points. Their bound scales linearly with H^7 where $H = 1/(1 - \gamma)$. Their result also applies to the “misspecified” case when the linear structure is only true up to a fixed error. In contrast to these results, we do not assume that the transition matrix has special structure; we make the weaker assumption that the optimal value function lies close to the span of the features.

Approximate Linear Programming The narrower context of the present work is the so-called approximate linear programming (ALP) approach to approximate planning in large MDPs, described in Section 1.2. The seminal work of de Farias and Van Roy [14] showed that the ALP solution’s error, compared to the optimal value function, is within a constant factor (involving $1/(1 - \gamma)$) of the best approximation error achievable by linear combinations of the given features. Unfortunately, as discussed earlier, the ALP has too many constraints to be tractable for large MDPs. Most subsequent work is therefore aimed at designing methods that keep the approximation guarantees without having to enumerate all the constraints. Schuurmans and Patrascu [34] and Guestrin et al. [17] propose using “constraint generation” for problems with additional structure (i.e., factorized transitions), while de Farias and Van Roy [15] propose randomly generating a subset of constraints from some a priori fixed distribution. All these methods require computation time that depends on uncontrolled quantities, such as the so-called induced width of a cost-network [17], or the discrepancy between the sampling distribution and the (unknown) optimal stationary distribution [15]. The fundamental difficulty is that when too many constraints are dropped, the linear program may become unbounded. To protect against this, de Farias and Van Roy [15] add an extra constraint on the optimization variables, but their bound then degrades to the *worst* approximation error over this constraint set.

Petrik and Zilberstein [31] demonstrate that the $1/(1 - \gamma)$ blow-up of the error in the bound of de Farias and Van Roy [14] can be tight. They also propose techniques to avoid it — one of them is to add extra constraints induced by short action sequences; another is to replace the hard constraints in the ALP with smooth ones with an associated Lagrange multiplier. Desai et al. [11] propose a specific way to choose the Lagrange multiplier, for which they also obtain error bounds and demonstrate improved empirical behavior. However, as they build on the work of de Farias and Van Roy [15], their results inherit the limitations of this latter work: the large number of constraints. Bhat et al. [5] extend

the work of Desai et al. [11] to nonparametric function approximation. Lakshminarayanan et al. [27] depart from constraint sampling and consider the error induced by linearly combining constraints. Petrik and Zilberstein [31], in addition to the above mentioned contributions, also give error bounds for the ALP obtained by replacing the transition matrix with a sample-average estimate.

The Dual Linear Program A parallel line of research aims to solve (an approximation of) the (Dual LP) optimization problem, in contrast to the aforementioned work focusing on (Primal LP). Recall from Section 1.2 that the dual variables μ are occupancy distributions over state-action pairs generated by policies — a common theme in these approaches is to approximate such distributions using low-dimensional “distribution features”. Wang et al. [41] introduce this idea in the context of *estimating* the occupancy distribution for a fixed policy rather than directly solving the *planning* problem of finding an optimal policy — the authors suggest using their estimation procedure alongside iterated policy improvement to find an optimal policy, but do not characterize the convergence rate or approximation error of the resulting algorithm. More recently, Abbasi-Yadkori et al. [1, 2] propose a stochastic gradient descent to be used on the Lagrangian derived from the dual LP and derive a policy suboptimality bound for the resulting poly-time algorithm; however, their results only apply under some restrictive conditions.

A major advantage of the (Dual LP) is that its solutions directly encode optimal policies (as discussed in Section 1.2) rather than just value functions. When the dual variables are approximated using “distribution features”, however, only a restricted class of policies can be represented. For example, when the distribution features are the occupancy measures of a given set of “base policies”, then solving the approximate dual LP means finding the best mixture of the base policies. Banijamali et al. [3] present an algorithm for this problem — under the additional assumption that the occupancy measures of the base policies have large overlap. They also show that, in general, this problem is NP-hard to even approximate — finding the best stochastic policy in a restricted class can be *harder* than finding an optimal policy of the MDP. We note in passing that the assumption of a restricted class of policies that contains a close-to-optimal policy can be considered complementary to our setting, where the optimal value function is close to the span of a given feature representation.

Primal-Dual Methods There has been significant interest in applying recent advances in primal-dual online optimization methods to planning in MDPs. Since the (Primal LP) optimizes value functions while the (Dual LP) optimizes occupancy measures (i.e. policies, indirectly), primal-dual optimization can be seen as an “actor-critic” approach that finds both policies and value functions simultaneously. Cogill [10] proposes solving the saddle-point form of the LP in Section 1.2, with no approximation and assuming full knowledge of the transition matrix. Chen and Wang [9] adopt the same approach but with stochastic updates using random samples of state transitions. Chen et al. [8] extend this idea to large MDPs using low-dimensional feature representations to approximate both the primal and dual variables. Bas-Serrano and Neu [4] identify a “coherence” condition on the primal and dual feature representations that is necessary to extract close-to-optimal policies from saddle-point solutions with function approximation — they also point out that, without such an assumption, the policy suboptimality bound of Chen et al. can scale with the number of states of the MDP (or worse). We can avoid this issue in our setting because we use the approximate solutions of (Saddle CoreLP) only to select one action, not an entire policy, unlike all these cited works.

6 Conclusions

We presented an approach to efficient online planning in large-scale γ -discounted MDPs in the presence of (i) a (relatively) weak d -dimensional feature representation; (ii) a core set of m states whose features’ convex hull covers the features of other states; and (iii) a stochastic simulator of the MDP. Our main contribution is an online planning algorithm that, for any target precision ε , achieves a value loss of $O(\varepsilon_{\text{approx}}/(1 - \gamma)^2 + \varepsilon)$, where $\varepsilon_{\text{approx}}$ (3) is the best achievable error in uniformly approximating the optimal value function of the MDP using the given feature representation. When the MDP has A actions per state, the algorithm’s runtime is $\text{poly}(1/\varepsilon, d, m, A, 1/(1 - \gamma))$, which is independent of the number of states in the MDP. Our work builds upon the approximation error bound of Lakshminarayanan et al. [27] for relaxations of the approximate linear program.

Du et al. [12] point out that it remains an open problem whether query-efficient planning is possible in large MDPs using only a simulator and features that have a small approximation error $\varepsilon_{\text{approx}}$,

with no additional assumptions. Our algorithm resolves this open problem in the special case when a small set of core states is available, i.e., when $m = \text{poly}(d)$. It remains an intriguing question whether this assumption can be removed without jeopardizing efficient planning. Other interesting questions are whether the results can be extended to smoothed ALPs [11], and whether the adaptive constraint generation of Petrik and Zilberstein [31] can be used to reduce the dependence on the planning horizon.

To achieve our results, we make several novel technical contributions: We slightly change the ALP approach of Lakshminarayanan et al. [27], adding extra constraints and using a saddle-point formulation. We then show that near-optimal action distributions can be extracted from approximate solutions of the saddle-point problem. We solve the saddle-point problem using a stochastic approximation algorithm, Stochastic Mirror-Prox [22] — a first-order primal-dual optimization method that uses stochastic gradient estimates, which in our case are provided by the simulator. We believe that these techniques and ideas can find applications in other problems beyond our work.

Acknowledgements

Csaba Szepesvári gratefully acknowledges funding from the Canada CIFAR AI Chairs Program, the Alberta Machine Intelligence Institute (Amii), and the Natural Sciences and Engineering Research Council of Canada (NSERC).

References

- [1] Abbasi-Yadkori, Y., Bartlett, P. L., Chen, X., and Malek, A. Large-scale Markov decision problems via the linear programming dual (Jan. 2019). [arXiv:1901.01992](https://arxiv.org/abs/1901.01992).
- [2] Abbasi-Yadkori, Y., Bartlett, P. L., and Malek, A. Linear programming for large-scale Markov decision problems. In *ICML* (2014). <http://proceedings.mlr.press/v32/malek14>.
- [3] Banijamali, E., Abbasi-Yadkori, Y., Ghavamzadeh, M., and Vlassis, N. Optimizing over a restricted policy class in MDPs. In *AISTATS* (2019). <http://proceedings.mlr.press/v89/banijamali19a>.
- [4] Bas-Serrano, J. and Neu, G. Faster saddle-point optimization for solving large-scale Markov decision processes. In *Conference on Learning for Dynamics and Control (LADC)* (2020).
- [5] Bhat, N., Farias, V., and Moallemi, C. C. Non-parametric approximate dynamic programming via the kernel method. In *NeurIPS* (2012). <https://papers.nips.cc/paper/4547-non-parametric-approximate-dynamic-programming-via-the-kernel-method>.
- [6] Blondel, V. D. and Tsitsiklis, J. N. A survey of computational complexity results in systems and control. *Automatica*, 36:1249–1274 (2000). doi:10.1016/S0005-1098(00)00050-9.
- [7] Boucherie, R. J. and van Dijk, N. M., eds. *Markov Decision Processes in Practice*, vol. 248 of *International Series in Operations Research & Management Science*. Springer (2017). doi:10.1007/978-3-319-47766-4.
- [8] Chen, Y., Li, L., and Wang, M. Scalable bilinear π learning using state and action features. In *ICML* (2018). <http://proceedings.mlr.press/v80/chen18e>.
- [9] Chen, Y. and Wang, M. Stochastic primal-dual methods and sample complexity of reinforcement learning (Dec. 2016). [arXiv:1612.02516](https://arxiv.org/abs/1612.02516).
- [10] Cogill, R. Primal-dual algorithms for discounted Markov decision processes. In *European Control Conference* (2015). doi:10.1109/ECC.2015.7330554.
- [11] Desai, V. V., Farias, V. F., and Moallemi, C. C. A smoothed approximate linear program. In *NeurIPS* (2009). <https://papers.nips.cc/paper/3799-a-smoothed-approximate-linear-program>.
- [12] Du, S. S., Kakade, S. M., Wang, R., and Yang, L. F. Is a good representation sufficient for sample efficient reinforcement learning? In *ICLR* (2020). [arXiv:1910.03016](https://arxiv.org/abs/1910.03016).
- [13] Du, S. S., Luo, Y., Wang, R., and Zhang, H. Provably efficient Q-learning with function approximation via distribution shift error checking oracle. In *NeurIPS* (2019). <https://papers.nips.cc/paper/9018-provably-efficient-q-learning-with-function-approximation-via-distribution-shift-error-checking-oracle>.
- [14] de Farias, D. P. and Van Roy, B. The linear programming approach to approximate dynamic programming. *Operations Research*, 51(6):850–865 (Dec. 2003). doi:10.1287/opre.51.6.850.24925.

- [15] de Farias, D. P. and Van Roy, B. On constraint sampling in the linear programming approach to approximate dynamic programming. *Mathematics of Operations Research*, 29:462–478 (2004). doi:10.1287/moor.1040.0094.
- [16] Feldman, Z. and Domshlak, C. Simple regret optimization in online planning for Markov decision processes. *Journal of Artificial Intelligence Research*, 51:165–205 (Sep. 2014). doi:10.1613/jair.4432.
- [17] Guestrin, C., Koller, D., Parr, R., and Venkataraman, S. Efficient solution algorithms for factored MDPs. *Journal of Artificial Intelligence Research*, 19:399–468 (2003). doi:10.1613/jair.1000.
- [18] Jiang, N., Krishnamurthy, A., Agarwal, A., Langford, J., and Schapire, R. E. Contextual decision processes with low Bellman rank are PAC-learnable. In *ICML* (2017). <http://proceedings.mlr.press/v70/jiang17c>.
- [19] Jin, C., Yang, Z., Wang, Z., and Jordan, M. I. Provably efficient reinforcement learning with linear function approximation. In *COLT* (2020). <http://proceedings.mlr.press/v125/jin20a.html>.
- [20] Juditsky, A. and Nemirovski, A. First-order methods for nonsmooth convex large-scale optimization, I: General purpose methods. In [37], chap. 5, pp. 121–148. doi:10.7551/mitpress/8996.003.0007.
- [21] Juditsky, A. and Nemirovski, A. First-order methods for nonsmooth convex large-scale optimization, II: Utilizing problem’s structure. In [37], chap. 6, pp. 149–183. doi:10.7551/mitpress/8996.003.0008.
- [22] Juditsky, A., Nemirovski, A., and Tauvel, C. Solving variational inequalities with Stochastic Mirror-Prox algorithm. *Stochastic Systems*, 1(1):17–58 (Jun. 2011). doi:10.1287/10-SSY011.
- [23] Kakade, S. and Langford, J. Approximately optimal approximate reinforcement learning. In *ICML* (2002). <https://homes.cs.washington.edu/~sham/papers/rl/aoarl.pdf>.
- [24] Kallenberg, L. Markov decision processes (2017). <https://goo.gl/yhvrph>. Lecture Notes.
- [25] Kearns, M., Mansour, Y., and Ng, A. Y. A sparse sampling algorithm for near-optimal planning in large Markov decision processes. *Machine Learning*, 49:193–208 (2002). doi:10.1023/A:1017932429737.
- [26] Kocsis, L. and Szepesvári, Cs. Bandit based Monte-Carlo planning. In *European Conference on Machine Learning* (2006). doi:10.1007/11871842_29.
- [27] Lakshminarayanan, C., Bhatnagar, S., and Szepesvári, Cs. A linearly relaxed approximate linear program for Markov decision processes. *IEEE Transactions on Automatic Control*, 63(4):1185–1191 (Apr. 2018). doi:10.1109/TAC.2017.2743163.
- [28] Lattimore, T., Szepesvári, Cs., and Weisz, G. Learning with good feature representations in bandits and in RL with a generative model. In *ICML* (2020). arXiv:1911.07676.
- [29] Mausam and Kolobov, A. *Planning with Markov Decision Processes: An AI Perspective*, vol. 17 of *Synthesis Lectures on Artificial Intelligence and Machine Learning*. Morgan & Claypool Publishers (2012). doi:10.2200/S00426ED1V01Y201206AIM017.
- [30] Munos, R. From bandits to Monte-Carlo tree search: The optimistic principle applied to optimization and planning. *Foundations and Trends® in Machine Learning*, 7(1):1–129 (2014). doi:10.1561/22000000038.
- [31] Petrik, M. and Zilberstein, S. Constraint relaxation in approximate linear programs. In *ICML* (2009). <https://icml.cc/Conferences/2009/papers/340.pdf>.
- [32] Puterman, M. L. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley Series in Probability and Statistics. Wiley-Interscience, Hoboken, NJ (1994). doi:10.1002/9780470316887.
- [33] Rust, J. Numerical dynamic programming in economics. In *Handbook of Computational Economics*, vol. 1, chap. 14, pp. 619–729. Elsevier, North Holland (1996). doi:10.1016/S1574-0021(96)01016-7.
- [34] Schuurmans, D. and Patrascu, R. Direct value-approximation for factored MDPs. In *NeurIPS* (2001). <http://papers.neurips.cc/paper/1981-direct-value-approximation-for-factored-mdps>.
- [35] Schweitzer, P. J. and Seidmann, A. Generalized polynomial approximations in Markovian decision processes. *Journal of Mathematical Analysis and Applications*, 110(2):568–582 (Sep. 1985). doi:10.1016/0022-247X(85)90317-8.

- [36] Singh, S. P., Jaakkola, T., and Jordan, M. I. Reinforcement learning with soft state aggregation. In *NeurIPS* (1995). <https://papers.nips.cc/paper/981-reinforcement-learning-with-soft-state-aggregation>.
- [37] Sra, S., Nowozin, S., and Wright, S. J., eds. *Optimization for Machine Learning*. Neural Information Processing. MIT Press, Cambridge, Massachusetts (2012). doi:10.7551/mitpress/8996.001.0001.
- [38] Sutton, R. S. Learning to predict by the methods of temporal differences. *Machine Learning*, 3(1):9–44 (1988). doi:10.1007/BF00115009.
- [39] Szepesvári, Cs. Efficient approximate planning in continuous space Markovian decision problems. *AI Communications*, 14(3):163–176 (2001). <https://www.ualberta.ca/~szepesva/papers/aicom.pdf>.
- [40] Van Roy, B. and Dong, S. Comments on the Du-Kakade-Wang-Yang lower bounds (2019). [arXiv:1911.07910](https://arxiv.org/abs/1911.07910).
- [41] Wang, T., Bowling, M., Schuurmans, D., and Lizotte, D. J. Stable dual dynamic programming. In *NeurIPS* (2008). <https://papers.nips.cc/paper/3179-stable-dual-dynamic-programming>.
- [42] Wen, Z. and Van Roy, B. Efficient exploration and value function generalization in deterministic systems. In *NeurIPS* (2013). <https://papers.nips.cc/paper/4972-efficient-exploration-and-value-function-generalization-in-deterministic-systems>.
- [43] White, D. J. A survey of applications of Markov decision processes. *The Journal of the Operational Research Society*, 44(11):1073–1096 (1993). doi:10.2307/2583870.
- [44] Yang, L. and Wang, M. Sample-optimal parametric Q-learning using linearly additive features. In *ICML* (2019). <http://proceedings.mlr.press/v97/yang19b>.
- [45] Yu, Y.-L. The strong convexity of von Neumann’s entropy (Jun. 2013). <http://www.cs.cmu.edu/~yaoliang/mynotes/sc.pdf>.
- [46] Zanette, A., Lazaric, A., Kochenderfer, M. J., and Brunskill, E. Limiting extrapolation in linear approximate value iteration. In *NeurIPS* (2019). <https://papers.nips.cc/paper/8799-limiting-extrapolation-in-linear-approximate-value-iteration>.

Appendix

Index of Notation

For the convenience of the reader, we have collected the most frequently used symbols and their meanings in the following table:

$\mathbb{R}, \mathbb{R}_+$	real numbers; non-negative real numbers.
$\mathbb{R}^d, \mathbb{R}^{m \times n}$	d -dimensional vectors; matrices of size $m \times n$.
$\mathbf{e}_i, \mathbf{0}, \mathbf{1}$	standard basis vector: $e_{i,i} = 1$ and $e_{i,j} = 0$ for $i \neq j$; constant zero or one vectors.
$\mathbf{a} \oplus \mathbf{b}$	concatenation of vectors: if $\mathbf{a} \in \mathbb{R}^n$ and $\mathbf{b} \in \mathbb{R}^m$, then $\mathbf{a} \oplus \mathbf{b} \in \mathbb{R}^{n+m}$.
$\mathcal{S}, \mathcal{A}, \mathcal{S}_*$	sets of states, actions, and core states (Section 1.1 and Assumption 2).
S, A, m	number of states $ \mathcal{S} $, actions $ \mathcal{A} $, and core states $ \mathcal{S}_* $, respectively.
s_0, s, s', a	planning state (Section 3) and other states $\in \mathcal{S}$; actions $\in \mathcal{A}$.
$\mathbf{P}, \mathbf{E}$	row-stochastic matrices in $\mathbb{R}_+^{S \times S}$ (Sections 1.1 and 1.2); $\mathbf{E}_{sa} = \mathbf{e}_s \in \mathbb{R}^S$.
$\mathbf{r}, \hat{\mathbf{r}}$	expected rewards $\in \mathbb{R}^{SA}$ (Section 1.1); random reward $\in [-1, 1]$ (Assumption 3).
$\gamma \in [0, 1)$	discount factor (Section 1.1).
v, v^*, v_π	value functions $\mathcal{S} \rightarrow \mathbb{R}$.
$\Delta_{\mathcal{S}}, \Delta_{\mathcal{A}}$	sets of probability distributions over states and actions.
$\boldsymbol{\mu}, \boldsymbol{\pi}, \boldsymbol{\pi}(s)$	probability distributions in $\Delta_{\mathcal{S}}$ and $\Delta_{\mathcal{A}}$, respectively; policy in $\mathcal{S} \rightarrow \Delta_{\mathcal{A}}$.
$\boldsymbol{\Phi}, \boldsymbol{\varphi}_s, \boldsymbol{\varphi}_0$	feature matrix $\in \mathbb{R}^{S \times d}$; state features $\in \mathbb{R}^d$; features of planning state $\boldsymbol{\varphi}_{s_0}$.
$\varepsilon_{\text{approx}}$	approximation error of $\boldsymbol{\Phi}$: $\min_{\boldsymbol{\theta} \in \mathbb{R}^d} \ \mathbf{v}^* - \boldsymbol{\Phi}\boldsymbol{\theta}\ _\infty$.
$\mathbf{W}, \mathbf{W}_*$	constraint matrices in $\{0, 1\}^{(1+m)A \times SA}$ and $\{0, 1\}^{mA \times SA}$ (Theorems 2 and 5).
$\lambda, \lambda_*, \boldsymbol{\theta}$	dual variables $\in \mathbb{R}_+^{(1+m)A}$ and $\in \mathbb{R}^{mA}$, respectively; primal variables $\in \mathbb{R}^d$.
$\Lambda, \Lambda_\gamma, \mathcal{B}$	dual spaces $\subset \mathbb{R}_+^{(1+m)A}$; primal space $\subset \mathbb{R}^d$ (Theorem 2 and Lemmas 8 and 11).
$\ \cdot\ _*$	dual norm of $\ \cdot\ $: defined by $\ \mathbf{u}\ _* = \sup_{\ \mathbf{x}\ =1} \langle \mathbf{u}, \mathbf{x} \rangle$.

A Proofs

A.1 Approximation Error for the Linearly Relaxed Approximate LP

We start by recalling and improving the approximation error bounds for the Linearly Relaxed Approximate Linear Program (LRALP $_\mu$) of Lakshminarayanan et al. [27].

Theorem 5. Suppose Assumptions 1 and 2 hold. Define the matrix $\mathbf{W}_* \in \{0, 1\}^{mA \times SA}$ with rows $[\mathbf{W}_*]_{sa} = \mathbf{e}_{sa} \in \mathbb{R}^{SA}$ ($s \in \mathcal{S}_*, a \in \mathcal{A}$). For any (possibly unnormalized) initial distribution $\boldsymbol{\mu} \in \mathbb{R}_+^S$,

$$V_{\text{LRALP}}(\boldsymbol{\mu}) := \min \{ \boldsymbol{\mu}^\top \boldsymbol{\Phi} \boldsymbol{\theta} \mid \boldsymbol{\theta} \in \mathbb{R}^d, \mathbf{W}_* \mathbf{r} + \mathbf{W}_* (\gamma \mathbf{P} - \mathbf{E}) \boldsymbol{\Phi} \boldsymbol{\theta} \leq \mathbf{0} \}. \quad (\text{LRALP}_\mu)$$

The value of (LRALP $_\mu$) is close to the optimal value of that initial distribution:

$$|V_{\text{LRALP}}(\boldsymbol{\mu}) - \boldsymbol{\mu}^\top \mathbf{v}^*| \leq \frac{10 \|\boldsymbol{\mu}\|_1 \varepsilon_{\text{approx}}}{1 - \gamma}.$$

This result follows from Lakshminarayanan et al. [27, Theorem IV.1], which we will not reproduce here for brevity. The error bound there is $2\|\boldsymbol{\mu}\|_1(3\varepsilon_{\text{approx}} + \|\mathbf{J}_{\text{ALP}}^* - \mathbf{J}_{\text{LRA}}^*\|_\infty)/(1 - \gamma)$, defining

$$\begin{aligned} J_{\text{ALP}}^*(s) &:= \min \{ \boldsymbol{\varphi}_s^\top \boldsymbol{\theta} \mid \boldsymbol{\theta} \in \mathbb{R}^d, \boldsymbol{\Phi} \boldsymbol{\theta} \geq \mathbf{v}^* \}, \\ J_{\text{LRA}}^*(s) &:= \min \{ \boldsymbol{\varphi}_s^\top \boldsymbol{\theta} \mid \boldsymbol{\theta} \in \mathbb{R}^d, \mathbf{W}_* \mathbf{E} \boldsymbol{\Phi} \boldsymbol{\theta} \geq \mathbf{W}_* \mathbf{E} \mathbf{v}^* \}. \end{aligned}$$

It only remains for us to bound $\|\mathbf{J}_{\text{ALP}}^* - \mathbf{J}_{\text{LRA}}^*\|_\infty$, improving upon Theorem IV.2 [27]:

Lemma 6. Under the conditions of Theorem 5, $\|\mathbf{J}_{\text{ALP}}^* - \mathbf{J}_{\text{LRA}}^*\|_\infty \leq 2\varepsilon_{\text{approx}}$.

Proof. By Assumption 1, the optimal value function is well-approximated by the feature representation; $\mathbf{v}^* = \boldsymbol{\Phi} \boldsymbol{\theta} + \boldsymbol{\delta}$ for some $\boldsymbol{\theta} \in \mathbb{R}^d$ and $\boldsymbol{\delta} \in \mathbb{R}^S$ with $\|\boldsymbol{\delta}\|_\infty \leq \varepsilon_{\text{approx}}$. By Assumption 2, $\boldsymbol{\Phi} = \mathbf{Z} \boldsymbol{\Phi}_*$, so $\mathbf{v}^* = \mathbf{Z} \boldsymbol{\Phi}_* \boldsymbol{\theta} + \boldsymbol{\delta}$. We use these facts after writing the linear program defining $J_{\text{ALP}}^*(s)$ in its dual form:

$$\begin{aligned} J_{\text{ALP}}^*(s) &= \max \{ \boldsymbol{\mu}^\top \mathbf{v}^* \mid \boldsymbol{\mu} \in \mathbb{R}_+^S, \boldsymbol{\mu}^\top \boldsymbol{\Phi} = \boldsymbol{\varphi}_s^\top \} \\ &= \max \{ \boldsymbol{\mu}^\top (\mathbf{Z} \boldsymbol{\Phi}_* \boldsymbol{\theta} + \boldsymbol{\delta}) \mid \boldsymbol{\mu} \in \mathbb{R}_+^S, \boldsymbol{\mu}^\top \mathbf{Z} \boldsymbol{\Phi}_* = \boldsymbol{\varphi}_s^\top \} \end{aligned}$$

By [Assumption 1](#), there is some $\eta \in \mathbb{R}^d$ such that $\Phi\eta = \mathbf{1}$. If $\mu^\top \Phi = \varphi_s^\top$, then $\|\mu\|_1 = \mu^\top \mathbf{1} = \mu^\top \Phi \eta = \varphi_s^\top \eta = 1$, which means that $\mu^\top \delta \leq \|\delta\|_\infty$. Replacing $\mu^\top \delta$ with $\|\delta\|_\infty$ in the objective increases its value; we move the resulting constant term out of the maximization:

$$\leq \|\delta\|_\infty + \max \{ \mu^\top Z \Phi_* \theta \mid \mu \in \mathbb{R}_+^S, \mu^\top Z \Phi_* = \varphi_s^\top \}$$

The objective and constraints of this maximization problem depend on μ only through $\mu^\top Z$. Thus we can replace $\mu^\top Z$ with $\mu_* \in \mathbb{R}_+^m$, which can only expand the feasible set of the maximization and increase its value:

$$\leq \|\delta\|_\infty + \max \{ \mu_*^\top \Phi_* \theta \mid \mu_* \in \mathbb{R}_+^m, \mu_*^\top \Phi_* = \varphi_s^\top \}$$

The matrix $U \in \{0, 1\}^{m \times S}$ with rows $[U_s]_{s \in \mathcal{S}_*} = e_s$ can be used to “select” the core state features from Φ , giving $\Phi_* = U\Phi$:

$$= \|\delta\|_\infty + \max \{ \mu_*^\top U \Phi \theta \mid \mu_* \in \mathbb{R}_+^m, \mu_*^\top U \Phi = \varphi_s^\top \}$$

By a similar argument as before, we see that $\|\mu_*^\top U\|_1 = 1$. We add $\mu_*^\top U \delta + \|\delta\|_\infty \geq 0$ to the objective (increasing its value), then move the constant out:

$$\begin{aligned} &\leq 2\|\delta\|_\infty + \max \{ \mu_*^\top U \Phi \theta + \mu_*^\top U \delta \mid \mu_* \in \mathbb{R}_+^m, \mu_*^\top U \Phi = \varphi_s^\top \} \\ &= 2\|\delta\|_\infty + \max \{ \mu_*^\top U v^* \mid \mu_* \in \mathbb{R}_+^m, \mu_*^\top U \Phi = \varphi_s^\top \} \\ &= 2\|\delta\|_\infty + \min \{ \varphi_s^\top \theta \mid \theta \in \mathbb{R}^d, U \Phi \theta \geq U v^* \}, \end{aligned}$$

where the last step is obtained by writing the dual of the linear program in the previous step. Now observe that the constraint $U \Phi \theta \geq U v^*$ is equivalent to the constraint $W_* E \Phi \theta \geq W_* E v^*$ in the definition of J_{LRA}^* — both of them require that $\varphi_s \theta \geq v_s^*$ for $s \in \mathcal{S}_*$. Thus we have shown that $J_{\text{ALP}}^*(s) - J_{\text{LRA}}^*(s) \leq 2\varepsilon_{\text{approx}}$ for all $s \in \mathcal{S}$. We also know that $J_{\text{ALP}}^*(s) \geq J_{\text{LRA}}^*(s)$, since $J_{\text{LRA}}^*(s)$ is a relaxation of $J_{\text{ALP}}^*(s)$. It follows that $\|J_{\text{ALP}}^* - J_{\text{LRA}}^*\|_\infty \leq 2\varepsilon_{\text{approx}}$. $\square$

A.2 Proof of [Theorem 2](#) — Approximation Error for CoreLP

Theorem 2 (CoreLP). Suppose [Assumptions 1](#) and [2](#) hold, $s_0 \in \mathcal{S}$, $\varphi_0 := \varphi_{s_0}$, $W \in \{0, 1\}^{(1+m)A \times SA}$ has rows $[W_{s_i a}]_{s_i \in \mathcal{S}_*, a \in \mathcal{A}} = e_{s_i a}$, and $\Lambda := \{\lambda \in \mathbb{R}_+^{(1+m)A} \mid \sum_{a \in \mathcal{A}} \lambda_{s_0 a} = 1\}$. Define

$$V^\dagger = \max \{ \lambda^\top W r \mid \lambda \in \Lambda, \varphi_0^\top + \lambda^\top W (\gamma P - E) \Phi = \mathbf{0} \}. \quad (\text{CoreLP})$$

Let $\lambda^\dagger \in \Lambda$ be a maximizer of [\(CoreLP\)](#) and let $\pi^\dagger \in \Delta_{\mathcal{A}}$ be given by $\pi^\dagger(a) = \lambda_{s_0 a}^\dagger$. Then

$$|V^\dagger - v^*(s_0)| \leq \frac{10\gamma\varepsilon_{\text{approx}}}{1-\gamma}, \quad v^*(s_0) - \sum_{a \in \mathcal{A}} \pi^\dagger(a) q^*(s_0, a) \leq \frac{20\gamma\varepsilon_{\text{approx}}}{1-\gamma}.$$

By the definition of $\Lambda \subset \mathbb{R}_+^{(1+m)A}$, we can decompose its elements as $\lambda = \pi \oplus \lambda_*$, with $\pi \in \Delta_{\mathcal{A}}$ as in the statement of the theorem and $\lambda_* \in \mathbb{R}_+^{mA}$ defined by $\lambda_{*,sa} = \lambda_{sa}$ for $s \in \mathcal{S}_*, a \in \mathcal{A}$ — in other words, $\Lambda \cong \Delta_{\mathcal{A}} \times \mathbb{R}_+^{mA}$. The main idea of the proof is that when λ is a solution of [\(CoreLP\)](#), then λ_* is a solution for the dual form of [\(LRALP \$_\mu\$ \)](#) from [Theorem 5](#). To make this connection between the two problems more precise, let us write the saddle-point forms of [\(LRALP \$_\mu\$ \)](#) and [\(CoreLP\)](#):

$$V_{\text{LRALP}}(\mu) = \max_{\lambda_* \in \mathbb{R}_+^{mA}} \min_{\theta \in \mathbb{R}^d} [g_\mu(\lambda_*, \theta) := \lambda_*^\top W_* r + \mu^\top \Phi \theta + \lambda_*^\top W (\gamma P - E) \Phi \theta] \quad (\text{Saddle LRALP}_\mu)$$

$$V^\dagger = \max_{\lambda \in \Lambda} \min_{\theta \in \mathbb{R}^d} [f(\lambda, \theta) := \lambda^\top W r + e_{s_0}^\top \Phi \theta + \lambda^\top W (\gamma P - E) \Phi \theta] \quad (\text{Saddle CoreLP})$$

Lemma 7 (Corresponding [\(LRALP \$_\mu\$ \)](#) and [\(CoreLP\)](#) solutions). Let $\lambda \in \Lambda \subset \mathbb{R}_+^{(1+m)A}$ be arbitrary and decompose it as $\lambda = \pi \oplus \lambda_*$, where $\pi \in \Delta_{\mathcal{A}}$ and $\lambda_* \in \mathbb{R}_+^{mA}$. Define the distribution $\mu_\pi \in \Delta_{\mathcal{S}}$ as

$$\mu_\pi^\top := \sum_{a \in \mathcal{A}} \pi(a) P_{s_0 a}, \quad \text{where } \pi(a) := \lambda_{s_0 a} \text{ for } a \in \mathcal{A}.$$

Then, for any $\theta \in \mathbb{R}^d$, and $g_\mu(\lambda_*, \theta)$ and $f(\lambda, \theta)$ as in [\(Saddle LRALP \$_\mu\$ \)](#) and [\(Saddle CoreLP\)](#),

$$f(\lambda, \theta) = \sum_{a \in \mathcal{A}} \pi(a) r(s_0, a) + g_{\mu_\pi}(\lambda_*, \theta), \quad \text{where } \lambda_{*,sa} = \lambda_{sa} \text{ for } s \in \mathcal{S}_*, a \in \mathcal{A}.$$

P_{s_0a} is the next-state distribution for action a at state s_0 — thus the distribution $\mu_\pi \in \Delta_S$ defined here is the expected next-state distribution when an action $a \sim \pi$ is taken at state s_0 . This lemma therefore connects solutions of (CoreLP) with (LRALP $_\mu$) when $\mu = \gamma\mu_\pi$ is the discounted next-state distribution for action $a \sim \pi$.

Proof of Lemma 7. Recall that W and W_* (defined in Theorems 2 and 5) are related — the rows of W_* correspond to state-action pairs in $S_* \times \mathcal{A}$, to which W adds A more rows corresponding to the actions at the current planning state s_0 . Thus

$$\lambda^\top W = \sum_{a \in \mathcal{A}} \pi_a e_{s_0a}^\top + \lambda_*^\top W_*, \quad (8)$$

which upon multiplying by r gives

$$\lambda^\top W r = \sum_{a \in \mathcal{A}} \pi_a r_{s_0a} + \lambda_*^\top W_* r. \quad (9)$$

Using (8) again,

$$\begin{aligned} e_{s_0}^\top + \lambda^\top W(\gamma P - E) &= e_{s_0}^\top + \sum_{a \in \mathcal{A}} \pi_a (\gamma P_{s_0a} - E_{s_0a}) + \lambda_*^\top W_*(\gamma P - E), \\ &= [e_{s_0}^\top - \sum_{a \in \mathcal{A}} \pi_a E_{s_0a}] + \gamma [\sum_{a \in \mathcal{A}} \pi_a P_{s_0a}] + \lambda_*^\top W_*(\gamma P - E). \end{aligned}$$

The first term is zero because $E_{s_0a} = e_{s_0}$ for all $a \in \mathcal{A}$, and the second term becomes $\gamma\mu_\pi$ when we substitute the definition of μ_π . We then multiply both sides by $\Phi\theta$:

$$e_{s_0}^\top \Phi\theta + \lambda^\top W(\gamma P - E)\Phi\theta = \gamma\mu_\pi^\top \Phi\theta + \lambda_*^\top W_*(\gamma P - E)\Phi\theta.$$

Adding this to (9) gives

$$\lambda^\top W r + e_{s_0}^\top \Phi\theta + \lambda^\top W(\gamma P - E)\Phi\theta = \sum_{a \in \mathcal{A}} \pi_a r_{s_0a} + \lambda_*^\top W_* r + \gamma\mu_\pi^\top \Phi\theta + \lambda_*^\top W_*(\gamma P - E)\Phi\theta,$$

where we substitute the definitions of $f(\lambda, \theta)$ and $g_{\gamma\mu_\pi}(\lambda_*, \theta)$ to get the desired result:

$$f(\lambda, \theta) = \sum_{a \in \mathcal{A}} \pi_a r_{s_0a} + g_{\gamma\mu_\pi}(\lambda_*, \theta). \quad \square$$

Proof of Theorem 2. Using the decomposition $\Lambda \cong \Delta_{\mathcal{A}} \times \mathbb{R}_+^{mA}$ in (Saddle CoreLP):

$$\begin{aligned} V^\dagger &= \max_{\pi \in \Delta_{\mathcal{A}}} \max_{\lambda_* \in \mathbb{R}_+^{mA}} \min_{\theta \in \mathbb{R}^d} f(\pi \oplus \lambda_*, \theta) && \text{(where } \pi \oplus \lambda_* = \lambda \in \Lambda) \\ &= \max_{\pi \in \Delta_{\mathcal{A}}} \max_{\lambda_* \in \mathbb{R}_+^{mA}} \min_{\theta \in \mathbb{R}^d} \left[\sum_{a \in \mathcal{A}} \pi_a r_{s_0a} + g_{\gamma\mu_\pi}(\lambda_*, \theta) \right] && \text{(using Lemma 7)} \\ &= \max_{\pi \in \Delta_{\mathcal{A}}} \left[\sum_{a \in \mathcal{A}} \pi_a r_{s_0a} + \max_{\lambda_* \in \mathbb{R}_+^{mA}} \min_{\theta \in \mathbb{R}^d} g_{\gamma\mu_\pi}(\lambda_*, \theta) \right] \\ &= \max_{\pi \in \Delta_{\mathcal{A}}} \left[q^\dagger(\pi) := \sum_{a \in \mathcal{A}} \pi_a r_{s_0a} + V_{\text{LRALP}}(\gamma\mu_\pi) \right]. && \text{(from (Saddle LRALP}_\mu)) \end{aligned}$$

We now turn our attention to bounding $V^\dagger$. From Theorem 5, we know that $|V_{\text{LRALP}}(\gamma\mu_\pi) - \gamma\mu_\pi^\top v^*| \leq 10\gamma\epsilon_{\text{approx}}/(1-\gamma)$ for any distribution over states $\mu_\pi \in \Delta_S$. Through a slight abuse of notation, we define $q^*(s_0, \pi) := \sum_a \pi_a r_{s_0a} + \gamma\mu_\pi^\top v^*$ as a generalization of the standard $q^*(s, a)$ value function to action distributions. Note that we will only need $q^*(s_0, \cdot)$, for which this abuse is ‘sensible’. Then for all $\pi \in \Delta_{\mathcal{A}}$,

$$|q^\dagger(\pi) - q^*(s_0, \pi)| = |V_{\text{LRALP}}(\gamma\mu_\pi) - \gamma\mu_\pi^\top v^*| \leq \frac{10\gamma\epsilon_{\text{approx}}}{1-\gamma}. \quad (10)$$

We also know that $v^*(s_0) = \max_{\pi \in \Delta_{\mathcal{A}}} q^*(s_0, \pi)$ (the equality happens with $\pi^* = e_{a^*}$ for an optimal action a^*). Hence,

$$|V^\dagger - v^*(s_0)| = |\max_{\pi \in \Delta_{\mathcal{A}}} q^\dagger(\pi) - \max_{\pi \in \Delta_{\mathcal{A}}} q^*(s_0, \pi)| \leq \max_{\pi \in \Delta_{\mathcal{A}}} |q^\dagger(\pi) - q^*(s_0, \pi)| \leq \frac{10\gamma\epsilon_{\text{approx}}}{1-\gamma},$$

where the last inequality follows from (10).

For the second part of the result, let $\lambda^\dagger$ be a maximizer of (CoreLP) and $\pi^\dagger$ be the action-distribution component (as before) so that $V^\dagger = q^\dagger(\pi^\dagger)$. Then, using again (10), combined with the last inequality,

$$\begin{aligned} \sum_{a \in \mathcal{A}} \pi^\dagger(a) q^*(s_0, a) &\equiv q^*(s_0, \pi^\dagger) \geq q^\dagger(\pi^\dagger) - \frac{10\gamma\epsilon_{\text{approx}}}{1-\gamma} \\ &= V^\dagger - \frac{10\gamma\epsilon_{\text{approx}}}{1-\gamma} \\ &\geq v^*(s_0) - \frac{20\gamma\epsilon_{\text{approx}}}{1-\gamma}. \end{aligned}$$

Reordering gives the desired result, namely that $v^*(s_0) - \sum_a \pi^\dagger(a) q^*(s_0, a) \leq 20\gamma\epsilon_{\text{approx}}/(1-\gamma)$. $\square$

A.3 Proof of Theorem 3 — Error Bounds for the CoreStoMP Algorithm

Theorem 3 (CoreStoMP). *Suppose Assumptions 1, 2, and 3 hold, and define*

$$B := \frac{(9/8)\sqrt{m}}{1-\gamma}, \quad C := \frac{(9/4)\sqrt{m(1+2\log A+2\gamma\log m)}}{(1-\gamma)^2}.$$

Let $\hat{\lambda}$ be the result of running Algorithm 1 for T iterations with the parameter B and the step size $\eta = C^{-1}\sqrt{2/7T}$, which requires $2T(1+(1+m)A)$ simulator queries. Define $\hat{\pi} \in \Delta_{\mathcal{A}}$ by $\hat{\pi}(a) = \hat{\lambda}_{s_0 a}$ (as in Theorem 2) and $a \sim \hat{\pi}$. Then

$$v^*(s_0) - \mathbb{E}[q^*(s_0, a)] \leq \frac{32\epsilon_{\text{approx}}}{1-\gamma} + \frac{21}{2(1-\gamma)^2} \sqrt{\frac{3m(1+2\log A+2\gamma\log m)}{T}}.$$

The proof of this theorem has two main ingredients: First, in Lemma 8, we show that approximate solutions of (Saddle CoreLP) can be used to recover near-optimal action distributions for the planning state s_0 — the approximation quality is measured by the *duality gap*. Second, in Lemma 11, we bound the expected duality gap of the Stochastic Mirror-Prox algorithm when specialized to our setting.

Lemma 8 (Approximate (Saddle CoreLP) solutions). *Suppose $\mathcal{B} \subset \mathbb{R}^d$ and $C_{\mathcal{B}} \geq 0$ are chosen such that, for any distribution over states $\mu \in \Delta_{\mathcal{S}}$, there is some $\theta \in \mathcal{B}$ that is feasible for (LRALP $_{\mu}$) and at most $C_{\mathcal{B}}$ -suboptimal. Define*

$$\Lambda_{\gamma} := \{\lambda \in \Lambda \mid \|\lambda\|_1 = 1/(1-\gamma)\}, \quad (11)$$

a subset of the set $\Lambda \subset \mathbb{R}_+^{(1+m)A}$ from Theorem 2. Define the $\mathcal{B}$ -bounded duality gap of an approximate solution of (Saddle CoreLP) as

$$\delta_{\mathcal{B}}(\hat{\lambda}, \hat{\theta}) := \max_{\lambda \in \Lambda_{\gamma}} f(\lambda, \hat{\theta}) - \inf_{\theta \in \mathcal{B}} f(\hat{\lambda}, \theta), \quad \text{where } \hat{\lambda} \in \Lambda \text{ and } \hat{\theta} \in \mathbb{R}^d. \quad (12)$$

For any $\hat{\lambda} \in \Lambda$ and $\hat{\theta} \in \mathbb{R}^d$, let $\hat{\pi}$ be the action distribution component of $\hat{\lambda}$, as in Theorem 2. Then

$$v^*(s_0) - \sum_{a \in \mathcal{A}} \hat{\pi}(a) q^*(s_0, a) \leq \frac{20\gamma\epsilon_{\text{approx}}}{1-\gamma} + \gamma C_{\mathcal{B}} + \delta_{\mathcal{B}}(\hat{\lambda}, \hat{\theta}).$$

This lemma generalizes the second result of Theorem 2 in two ways: First, the Stochastic Mirror-Prox algorithm does not produce exact solutions of (Saddle CoreLP); the optimization error is measured by the duality gap — here we see the effect of a non-zero duality gap on the resulting action distribution. Second, the primal variables θ in (Saddle CoreLP) have the unbounded domain $\mathbb{R}^d$, whereas the Stochastic Mirror-Prox algorithm requires the optimization domain to have a bounded *diameter*; see Proof of Lemma 11 — Stochastic Mirror-Prox. This lemma shows that restricting θ to a large-enough bounded set $\mathcal{B}$ only incurs an additional $C_{\mathcal{B}}$ error. Indeed, the second issue is related to the first — an *unbounded* form of the duality gap would be infinite for *any* approximate solution, making it useless as a measure of optimization accuracy; the $\mathcal{B}$ -bounded duality gap therefore addresses both these issues:

Claim 9. For any $\hat{\lambda} \in \Lambda$, $\hat{\theta} \in \mathbb{R}^d$, $\theta \in \mathcal{B} \subset \mathbb{R}^d$, and $\delta_{\mathcal{B}}(\hat{\lambda}, \hat{\theta})$ being the $\mathcal{B}$ -bounded duality gap (12),

$$V^\dagger \leq f(\hat{\lambda}, \theta) + \delta_{\mathcal{B}}(\hat{\lambda}, \hat{\theta}).$$

Proof. Let $\lambda^* \in \Lambda \subset \mathbb{R}_+^{(1+m)A}$ be a maximizer of (CoreLP) — this exists because the optimization is bounded (Theorem 2). Then

$$\begin{aligned} \mathbf{0} &= \varphi_0^\top + \lambda^{*\top} \mathbf{W}(\gamma \mathbf{P} - \mathbf{E}) \Phi && \text{(since } \lambda^* \text{ is feasible for (CoreLP))} \\ &= \varphi_0^\top \hat{\theta} + \lambda^{*\top} \mathbf{W}(\gamma \mathbf{P} - \mathbf{E}) \Phi \hat{\theta}. && \text{(multiplying by } \hat{\theta} \in \mathbb{R}^d \end{aligned} \quad (13)$$

Since, λ^* is a maximizer of (CoreLP), $V^\dagger = \lambda^{*\top} \mathbf{W} \mathbf{r}$:

$$\begin{aligned} V^\dagger &= \lambda^{*\top} \mathbf{W} \mathbf{r} + \varphi_0^\top \hat{\theta} + \lambda^{*\top} \mathbf{W}(\gamma \mathbf{P} - \mathbf{E}) \Phi \hat{\theta} && \text{(adding } V^\dagger \text{ on l.h.s. and } \lambda^{*\top} \mathbf{W} \mathbf{r} \text{ on r.h.s.)} \\ &= f(\lambda^*, \hat{\theta}). && \text{(definition of } f \text{ from (Saddle CoreLP))} \end{aligned} \quad (14)$$

Assumption 1 tells us that $\Phi \eta = \mathbf{1}$ for some $\eta \in \mathbb{R}^d$ — multiplying (13) by η , we see that λ^* must satisfy $1 + \gamma \|\lambda^*\|_1 = \|\lambda^*\|_1$, as does any other feasible solution of (CoreLP). In particular, this means that $\|\lambda^*\|_1 = 1/(1 - \gamma)$ and so $\lambda^* \in \Lambda_\gamma$. Using the definition of $\delta_{\mathcal{B}}$ from (12),

$$\begin{aligned} \delta(\hat{\lambda}, \hat{\theta}) &\geq f(\lambda^*, \hat{\theta}) - f(\hat{\lambda}, \theta) && \text{(since } \lambda^* \in \Lambda_\gamma \text{ and } \theta \in \mathcal{B}) \\ &= V^\dagger - f(\hat{\lambda}, \theta). && \text{(using (14))} \end{aligned} \quad \square$$

Claim 10. For any $\hat{\lambda}_* \in \mathbb{R}_+^{m_A}$ and distribution over states $\mu \in \Delta_{\mathcal{S}}$, suppose $\theta \in \mathbb{R}^d$ is feasible for (LRALP $_\mu$) and at most $C_{\mathcal{B}}$ -suboptimal. Then, with g_μ being the objective function of (Saddle LRALP $_\mu$),

$$V_{\text{LRALP}}(\gamma \mu) \geq g_{\gamma \mu}(\hat{\lambda}_*, \theta) - \gamma C_{\mathcal{B}}.$$

Proof. Since θ is feasible for (LRALP $_\mu$), $\mathbf{W}_* \mathbf{r} + \mathbf{W}_*(\gamma \mathbf{P} - \mathbf{E}) \Phi \theta \leq \mathbf{0}$. Multiplying both sides of this inequality by $\hat{\lambda}_*^\top \geq \mathbf{0}$,

$$\begin{aligned} \hat{\lambda}_*^\top \mathbf{W}_* \mathbf{r} + \hat{\lambda}_*^\top \mathbf{W}_*(\gamma \mathbf{P} - \mathbf{E}) \Phi \theta &\leq 0 \\ \hat{\lambda}_*^\top \mathbf{W}_* \mathbf{r} + \gamma \mu^\top \Phi \theta + \hat{\lambda}_*^\top \mathbf{W}_*(\gamma \mathbf{P} - \mathbf{E}) \Phi \theta &\leq \gamma \mu^\top \Phi \theta && \text{(adding } \gamma \mu^\top \Phi \theta \text{ to both sides)} \\ g_{\gamma \mu}(\hat{\lambda}_*, \theta) &\leq \gamma \mu^\top \Phi \theta. && \text{(definition of } g \text{ from (Saddle LRALP}_\mu)) \end{aligned}$$

Note that the choice of μ does not affect the constraints of (LRALP $_\mu$), only its objective function — thus θ is feasible for the problem defining $V_{\text{LRALP}}(\gamma \mu)$ and is $\gamma C_{\mathcal{B}}$ suboptimal: $\gamma \mu^\top \Phi \theta \leq V_{\text{LRALP}}(\gamma \mu) + \gamma C_{\mathcal{B}}$. Substituting this into the last inequality and rearranging gives the desired result. $\square$

Proof of Lemma 8. As in Lemma 7, we write $\hat{\lambda} = \hat{\pi} \oplus \hat{\lambda}_*$ with $\hat{\pi} \in \Delta_{\mathcal{A}}$ and $\hat{\lambda}_* \in \mathbb{R}_+^{m_A}$ and define $\mu_{\hat{\pi}} = \sum_a \hat{\pi}_a \mathbf{P}_{s_0 a}$. By our assumption, there is some $\theta \in \mathcal{B}$ that is feasible and at most $C_{\mathcal{B}}$ -suboptimal for (LRALP $_\mu$) with $\mu = \mu_{\hat{\pi}}$ — this allows us to apply Claims 9 and 10 below:

$$\begin{aligned} v^*(s_0) &\leq V^\dagger + \frac{10\gamma \varepsilon_{\text{approx}}}{1 - \gamma} && \text{(Theorem 2)} \\ &\leq f(\hat{\lambda}, \theta) + \delta_{\mathcal{B}}(\hat{\lambda}, \hat{\theta}) + \frac{10\gamma \varepsilon_{\text{approx}}}{1 - \gamma} && \text{(Claim 9)} \\ &= \sum_{a \in \mathcal{A}} \hat{\pi}_a r_{s_0 a} + g_{\gamma \mu_{\hat{\pi}}}(\hat{\lambda}_*, \theta) + \delta_{\mathcal{B}}(\hat{\lambda}, \hat{\theta}) + \frac{10\gamma \varepsilon_{\text{approx}}}{1 - \gamma} && \text{(Lemma 7)} \\ &\leq \sum_{a \in \mathcal{A}} \hat{\pi}_a r_{s_0 a} + V_{\text{LRALP}}(\gamma \mu_{\hat{\pi}}) + \gamma C_{\mathcal{B}} + \delta_{\mathcal{B}}(\hat{\lambda}, \hat{\theta}) + \frac{10\gamma \varepsilon_{\text{approx}}}{1 - \gamma} && \text{(Claim 10)} \\ &\leq \sum_{a \in \mathcal{A}} \hat{\pi}_a r_{s_0 a} + \gamma \mu_{\hat{\pi}}^\top v^* + \gamma C_{\mathcal{B}} + \delta_{\mathcal{B}}(\hat{\lambda}, \hat{\theta}) + \frac{20\gamma \varepsilon_{\text{approx}}}{1 - \gamma} && \text{(Theorem 5)} \\ &= \sum_{a \in \mathcal{A}} \hat{\pi}_a q^*(s_0, a) + \gamma C_{\mathcal{B}} + \delta_{\mathcal{B}}(\hat{\lambda}, \hat{\theta}) + \frac{20\gamma \varepsilon_{\text{approx}}}{1 - \gamma}, \end{aligned}$$

where the last step used $\sum_a \hat{\pi}_a r_{s_0 a} + \gamma \mu_{\hat{\pi}}^\top v^* = \sum_a \hat{\pi}_a (r_{s_0 a} + \gamma \mathbf{P}_{s_0 a} v^*) = \sum_a \hat{\pi}_a q^*(s_0, a)$. Rearranging the inequality completes the proof:

$$v^*(s_0) - \sum_{a \in \mathcal{A}} \hat{\pi}_a q^*(s_0, a) \leq \frac{20\gamma \varepsilon_{\text{approx}}}{1 - \gamma} + \gamma C_{\mathcal{B}} + \delta_{\mathcal{B}}(\hat{\lambda}, \hat{\theta}). \quad \square$$

The following lemma bounds the expected duality gap of the Stochastic Mirror-Prox algorithm when applied to our setting. We defer the proof to [Appendix A.4](#).

Lemma 11 (Stochastic Mirror-Prox). *Using the constants B and C from [Theorem 3](#), define*

$$\mathcal{B} := \{\theta \in \mathbb{R}^d \mid \|\Phi_* \theta\|_2 \leq B\}, \quad (15)$$

and let $\delta_{\mathcal{B}}(\hat{\lambda}, \hat{\theta})$ be the $\mathcal{B}$ -bounded duality gap (12). Then the results of running [Algorithm 1](#) for T iterations satisfy

$$\varepsilon_{\text{opt}} := \mathbb{E}[\delta(\hat{\lambda}, \hat{\theta})] \leq \frac{14C}{\sqrt{3T}}.$$

Proof of [Theorem 3](#). First, observe that $v^*(s_0) - q^*(s_0, a) \leq 2/(1 - \gamma)$ for any action a , since all the rewards lie in $[-1, 1]$ by [Assumption 3](#). Thus, if $\varepsilon_{\text{approx}} > 1/16$ then $32\varepsilon_{\text{approx}}/(1 - \gamma) > 2/(1 - \gamma)$ and the result is trivially true. From now on, we will assume that $\varepsilon_{\text{approx}} \leq 1/16$.

To prove our result, we will combine [Lemmas 8](#) and [11](#), for which we need to show that the set $\mathcal{B}$ defined in (15) satisfies the requirements of [Lemma 8](#). Specifically, we need to show that $\mathcal{B}$ contains a feasible solution of the linear program (LRALP $_{\mu}$) with sub-optimality bounded by a constant $C_{\mathcal{B}}$. Note that the constraints of (LRALP $_{\mu}$) do not depend on μ , only the objective function, so the choice of μ does not affect feasibility.

Since (LRALP $_{\mu}$) is a relaxation of the ALP (see [Section 1.2](#)), any feasible solution of the ALP is also feasible for (LRALP $_{\mu}$). de Farias and Van Roy [14, Theorem 2] show that the ALP has a feasible solution $\Phi\theta$ that is close to v^* — more precisely

$$\|\Phi\theta - v^*\|_{\infty} \leq \frac{2\varepsilon_{\text{approx}}}{1 - \gamma}. \quad (16)$$

Since $\|v^*\|_{\infty} \leq 1/(1 - \gamma)$, we must have $\|\Phi\theta\|_{\infty} \leq (1 + 2\varepsilon_{\text{approx}})/(1 - \gamma)$. It follows that

$$\|\Phi_* \theta\|_2 \leq \sqrt{m} \|\Phi_* \theta\|_{\infty} = \sqrt{m} \|\Phi\theta\|_{\infty} \leq \frac{(1 + 2\varepsilon_{\text{approx}})\sqrt{m}}{1 - \gamma} \leq \frac{(9/8)\sqrt{m}}{1 - \gamma} = B,$$

where the first inequality is a property of the 2-norm, the next equality is thanks to [Assumption 2](#), and the last inequality is because we assumed $\varepsilon_{\text{approx}} \leq 1/16$ — this shows that $\theta \in \mathcal{B}$. We also have

$$\mu^{\top} \Phi\theta - \mu^{\top} v^* \leq \|\mu\|_1 \|\Phi\theta - v^*\|_{\infty} \leq \frac{2\varepsilon_{\text{approx}}}{1 - \gamma} \quad (\text{using (16) and } \|\mu\|_1 = 1)$$

$$|\mu^{\top} v^* - V_{\text{LRALP}}(\mu)| \leq \frac{10\varepsilon_{\text{approx}}}{1 - \gamma} \quad (\text{using [Theorem 5](#) and } \|\mu\|_1 = 1)$$

We get the value of $C_{\mathcal{B}}$ by putting these two bounds together:

$$\mu^{\top} \Phi\theta - V_{\text{LRALP}}(\mu) \leq \frac{12\varepsilon_{\text{approx}}}{1 - \gamma} =: C_{\mathcal{B}}.$$

[Lemma 8](#) applied to $\mathcal{B}$ with this value of $C_{\mathcal{B}}$ gives

$$v^*(s_0) - \sum_{a \in \mathcal{A}} \hat{\pi}(a) q^*(s_0, a) \leq \frac{32\gamma\varepsilon_{\text{approx}}}{1 - \gamma} + \delta_{\mathcal{B}}(\hat{\lambda}, \hat{\theta}).$$

Taking expectations on both sides and substituting the value of $\varepsilon_{\text{opt}} = \mathbb{E}[\delta_{\mathcal{B}}(\hat{\lambda}, \hat{\theta})]$ from [Lemma 11](#),

$$v^*(s_0) - \mathbb{E}[q^*(s_0, a)] \leq \frac{32\gamma\varepsilon_{\text{approx}}}{1 - \gamma} + \frac{14C}{\sqrt{3T}}.$$

We drop the γ factor from the leading term and plug in the value of C from the statement of [Theorem 3](#) to finish the proof. $\square$

A.4 Proof of [Lemma 11](#) — Stochastic Mirror-Prox

Lemma 11 (Stochastic Mirror-Prox). *Using the constants B and C from [Theorem 3](#), define*

$$\mathcal{B} := \{\theta \in \mathbb{R}^d \mid \|\Phi_* \theta\|_2 \leq B\}, \quad (15)$$

and let $\delta_{\mathcal{B}}(\hat{\lambda}, \hat{\theta})$ be the $\mathcal{B}$ -bounded duality gap (12). Then the results of running [Algorithm 1](#) for T iterations satisfy

$$\varepsilon_{\text{opt}} := \mathbb{E}[\delta(\hat{\lambda}, \hat{\theta})] \leq \frac{14C}{\sqrt{3T}}.$$

Throughout this section, we will use the definitions of Λ_γ and $\mathcal{B}$ from (11) and (15), respectively. We will also define the composite space $Z := \Lambda_\gamma \times \mathcal{B}$. We will use the norm $\|\lambda\|_1$ for $\lambda \in \Lambda_\gamma$, whose dual norm is $\|\cdot\|_\infty$. For $\theta \in \mathcal{B}$ we will use the norm $\|\theta\| \equiv \|\Phi\theta\|_2$ — the corresponding dual norm enjoys the convenient bound $\|\Phi_*^\top \mathbf{u}\|_* = \sup_{\|\theta\| \leq 1} \mathbf{u}^\top \Phi_* \theta \leq \|\mathbf{u}\|_2 \leq \|\mathbf{u}\|_1$ for any vector $\mathbf{u}$.⁴ The last inequality is due a general property of p -norms: $\|\mathbf{u}\|_p \leq \|\mathbf{u}\|_q$ whenever $\infty \geq p \geq q \geq 1$.

A.4.1 Lipschitz Constants

Our first step will be to bound the Lipschitz constants associated with f , the objective function of (CoreLP). In other words, we are looking for bounds on $\|f_\lambda(\theta)\|_\infty$ and $\|f_\theta(\lambda)\|_*$.

First, for any $\theta \in \mathcal{B}$ we have

$$\begin{aligned} \|f_\lambda(\theta)\|_\infty &= \|\mathbf{W}\mathbf{r} + \mathbf{W}(\gamma\mathbf{P} - \mathbf{E})\Phi\theta\|_\infty \\ &\leq \|\mathbf{W}\mathbf{r}\|_\infty + \|\mathbf{W}(\gamma\mathbf{P} - \mathbf{E})\Phi\theta\|_\infty && \text{(by the triangle inequality)} \\ &\leq 1 + \max_i \|\mathbf{W}_i(\gamma\mathbf{P} - \mathbf{E})\|_1 \|\Phi\theta\|_\infty. && \text{(by definition of } \|\cdot\|_\infty \text{ and Hölder's inequality)} \end{aligned}$$

Now, by the property of norms, $\|\Phi\theta\|_\infty \leq \|\Phi\theta\|_2 \leq B$. Secondly, $\mathbf{W}_i\mathbf{P}$ and $\mathbf{W}_i\mathbf{E}$ are probability distributions, so $\|\mathbf{W}_i(\gamma\mathbf{P} - \mathbf{E})\|_1 \leq \gamma\|\mathbf{W}_i\mathbf{P}\|_1 + \|\mathbf{W}_i\mathbf{E}\|_1 = 1 + \gamma$ and

$$\|f_\lambda(\theta)\|_\infty \leq 1 + (1 + \gamma)B \leq 2B. \quad (\text{since } B \geq 1/(1 - \gamma))$$

For the other gradient, we use the bound on dual norms mentioned above:

$$\begin{aligned} \|f_\theta(\lambda)\|_* &= \|(\varphi_{s_0}^\top + \lambda^\top \mathbf{W}(\gamma\mathbf{P} - \mathbf{E}))\Phi\|_* \\ &\leq \|\mathbf{e}_{s_0}^\top\|_1 + \|\lambda^\top \mathbf{W}(\gamma\mathbf{P} - \mathbf{E})\|_1 \\ &\leq 1 + \frac{1 + \gamma}{1 - \gamma} = \frac{2}{1 - \gamma}, \end{aligned}$$

where the last inequality uses the fact that $(1 - \gamma)\lambda$ is a probability distribution, as are the rows of $\mathbf{W}$, $\mathbf{P}$, and $\mathbf{E}$.

A.4.2 Gradient Estimator Variance

Next, we will bound the variance in the stochastic estimators $\hat{f}_\lambda(\theta)$ and $\hat{f}_\theta(\lambda)$ defined in (6) and (7), respectively, compared to the true gradients $f_\lambda(\theta)$ and $f_\theta(\lambda)$ defined in (4) and (5), respectively.

First, we bound $\mathbb{E}[\|\hat{f}_\lambda(\theta) - f_\lambda(\theta)\|_\infty^2]$ for any $\theta \in \mathcal{B}$ by bounding its components. For any state $s \in \mathcal{S}_+$, action $a \in \mathcal{A}$, and reward $\hat{r}$, we have

$$\begin{aligned} &|[\hat{f}_\lambda(\theta)]_{sa} - [f_\lambda(\theta)]_{sa}| \\ &= |(\hat{r} + \gamma\varphi_{s'}^\top \theta - \varphi_s^\top \theta) - (r_{sa} + \gamma\mathbf{P}_{sa}\Phi\theta - \varphi_s^\top \theta)| && \text{(for some random } s' \sim \mathbf{P}_{sa}) \\ &\leq |\hat{r} - r_{sa}| + \gamma|\varphi_{s'}^\top \theta - \mathbf{P}_{sa}\Phi\theta| && \text{(by the triangle inequality)} \\ &\leq 2 + \gamma|(\mathbf{e}_{s'}^\top - \mathbf{P}_{sa})\Phi\theta| && \text{(bounded rewards)} \\ &\leq 2 + \gamma\|\mathbf{e}_{s'}^\top - \mathbf{P}_{sa}\|_1 \|\Phi\theta\|_\infty && \text{(using Hölder's inequality)} \\ &\leq 2 + 2\gamma B && \text{(since } \|\Phi\theta\|_\infty \leq \|\Phi\theta\|_2 \leq B) \\ &\leq 2B. && \text{(since } B \geq 1/(1 - \gamma)) \end{aligned}$$

It follows that $\|\hat{f}_\lambda(\theta) - f_\lambda(\theta)\|_\infty^2 \leq (2B)^2$, and the same bound must hold for the expectation.

We will now bound the other gradient, using the following property of Euclidean norms: for any vector-valued random variable $\mathbf{h}$ with mean $\bar{\mathbf{h}}$, $\mathbb{E}[\|\mathbf{h} - \bar{\mathbf{h}}\|^2] = \mathbb{E}[\|\mathbf{h}\|^2] - \|\bar{\mathbf{h}}\|^2 \leq \mathbb{E}[\|\mathbf{h}\|^2]$. Then, for a random choice of state $s \in \mathcal{S}_+$, action $a \in \mathcal{A}$, and next state $s' \sim \mathbf{P}_{sa}$:

$$\begin{aligned} \mathbb{E}[\|\hat{f}_\theta(\lambda) - f_\theta(\lambda)\|_*^2] &= \mathbb{E}\left[\|(\varphi_0^\top + \|\lambda\|_1(\gamma\varphi_{s'} - \varphi_s)) - (\varphi_0^\top + \lambda^\top \mathbf{W}(\gamma\mathbf{P} - \mathbf{E})\Phi)\|_*^2\right] \\ &= \left(\frac{1}{1 - \gamma}\right)^2 \mathbb{E}\left[\|(\gamma\varphi_{s'} - \varphi_s) - (\lambda/\|\lambda\|_1)^\top \mathbf{W}(\gamma\mathbf{P} - \mathbf{E})\Phi\|_*^2\right] \end{aligned}$$

⁴More generally, $\|\xi\|_* = \inf\{\|\eta\|_2 \mid \xi^\top = \eta^\top \Phi_*\}$, which is non-zero when $\xi \neq \mathbf{0}$ and Φ_* has full column rank.

Now, since $(s, a) \sim \lambda / \|\lambda\|_1$ and $s' \sim \mathbf{P}_{sa}$, we have $\mathbb{E}[\gamma \boldsymbol{\varphi}_{s'} - \boldsymbol{\varphi}_s] = (\lambda / \|\lambda\|_1)^\top \mathbf{W}(\gamma \mathbf{P} - \mathbf{E}) \boldsymbol{\Phi}$, so by the above property of variance for vector-valued random variables,

$$\begin{aligned} &\leq \left(\frac{1}{1-\gamma} \right)^2 \mathbb{E}[\|\gamma \boldsymbol{\varphi}_{s'} - \boldsymbol{\varphi}_s\|_*^2] \\ &= \left(\frac{1}{1-\gamma} \right)^2 \mathbb{E}[\|(\gamma \mathbf{e}_{s'} - \mathbf{e}_s) \boldsymbol{\Phi}\|_*^2] \\ &\leq \left(\frac{1+\gamma}{1-\gamma} \right)^2 \leq \left(\frac{2}{1-\gamma} \right)^2. \end{aligned}$$

A.4.3 Distance-Generating Functions

The Stochastic Mirror-Prox algorithm requires strongly convex *distance-generating functions* for Λ_γ and $\mathcal{B}$ with respect to their respective norms. A function $\omega : \mathcal{X} \rightarrow \mathbb{R}$ (with domain $\mathcal{X} \subset \mathbb{R}^n$) is said to be σ -strongly convex (where $\sigma > 0$ is called the modulus of convexity) with respect to a norm $\|\cdot\|$ on $\mathcal{X}$ if any of the following conditions hold for all $\mathbf{x}, \mathbf{y} \in \mathcal{X}$

- (i) For all $\alpha \in [0, 1]$, $\alpha \omega(\mathbf{x}) + (1-\alpha) \omega(\mathbf{y}) \geq \omega(\alpha \mathbf{x} + (1-\alpha) \mathbf{y}) + \sigma \alpha(1-\alpha) \|\mathbf{x} - \mathbf{y}\|^2/2$.
- (ii) ω is convex and $\omega(\mathbf{x}) \geq \omega(\mathbf{y}) + \langle \nabla \omega(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle + \sigma \|\mathbf{x} - \mathbf{y}\|^2/2$.
- (iii) $\mathcal{X}$ is convex and $\langle \nabla \omega(\mathbf{x}) - \nabla \omega(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle \geq \sigma \|\mathbf{x} - \mathbf{y}\|^2$.

Condition (i) is the definition of strong convexity; note that it reduces to convexity when $\sigma = 0$. Conditions (ii) and (iii) are equivalent to the definition under appropriate differentiability conditions on ω that hold in our setting and when $\mathbf{x}, \mathbf{y}$ are in the interior of $\mathcal{X}$; see Yu [45] for details. Juditsky et al. [22] uses “strongly convex” to mean that a function is 1-strongly convex according to condition (iii).

Define the *divergence function*:

$$D_\omega(\mathbf{x}, \mathbf{y}) := \omega(\mathbf{x}) - \omega(\mathbf{y}) - \langle \nabla \omega(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle.$$

When ω is convex, one can see that D_ω is always non-negative, and by condition (ii) the σ -strong convexity of ω is equivalent to $D_\omega(\mathbf{x}, \mathbf{y}) \geq \sigma \|\mathbf{x} - \mathbf{y}\|^2/2$. We will use this equivalence to establish the strong convexity of our distance-generating functions below. An important operation related to distance-generating functions is the *proximal projection* onto $\mathcal{X}$ with respect to ω :

$$\Pi_\omega(\mathbf{x}, \boldsymbol{\xi}) := \arg \min_{\mathbf{y} \in \mathcal{X}} D_\omega(\mathbf{y}, \mathbf{x}) + \langle \boldsymbol{\xi}, \mathbf{y} \rangle, \quad \boldsymbol{\xi} \in \mathbb{R}^n.$$

The *center* of $\mathcal{X}$ with respect to ω is defined as $\mathbf{x}_0 := \arg \min_{\mathbf{x} \in \mathcal{X}} \omega(\mathbf{x})$ and the *diameter* of $\mathcal{X}$ is $\Omega_{\mathcal{X}} := \sup_{\mathbf{y} \in \mathcal{X}} \sqrt{2D_\omega(\mathbf{y}, \mathbf{x}_0)}$.

A strongly convex distance-generating function can be thought of as a generalization of the squared norm $\|\cdot\|^2$ — the corresponding divergence generalizes the squared distance function $\|\mathbf{x} - \mathbf{y}\|^2$; unlike the squared distance, however, the divergence may not be symmetric. Indeed, when $\|\cdot\|$ is an Euclidean norm, and *only* for such norms, the function $\omega(\mathbf{x}) = \|\mathbf{x}\|^2/2$ is 1-strongly convex [45, Proposition 2]. In this special case, the divergence is $D_\omega(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\|^2/2$ and the proximal projection is simply the Euclidean projection: $\Pi_\omega(\mathbf{x}, \boldsymbol{\xi}) = \arg \min_{\mathbf{y} \in \mathcal{X}} \|\mathbf{x} + \boldsymbol{\xi} - \mathbf{y}\|$.

Thus, since the domain $\mathcal{B}$ of the primal variables $\boldsymbol{\theta}$ is equipped with the Euclidean norm $\|\boldsymbol{\theta}\| \equiv \|\boldsymbol{\Phi} \boldsymbol{\theta}\|_2$, we will use the 1-strongly convex distance-generating function $\omega_{\mathcal{B}}(\boldsymbol{\theta}) = \|\boldsymbol{\theta}\|^2/2$. Since $\mathcal{B}$ is the Euclidean ball under this norm, the center of $\mathcal{B}$ is $\boldsymbol{\theta}_0 = \mathbf{0}$ and its “diameter” (actually the radius, in this case) is $\Omega_{\mathcal{B}} = B$; $\boldsymbol{\theta}_0$ is used as the initial value of $\boldsymbol{\theta}$ in Algorithm 1. The proximal projection is

$$\Pi_{\mathcal{B}}(\boldsymbol{\theta}, \boldsymbol{\xi}) = \arg \min_{\boldsymbol{\theta}' \in \mathcal{B}} \|\boldsymbol{\theta} + \boldsymbol{\xi} - \boldsymbol{\theta}'\| = \frac{\boldsymbol{\theta} + \boldsymbol{\xi}}{\max\{1, \|\boldsymbol{\theta} + \boldsymbol{\xi}\|/B\}}.$$

For the dual variables $\boldsymbol{\lambda} \in \Lambda_\gamma$, our distance-generating function is a modification of the *unnormalized negentropy* $h(\boldsymbol{\lambda}) = \sum_i \lambda_i (\log \lambda_i - 1)$. It is well-known that this function is 1-strongly convex on the set $\{\boldsymbol{\lambda} \geq \mathbf{0} \mid \|\boldsymbol{\lambda}\|_1 \leq 1\}$ [e.g., 45, Theorem 5]. To achieve 1-strong convexity on Λ_γ (where $\|\boldsymbol{\lambda}\|_1 = 1/(1-\gamma) > 1$), we use a modified form of this function:

$$h_\gamma(\boldsymbol{\lambda}) := \frac{h((1-\gamma)\boldsymbol{\lambda})}{(1-\gamma)^2}.$$

It follows that $\nabla h_\gamma(\lambda) = [\nabla h((1-\gamma)\lambda)]/(1-\gamma)$. Thus, defining $D_\Lambda(\lambda', \lambda) := D_{h_\gamma}(\lambda', \lambda)$, we have

$$\begin{aligned} D_\Lambda(\lambda', \lambda) &= \left(\frac{1}{1-\gamma} \right) \left[\frac{h((1-\gamma)\lambda')}{1-\gamma} - \frac{h((1-\gamma)\lambda)}{1-\gamma} - \langle \nabla h((1-\gamma)\lambda), \lambda' - \lambda \rangle \right] \\ &= \left(\frac{1}{1-\gamma} \right) \left[\frac{h((1-\gamma)\lambda')}{1-\gamma} - \frac{h((1-\gamma)\lambda)}{1-\gamma} - \frac{\langle \nabla h((1-\gamma)\lambda), (1-\gamma)\lambda' - (1-\gamma)\lambda \rangle}{1-\gamma} \right] \end{aligned}$$

Now, since $(1-\gamma)\lambda, (1-\gamma)\lambda' \in \Delta_{\mathcal{S}_* \times \mathcal{A}}$ and h is strongly convex on this set, we have

$$D_\Lambda(\lambda', \lambda) \geq \left(\frac{1}{1-\gamma} \right)^2 \frac{\|(1-\gamma)\lambda - (1-\gamma)\lambda'\|_1^2}{2} = \frac{\|\lambda - \lambda'\|_1^2}{2}.$$

Thus we have shown that h_γ is 1-strongly convex on Λ_γ . By the properties of the negentropy function, we can verify that $h_\gamma(\lambda)$ is minimized for $\lambda_0 = \mathbf{1}_{\mathcal{A}}/A \oplus \gamma \mathbf{1}_{mA}/(1-\gamma)mA$, i.e. the uniform distribution over actions concatenated with the scaled uniform distribution over state-action pairs in $\mathcal{S}_* \times \mathcal{A}$ — this value is used as the initializer in [Algorithm 1](#). Conversely, $h_\gamma(\lambda)$ is maximized for $\tilde{\lambda} = \mathbf{e}_a \oplus \gamma \mathbf{e}_{s'a'}/(1-\gamma)$, i.e. when $\tilde{\lambda}$ is concentrated on (s_0, a) for some $a \in \mathcal{A}$ and some $(s', a') \in \mathcal{S}_* \times \mathcal{A}$. We can verify through a short calculation that

$$\begin{aligned} D_\Lambda(\tilde{\lambda}, \lambda_0) &= \frac{\ell}{(1-\gamma)^2}, & \text{where } \ell &:= \log A + \gamma \log m \\ \Omega_\Lambda &= \sqrt{2D_\Lambda(\tilde{\lambda}, \lambda_0)} = \frac{\sqrt{2\ell}}{1-\gamma}. \end{aligned}$$

Finally, the proximal projection onto Λ_γ with respect to h_γ is

$$\Pi_\Lambda(\lambda, \rho) = \frac{\tilde{\lambda}_0}{\|\tilde{\lambda}_0\|_1} \oplus \frac{\gamma \tilde{\lambda}_*}{(1-\gamma)\|\tilde{\lambda}_*\|_1}, \quad \text{where } \tilde{\lambda} := \exp(\log \lambda + \rho),$$

where $\tilde{\lambda}_{s_0} := [\tilde{\lambda}_{s_0 a}]_{a \in \mathcal{A}}$ and $\tilde{\lambda}_* := [\tilde{\lambda}_{s a}]_{s \in \mathcal{S}_*, a \in \mathcal{A}}$, so that $\tilde{\lambda} = \tilde{\lambda}_{s_0} \oplus \tilde{\lambda}_*$.

A.4.4 The Composite Space

We will now gather together the preceding results and use them to construct a norm and distance-generating function on the composite optimization domain $Z = \Lambda_\gamma \times \mathcal{B} \subset \mathbb{R}^{(1+m)A} \oplus \mathbb{R}^d$. We closely follow the construction of Juditsky et al. [22, §4.2]. First, we define the squared norm:

$$\|\lambda \oplus \theta\|^2 := \Omega_\Lambda^{-2} \|\lambda\|_1^2 + \Omega_{\mathcal{B}}^{-2} \|\theta\|^2 = \left(\frac{1-\gamma}{\sqrt{2\ell}} \|\lambda\|_1 \right)^2 + \left(\frac{1}{B} \|\theta\| \right)^2.$$

The corresponding squared dual norm is

$$\|\rho \oplus \xi\|_*^2 := \Omega_\Lambda^2 \|\lambda\|_\infty^2 + \Omega_{\mathcal{B}}^2 \|\theta\|_*^2 = \left(\frac{\sqrt{2\ell}}{1-\gamma} \|\lambda\|_\infty \right)^2 + (B \|\theta\|)^2.$$

Define the operator $F(\lambda, \theta) := -f_\lambda(\theta) \oplus f_\theta(\lambda)$. Its Lipschitz constant with respect to this norm is

$$\begin{aligned} \|F(\lambda, \theta)\|_* &= \sqrt{\frac{2\ell \|f_\lambda(\theta)\|_\infty^2}{(1-\gamma)^2} + B^2 \|f_\theta(\lambda)\|_*^2} \\ &\leq \sqrt{\frac{8\ell B^2}{(1-\gamma)^2} + \left(\frac{2B}{1-\gamma} \right)^2} \\ &= \frac{2B\sqrt{1+2\ell}}{1-\gamma} = C. \end{aligned}$$

Similarly, define the estimator $\hat{F}(\lambda, \theta) := -\hat{f}_\lambda(\theta) \oplus \hat{f}_\theta(\lambda)$. Its variance enjoys

$$\mathbb{E}[\|\hat{F}(\lambda, \theta) - F(\lambda, \theta)\|_*^2] \leq C^2.$$

We do not repeat the calculation because it is identical to the previous one, since the variances of our estimators have the same bounds as their squared Lipschitz constants.

Finally, we construct the composite distance-generating function:

$$\omega_Z(\lambda, \theta) := \frac{(1 - \gamma)^2 h_\gamma(\lambda)}{2\ell} + \frac{\|\theta\|^2}{2B^2}.$$

We can verify that this function is 1-strongly convex on Z with respect to the norm defined above, and that the diameter of Z under this function is $\Omega \leq \sqrt{2}$.

Proof of Lemma 11. We apply the result of Juditsky et al. [22, Corollary 1] to our setting, where the Lipschitz constant of F is C , the variance of $\hat{F}$ is C^2 , and the diameter of Z is $\Omega \leq \sqrt{2}$. Then the result tells us that a suitable learning rate is $\eta = C^{-1}\sqrt{2/7T}$, as specified in Theorem 3, and the resulting bound on the expected duality gap after T iterations is

$$\varepsilon_{\text{opt}} := \mathbb{E}[\delta_{\mathcal{B}}(\hat{\lambda}, \hat{\theta})] \leq \frac{14C}{\sqrt{3T}}. \quad \square$$