

ALGEBRAIC OPTIMIZATION OF SEQUENTIAL DECISION PROBLEMS

MAREIKE DRESSLER, MARINA GARROTE-LÓPEZ, GUIDO MONTÚFAR, JOHANNES MÜLLER,
AND KEMAL ROSE

ABSTRACT. We study the optimization of the expected long-term reward in finite partially observable Markov decision processes over the set of stationary stochastic policies. In the case of deterministic observations, also known as state aggregation, the problem is equivalent to optimizing a linear objective subject to quadratic constraints. We characterize the feasible set of this problem as the intersection of a product of affine varieties of rank one matrices and a polytope. Based on this description, we obtain bounds on the number of critical points of the optimization problem. Finally, we conduct experiments in which we solve the KKT equations or the Lagrange equations over different boundary components of the feasible set, and compare the result to the theoretical bounds and to other constrained optimization methods.

1. INTRODUCTION

Solving sequential decision problems has a long-standing history in computer science, economics, mathematics, and statistics [Bel57, How60, Che68]. Such problems include the optimal control of robots, machine maintenance, search problems, and inventory problems, which can be formulated in continuous or discrete time, space, and control variables [Whi88, Bel66]. A sequential decision problem is particularly challenging if only partial information about the true state of the system is available to the acting agent.

Partially observable Markov decision processes (POMDPs) offer a model for sequential decision-making under state uncertainty. Here, at every time step the agent selects an action and receives an instantaneous reward depending on the selected action and the current state, which in turn influence the state at the next time step. However, the agent selects its actions based on observations that might not fully reveal the underlying state. We study stochastic action selection mechanisms that do not depend on the prior history of observations but only on the current observation, which are known as memoryless, stationary, or reactive policies.

A common measure for the performance of a policy is the expectation of the instantaneous rewards accumulated over time and discounted into the future. We will refer to this measure simply as the reward function. Identifying a policy that maximizes the reward is a challenging task since it is a nonconcave function that can exhibit non global strict local optima [BR19]. Indeed it has been shown that this optimization problem is NP-hard in general [VLB12]. A common approach are local optimization procedures, such as policy gradient methods [SMSM99, AYA18]. Whereas global optimality guarantees for gradient methods in fully observable systems have been given in [BR19], for general POMDPs we do not have such guarantees.

2020 *Mathematics Subject Classification.* 62R01, 90C23, 90C40.

Key words and phrases. Partially observable Markov decision process, algebraic degree, polynomial optimization, state aggregation, state-action frequencies.

Various approaches have been suggested to study the geometry underlying the optimization problem. A classic line of works has established that in the fully observable case (where the observation fully identifies the underlying state), the optimization problem is equivalent to a linear program over a polytope of feasible state-action frequencies [Der70, Kal94]. These studies have been complemented by the characterization of the set of feasible value functions of a Markov decision process as a finite union of polytopes [DTLR⁺19, WKZ⁺22, WDL22]. However, for partially observable systems the geometry of the reward optimization problem is more complex. The problem can be formulated as a quadratically constrained linear program with the policy and the value function as search variables [ABZ06]. More recently, the set of feasible state-action frequencies was described as a union of convex sets in [MGZA15] and as a semialgebraic set in [MM22a], who also provided a method for computing the polynomial constraints. This yields a polynomially constrained linear program with the state-action frequencies as search variables. The possible advantages of taking this constrained optimization perspective in state-action space were recently studied in [MM22b] using interior point methods.

Related approaches have been proposed in other settings as well. In continuous time and space, a convex relaxation of linear quadratic control problems based on state-action frequencies has been proposed and studied in [LHPT08]. In [Ney03] the graphs of different stochastic games are described as semialgebraic sets, where (generalized) Nash equilibria including a convex relaxation for their computation have been studied with algebraic tools in [NT21, PS22].

In this paper we study finite POMDPs and build upon the recent work [MM22a], which expresses reward optimization in POMDPs with memoryless stochastic policies as a linear program with polynomial constraints; that is, we are concerned with the optimization of a linear function over a nonconvex semialgebraic set. We focus on the case of deterministic observations, where, as we will see, the polynomial constraints are quadratic and can be written as a sum of certain 2×2 minors (Theorem 6). By investigating the geometry of the semialgebraic set, we determine upper bounds on the number of (complex) critical points of the reward optimization problem, i.e., its algebraic degree (Theorem 11). We then provide a computational method that solves the optimization problem by computing the critical points via the Karush-Kuhn-Tucker conditions, whereby we identify ways to reduce the combinatorial complexity of the problem by focusing on relevant boundary components (Theorem 9, [MR17]). We implement this approach using numerical algebra methods [BT18] that automatically certify the correctness of the results [BRT21]. We use a convex relaxation of the polynomial problem to certify the global optimality of the results. Moreover, we observe that in specific instances this numerical algebraic approach leads to superior results than two commonly used optimization methods. Finally, we compare the number of critical points obtained in numerical experiments with our theoretical bounds.

The paper is organized as follows. In Section 2, we introduce partially observable Markov decision processes and related notation. In Section 3, we describe the geometry of the feasible set and its defining (in)equalities for the reward optimization problem in POMDPs with deterministic observations. In Section 4, we provide an upper bound on the number of critical points for the problem we are considering. Finally, in Section 5, we use the description of reward maximization as a constrained polynomial optimization problem to numerically solve the critical equations.

Notation: For a finite set \mathcal{X} we denote the free linear space over \mathcal{X} by $\mathbb{R}^{\mathcal{X}} = \{f: \mathcal{X} \rightarrow \mathbb{R}\}$ and the simplex of probability distributions over \mathcal{X} as $\Delta_{\mathcal{X}} = \{\mu \in \mathbb{R}^{\mathcal{X}} : \sum_x \mu_x = 1 \text{ and } \mu \geq 0\}$. The *conditional probability polytope* consisting of all column-stochastic matrices¹ in $\mathbb{R}^{\mathcal{Y} \times \mathcal{X}}$ is the product $\Delta_{\mathcal{Y}}^{\mathcal{X}} = \Delta_{\mathcal{Y}} \times \cdots \times \Delta_{\mathcal{Y}}$. We call the elements of this set *conditional probability distributions* or *Markov kernels* from \mathcal{X} to \mathcal{Y} . Given a Markov kernel $Q \in \Delta_{\mathcal{Y}}^{\mathcal{X}}$, the conditional probability $Q(y|x)$ is the entry Q_{yx} . Note that a composition of Markov kernels is matrix multiplication. For a probability distribution $p \in \Delta_{\mathcal{X}}$ and a Markov kernel $Q \in \Delta_{\mathcal{Y}}^{\mathcal{X}}$ we denote their composition into a joint probability distribution by $p * Q \in \Delta_{\mathcal{X} \times \mathcal{Y}}$ and define it as $p * Q = \text{diag}(p)Q^T$, that is, with entries $(p * Q)(x, y) := p(x)Q(y|x)$. For a subset $A \subseteq \mathcal{X}$ we denote the complement $\mathcal{X} \setminus A$ of A in \mathcal{X} by A^c .

2. PARTIALLY OBSERVABLE MARKOV DECISION PROCESSES

Partially observable Markov decision processes provide a powerful model to describe sequential decision making problems with state uncertainty.

Definition 1. A finite *partially observable Markov decision process* or shortly *POMDP* is a tuple $(\mathcal{S}, \mathcal{O}, \mathcal{A}, \alpha, \beta, r)$, where \mathcal{S} , \mathcal{O} , and \mathcal{A} are finite sets called the *state*, *observation*, and *action space* respectively and $\alpha \in \Delta_{\mathcal{S}}^{\mathcal{S} \times \mathcal{A}}$ and $\beta \in \Delta_{\mathcal{O}}^{\mathcal{S}}$ are Markov kernels, which we call the *transition* and *observation kernel* respectively. Furthermore, we consider an *instantaneous reward vector* $r \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$. We denote the cardinalities of \mathcal{S} , \mathcal{A} , and \mathcal{O} by $n_{\mathcal{S}}$, $n_{\mathcal{A}}$, and $n_{\mathcal{O}}$.

From a modeling perspective, $\alpha(s'|s, a)$ is the probability of transitioning from state s to state s' upon taking action a , and $\beta(o|s)$ is the probability of making the observation o if the system is in state s . The entry r_{sa} corresponds to an instantaneous reward received upon selecting action a in state s .

A *(memoryless stochastic) policy* is a column-stochastic matrix $\pi \in \Delta_{\mathcal{A}}$ from the set of observations to the set of actions. The entry $\pi(a|o)$ is the probability with which action $a \in \mathcal{A}$ is selected given the observation $o \in \mathcal{O}$. A policy can be interpreted as a randomized decision rule that encodes which action should be taken, based on the current observation. Every policy $\pi \in \Delta_{\mathcal{A}}$ defines a transition kernel $P_{\pi \circ \beta} \in \Delta_{\mathcal{S} \times \mathcal{A}}$ with entries

$$P_{\pi \circ \beta}(s', a'|s, a) := \alpha(s'|s, a)(\pi \circ \beta)(a'|s'),$$

which we call the *state-action transition kernel* associated with $\pi \circ \beta$ and α . Given an initial distribution, the state-action transition kernel $P_{\pi \circ \beta}$ defines a Markov process on the state-action space $\mathcal{S} \times \mathcal{A}$.

One is particularly interested in the probability that the Markov process assigns to any given state-action pair, averaged over time, whereby it is convenient to discount events at larger times t by weighting them by $(1 - \gamma)\gamma^t$ for a *discount factor* $\gamma \in (0, 1)$. Given an initial state distribution $\mu \in \Delta_{\mathcal{S}}$ and a discount factor $\gamma \in (0, 1)$, one thus defines the *(discounted) state-action frequency* associated with policy $\pi \in \Delta_{\mathcal{A}}$ as

$$\eta^{\pi} := (1 - \gamma) \sum_{t \geq 0} \gamma^t P_{\pi \circ \beta}^t(\mu * (\pi \circ \beta)) = (1 - \gamma)(I - \gamma P_{\pi \circ \beta})^{-1}(\mu * (\pi \circ \beta)),$$

¹We choose to work with column-stochastic rather than row-stochastic matrices to have $Q_{yx} = Q(y|x)$, which makes composition of two Markov kernels $Q_1 \circ Q_2$ equivalent to matrix multiplication $Q_1 Q_2$.

where I is the identity matrix; see [Der70, Kal94]. We further define the map

$$(1) \quad \begin{aligned} \Phi: \Delta_{\mathcal{A}}^{\mathcal{O}} &\rightarrow \Delta_{\mathcal{S} \times \mathcal{A}} \\ \pi &\mapsto \eta^{\pi} = (1 - \gamma)(I - \gamma P_{\pi \circ \beta})^{-1}(\mu * (\pi \circ \beta)). \end{aligned}$$

Elementary calculations show $\eta^{\pi}(a|s) = (\pi \circ \beta)(a|s)$. We denote the state-marginal of η^{π} by $\rho_s^{\pi} = \sum_{a \in \mathcal{A}} \eta_{sa}^{\pi}$ and refer to it as the *state frequency*. By definition of conditional probability distributions it holds that

$$(2) \quad \eta_{sa}^{\pi} = \eta^{\pi}(a|s)\rho_s^{\pi} = (\pi \circ \beta)(a|s)\rho_s^{\pi}.$$

Finally, as a measure for the performance of policies, we introduce the *reward function*²:

$$(3) \quad R(\pi) := \sum_{s \in \mathcal{S}, a \in \mathcal{A}} r_{sa} \Phi(\pi)_{sa} = \langle r, \Phi(\pi) \rangle_{\mathcal{S} \times \mathcal{A}}.$$

The reward function R is a widely used criterion to evaluate the performance of a policy. It is equal to the expected value $\mathbb{E}[(1 - \gamma) \sum_{t \geq 0} \gamma^t r(s_t, a_t)]$ of the (discounted) accumulated instantaneous rewards along state-action trajectories distributed according to the Markov process with transition kernel $P_{\pi \circ \beta}$ and initial state-action distribution $\mu * (\pi \circ \beta)$. We refer to standard textbooks for an in-depth discussion [How60, Der70, Put14].

We consider the following *reward optimization problem (ROP)*, which is the standard problem in (discounted) Markov decision processes:

$$(ROP) \quad \text{maximize } R(\pi) \quad \text{subject to } \pi \in \Delta_{\mathcal{A}}^{\mathcal{O}}.$$

In this work, we focus on **deterministic observations** $\beta \in \Delta_{\mathcal{O}}^{\mathcal{S}} \cap \{0, 1\}^{\mathcal{O} \times \mathcal{S}}$, where we can identify the observation kernel with a deterministic mapping $g_{\beta}: \mathcal{S} \rightarrow \mathcal{O}$. We denote the fibers of g_{β} by $S_o := \{s \in \mathcal{S} : g_{\beta}(s) = o\}$ and their cardinality by $d_o := |S_o|$. Note that the fibers S_o are a disjoint partition of the states \mathcal{S} and hence $(d_o)_{o \in \mathcal{O}}$ is a partition of $n_{\mathcal{S}}$, i.e., $\sum_{o \in \mathcal{O}} d_o = n_{\mathcal{S}}$. This special type of partial observability is known in the literature as state-aggregation.

Example 2. Let $(\mathcal{S}, \mathcal{O}, \mathcal{A}, \alpha, \beta, r)$ be a POMDP with state space $\mathcal{S} = \{s_1, s_2, s_3\}$, action space $\mathcal{A} = \{a_1, a_2\}$, and observation space $\mathcal{O} = \{o_1, o_2\}$. We consider the (deterministic) transitions depicted in Figure 1, which correspond to the column stochastic matrix

$$\alpha = \begin{matrix} & s_1, a_1 & s_1, a_2 & s_2, a_1 & s_2, a_2 & s_3, a_1 & s_3, a_2 \\ \begin{matrix} s_1 \\ s_2 \\ s_3 \end{matrix} & \left(\begin{matrix} 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 & 1 & 0 \end{matrix} \right) \end{matrix} \in \Delta_{\mathcal{S}}^{\mathcal{A} \times \mathcal{A}}.$$

Assume the agent cannot distinguish the states s_1 and s_2 , so that the observation kernel is

$$\beta = \begin{matrix} & s_1 & s_2 & s_3 \\ \begin{matrix} o_1 \\ o_2 \end{matrix} & \left(\begin{matrix} 1 & 1 & 0 \\ 0 & 0 & 1 \end{matrix} \right) \end{matrix} \in \Delta_{\mathcal{O}}^{\mathcal{S}}.$$

Hence, we are optimizing over stochastic matrices

$$\pi = \begin{matrix} & o_1 & o_2 \\ \begin{matrix} a_1 \\ a_2 \end{matrix} & \left(\begin{matrix} \pi_{a_1 o_1} & \pi_{a_1 o_2} \\ \pi_{a_2 o_1} & \pi_{a_2 o_2} \end{matrix} \right) \end{matrix} \in \Delta_{\mathcal{A}}^{\mathcal{O}}.$$

²More precisely, this is the infinite-horizon expected discounted reward function.

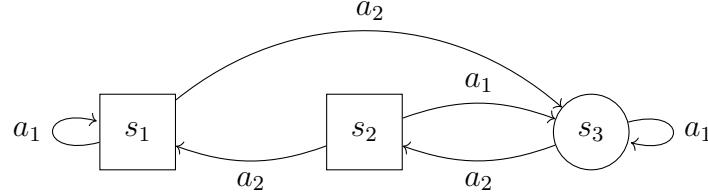


FIGURE 1. Transition graph of Example 2; states s_1, s_2 lead to observation o_1 , and s_3 leads to observation o_2 .

Further, we consider a uniform initial distribution $\mu \in \Delta_S$ and a discount factor $\gamma = 1/2$. Finally, let us assume the instantaneous reward vector is $r(s, a) = \delta_{s_1 s}$, which corresponds to a reward of +1 obtained in state s_1 . Combining the Neumann series with Cramer's rule (see [MM22a]) one sees that the reward function R is a rational function with the explicit expression $R(\pi) = \frac{f(\pi)}{2g(\pi)} - \frac{1}{2}$, where f and g are determinantal polynomials given by

$$(4) \quad \begin{aligned} f(\pi) &= \det \begin{pmatrix} -0.49\pi_{a_1 o_1} + 0.5\pi_{a_2 o_1} + 1 & -0.99\pi_{a_2 o_2} \\ -0.5\pi_{a_1 o_1} - 0.49\pi_{a_2 o_1} & -0.99\pi_{a_1 o_2} + 1 \end{pmatrix} \\ &= \pi_{a_1 o_1}^2 \pi_{a_2 o_2} - 2\pi_{a_1 o_1} \pi_{a_2 o_1} \pi_{a_1 o_2} - 2\pi_{a_2 o_1}^2 \pi_{a_1 o_2} - \pi_{a_2 o_1}^2 \pi_{a_2 o_2} + 4\pi_{a_1 o_1} \pi_{a_2 o_1} \\ &\quad + 2\pi_{a_1 o_1} \pi_{a_1 o_2} - 6\pi_{a_1 o_1} \pi_{a_2 o_2} + 4\pi_{a_2 o_1}^2 - 4\pi_{a_2 o_1} \pi_{a_1 o_2} - 4\pi_{a_1 o_1} + 8\pi_{a_2 o_1} - 12\pi_{a_1 o_2} + 24 \end{aligned}$$

and

$$(5) \quad \begin{aligned} g(\pi) &= \det \begin{pmatrix} -0.99\pi_{a_1 o_1} + 1 & -0.99\pi_{a_2 o_2} \\ -0.99\pi_{a_2 o_1} & -0.99\pi_{a_1 o_2} + 1 \end{pmatrix} \\ &= 3\pi_{a_1 o_1}^2 \pi_{a_2 o_2} - 3\pi_{a_2 o_1}^2 \pi_{a_2 o_2} + 6\pi_{a_1 o_1} \pi_{a_1 o_2} - 6\pi_{a_1 o_1} \pi_{a_2 o_2} - 12\pi_{a_1 o_1} - 12\pi_{a_1 o_2} + 24. \end{aligned}$$

The reward function is to be optimized over the observation policy, that is, we have

$$(6) \quad \text{maximize } R(\pi) \quad \text{subject to } \begin{cases} \pi_{oa} \geq 0 & \text{for all } o \in \mathcal{O}, a \in \mathcal{A}, \\ \sum_{a \in \mathcal{A}} \pi_{oa} = 1 & \text{for all } o \in \mathcal{O}. \end{cases}$$

◇

3. THE GEOMETRY OF REWARD OPTIMIZATION

In this section, we discuss the formulation of the reward optimization problem as a polynomially constrained linear program from [MM22a]. For deterministic observations we provide a new description of the feasible state-action frequencies as the intersection of a product of varieties of rank-one matrices, an affine space, and the simplex (see Theorem 6).

Clearly, optimizing $R(\pi) = \langle r, \Phi(\pi) \rangle$ over $\Delta_{\mathcal{A}}^{\mathcal{O}}$ is equivalent to the *reward maximization problem in the state-action space (ROPSA)*:

$$(ROPSA) \quad \text{maximize } \langle r, \eta \rangle \quad \text{subject to } \eta \in \Phi(\Delta_{\mathcal{A}}^{\mathcal{O}}).$$

By definition, the feasible set $\Phi(\Delta_{\mathcal{A}}^{\mathcal{O}})$ is a subset of the probability simplex $\Delta_{S \times \mathcal{A}}$. Cramer's rule implies that the parametrization Φ is a rational map and hence, by the Tarski-Seidenberg theorem, the range $\Phi(\Delta_{\mathcal{A}}^{\mathcal{O}})$ is semialgebraic. Next we discuss the solution of the implicitization problem for the parametric set $\Phi(\Delta_{\mathcal{A}}^{\mathcal{O}})$ as recently given by [MM22a], i.e., a representation of this set as the solution set to a list of polynomial (in)equalities.

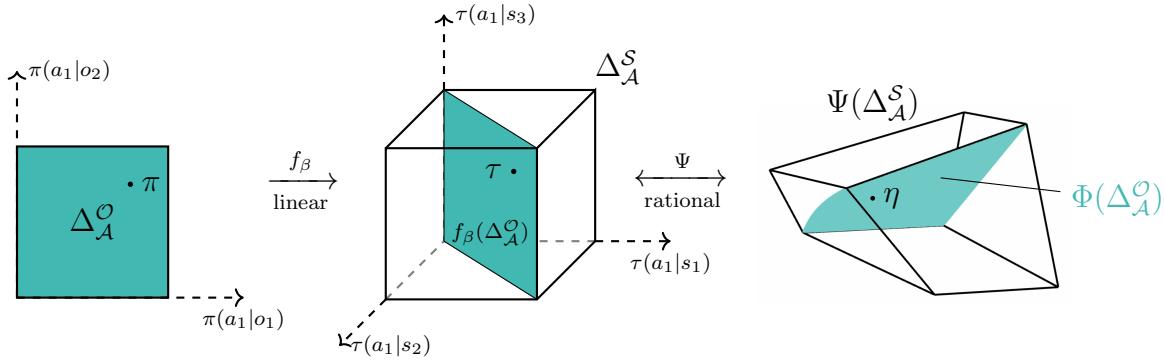


FIGURE 2. Shown is the policy polytope Δ_A^O (left), the effective policy polytope $f_\beta(\Delta_A^O)$ within the state policy polytope Δ_A^S (middle), and the set of feasible state-action frequencies $\Phi(\Delta_A^O)$ within the state-action polytope $\Psi(\Delta_A^S)$ (right). Note that $\Phi(\Delta_A^O)$ is the nonlinear solution set of the constraints given in equation (10) and $\Psi(\Delta_A^S)$ is a three-dimensional polytope that is combinatorially equivalent to the cube.

The mapping Φ can be seen as a composition $\Psi \circ f_\beta$ of a linear and non-linear map, illustrated in Figure 2 for the POMDP of Example 2, with

$$\begin{aligned} f_\beta: \Delta_A^O &\longrightarrow \Delta_A^S \\ \pi &\longmapsto \tau = \pi \circ \beta \end{aligned} \quad \text{and} \quad \begin{aligned} \Psi: \Delta_A^S &\longrightarrow \Delta_{S \times A} \\ \tau &\longmapsto \eta = (1 - \gamma)(I - \gamma P_\tau)^{-1}(\mu * \tau). \end{aligned}$$

We recall the following classic result.

Proposition 3 (The state-action polytope of Markov decision processes, [Der70]). *The image $\Psi(\Delta_A^S)$ is a polytope given by $\Psi(\Delta_A^S) = \mathcal{L} \cap \Delta_{S \times A}$, where*

$$(7) \quad \mathcal{L} := \left\{ \eta \in \mathbb{R}^{S \times A} : \ell_s(\eta) = 0 \text{ for all } s \in \mathcal{S} \right\},$$

$$\text{and } \ell_s(\eta) := \sum_a \eta_{sa} - \gamma \sum_{s',a'} \eta_{s'a'} \alpha(s|s', a') - (1 - \gamma)\mu_s.$$

In particular, the set of state-action frequencies $\Psi(\Delta_A^S)$ of a fully observable Markov decision process forms a polytope, referred to as the *state-action polytope*. The constraints encoded in \mathcal{L} describe a generalized stationarity property of the state-action frequencies, recovering stationarity in the limit where the discount factor is $\gamma = 1$. In order to relate the space of state policies to state-action frequencies we make the following assumption.

Assumption 4 (Positivity). For every $s \in \mathcal{S}$ and $\pi \in \Delta_A^O$, we assume that $\sum_a \eta_{sa} > 0$.

This assumption is satisfied, for example, if the system is ergodic or if the initial distribution $\mu \in \Delta_{\mathcal{S}}$ has full support, i.e., has only strictly positive entries. This can be seen by interpreting $\sum_a \eta_{sa}$ as a weighted average of the time spent in state s when following the policy π . An important consequence of this assumption is that the state policies τ and the state-action frequencies η are in one-to-one correspondence, whereby the state policies can easily be computed from the state-action frequencies by conditioning.

Proposition 5 ([MM22a]). *Under Assumption 4, the mapping $\Psi: \Delta_{\mathcal{A}}^{\mathcal{S}} \rightarrow \Psi(\Delta_{\mathcal{A}}^{\mathcal{S}})$ is rational and bijective with rational inverse given by conditioning*

$$\begin{aligned} \Gamma: \Psi(\Delta_{\mathcal{A}}^{\mathcal{S}}) &\longrightarrow \Delta_{\mathcal{A}}^{\mathcal{S}} \\ \eta &\longmapsto \tau, \quad \text{where } \tau_{as} = \frac{\eta_{sa}}{\sum_{a'} \eta_{sa'}}. \end{aligned}$$

The function Ψ is defined everywhere on $\Delta_{\mathcal{A}}^{\mathcal{S}}$ and bijectively identifies the defining inequalities of the polytope $f_{\beta}(\Delta_{\mathcal{A}}^{\mathcal{O}})$ within $\Delta_{\mathcal{A}}^{\mathcal{S}}$ with the defining inequalities of $\Phi(\Delta_{\mathcal{A}}^{\mathcal{O}})$ within $\Psi(\Delta_{\mathcal{A}}^{\mathcal{S}})$ via the pullback along Γ . This relates the geometry of $\Phi(\Delta_{\mathcal{A}}^{\mathcal{O}})$ and $f_{\beta}(\Delta_{\mathcal{A}}^{\mathcal{O}})$. The defining inequalities of $f_{\beta}(\Delta_{\mathcal{A}}^{\mathcal{O}})$ can be computed algorithmically, see e.g., [JKM04]. As we demonstrate in what follows, for deterministic observations the defining inequalities can be given in closed form. In particular, we show the following:

Theorem 6 (Feasible state-action frequencies). *For deterministic observation β the set of feasible state-action frequencies $\Phi(\Delta_{\mathcal{A}}^{\mathcal{O}})$ is the intersection $\mathcal{L} \cap \mathcal{X} \cap \Delta_{\mathcal{S} \times \mathcal{A}}$ of the linear space \mathcal{L} defined in (7), the product of real determinantal varieties*

$$\mathcal{X} := \left\{ \eta \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}} : \eta_{sa} \eta_{s'a'} - \eta_{sa'} \eta_{s'a} = 0 \ \forall a, a' \in \mathcal{A} \text{ and } s, s' \in \mathcal{S} \text{ with } g_{\beta}(s) = g_{\beta}(s') \right\},$$

and the probability simplex $\Delta_{\mathcal{S} \times \mathcal{A}}$. In particular, the only inequalities are of the form $\eta \geq 0$.

We call $\mathcal{L} \cap \mathcal{X}$ the *state-aggregation variety*. Note that \mathcal{X} is determined by the condition that for every observation o the $d_o \times n_{\mathcal{A}}$ submatrix $(\eta_{sa})_{s \in S_o, a \in \mathcal{A}}$ of η , consisting of all entries η_{sa} with $\beta(s) = o$, has rank one. In particular, the projective variety associated to \mathcal{X} is a join of Segre varieties.

Proof of Theorem 6. Proposition 3 provides a description of the polytope $\Psi(\Delta_{\mathcal{A}}^{\mathcal{S}})$ as the intersection $\mathcal{L} \cap \Delta_{\mathcal{S} \times \mathcal{A}}$ so we are left with finding defining equations for $\Psi(f_{\beta}(\Delta_{\mathcal{A}}^{\mathcal{O}}))$ in $\Psi(\Delta_{\mathcal{A}}^{\mathcal{S}})$. To do this, observe that the polytope $f_{\beta}(\Delta_{\mathcal{A}}^{\mathcal{O}})$ consists of those elements $\tau \in \Delta_{\mathcal{A}}^{\mathcal{S}}$ satisfying the linear equations $\tau_{as} - \tau_{as'}$ for all $a \in \mathcal{A}$, $s, s' \in \mathcal{S}$ such that $g_{\beta}(s) = g_{\beta}(s')$. In other words, all columns of τ indexed by states with equal observations coincide. Fix an action $a_o \in \mathcal{A}$ and a state $s_o \in S_o$ for each observation $o \in \mathcal{O}$. Then the non-redundant defining equalities of $f_{\beta}(\Delta_{\mathcal{A}}^{\mathcal{O}})$ are given by

$$l_{sa}^o(\tau) := \tau_{as} - \tau_{as_o} = 0,$$

for all observations $o \in \mathcal{O}$, actions $a \in \mathcal{A} \setminus \{a_o\}$, and states in the fiber $s \in S_o \setminus \{s_o\}$. These equations determine the range of f_{β} as a function $\mathbb{R}^{\mathcal{A} \times \mathcal{O}} \rightarrow \mathbb{R}^{\mathcal{A} \times \mathcal{S}}$ (corresponding to the set \mathcal{U} in [MM22a, Theorem 12]). After applying the pullback Γ^* of the conditioning map Γ to the linear functions l_{sa}^o we get the rational equations

$$(8) \quad (\Gamma^* l_{sa}^o)(\eta) = l_{sa}^o(\Gamma(\eta)) = \eta_{sa} \left(\sum_{a' \in \mathcal{A}} \eta_{sa'} \right)^{-1} - \eta_{s_o a} \left(\sum_{a' \in \mathcal{A}} \eta_{s_o a'} \right)^{-1} = 0,$$

which we rephrase as the vanishing of the polynomials

$$(9) \quad p_{sa}^o(\eta) := \eta_{sa} \sum_{a' \in \mathcal{A}} \eta_{s_o a'} - \eta_{s_o a} \sum_{a' \in \mathcal{A}} \eta_{sa'} = \sum_{a' \in \mathcal{A} \setminus \{a\}} (\eta_{sa} \eta_{s_o a'} - \eta_{s_o a} \eta_{sa'}).$$

These are defining polynomial equations of $\Psi(f_{\beta}(\Delta_{\mathcal{A}}^{\mathcal{O}}))$ in $\Psi(\Delta_{\mathcal{A}}^{\mathcal{S}})$. Let now \mathcal{W} be the variety determined by the equations (8). It remains to show $\mathcal{X} \cap \mathcal{L} \cap \Delta_{\mathcal{S} \times \mathcal{A}} = \mathcal{W} \cap \mathcal{L} \cap \Delta_{\mathcal{S} \times \mathcal{A}}$. Since

p_{sa}^o is a linear combination of 2×2 minors, we have the inclusion $\mathcal{X} \subseteq \mathcal{W}$. On the other hand, equation (8) implies the linear dependence of the two vectors

$$(\eta_{sa})_a, (\eta_{soa})_a \in \mathbb{R}^{\mathcal{A}}$$

for every observation o and state $s \in S_o$. Consequently, every 2×2 minor in the definition of \mathcal{X} vanishes on $\mathcal{W} \cap \mathcal{L} \cap \Delta_{\mathcal{S} \times \mathcal{A}}$. This shows the desired inclusion

$$\mathcal{W} \cap \mathcal{L} \cap \Delta_{\mathcal{S} \times \mathcal{A}} \subseteq \mathcal{X} \cap \mathcal{L} \cap \Delta_{\mathcal{S} \times \mathcal{A}},$$

which finishes the proof. \square

Hence by Theorem 6, in the case of a deterministic observation kernel all inequalities are linear and the equalities are either linear or 2×2 principal minors. This is in contrast to the case of general observation kernels, where nonlinear defining inequalities appear and the polynomial constraints might be of higher degree (see [MM22a, Theorem 16]). Since all defining equalities of \mathcal{X} are binomial, it is a toric variety. The following monomial parametrization of \mathcal{X} can be inferred from the discussion of the family of state-frequencies and equation (2):

$$\begin{aligned} \mathbb{R}^{\mathcal{S}} \times \mathbb{R}^{\mathcal{A} \times \mathcal{O}} &\longrightarrow \mathcal{X} \\ (\rho, \pi) &\longmapsto \eta, \quad \text{where } \eta(s, a) = \pi(a|g_\beta(s))\rho(s). \end{aligned}$$

The following characterization of the set of feasible state-action frequencies with fewer equations will be useful later.

Corollary 7 (Alternative characterization of feasible state-action frequencies). *For deterministic observation β , fix an arbitrary action $a_o \in \mathcal{A}$ and an arbitrary state $s_o \in S_o$ for every $o \in \mathcal{O}$. The set of feasible state-action frequencies $\Phi(\Delta_{\mathcal{A}}^{\mathcal{O}})$ can be described as the intersection $\mathcal{L} \cap \mathcal{Y} \cap \Delta_{\mathcal{S} \times \mathcal{A}}$, where*

$$\mathcal{Y} := \{\eta \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}} : p_{sa}^o(\eta) = 0 \text{ for all } o \in \mathcal{O}, a \in \mathcal{A} \setminus \{a_o\}, s \in S_o \setminus \{s_o\}\},$$

and the polynomials p_{sa}^o are given in (9). \mathcal{Y} is a complete intersection of these polynomials.

Proof. This follows directly from the proof of Theorem 6. \square

Example 8. We continue Example 2 from above. The defining (in)equalities of $\Phi(\Delta_{\mathcal{A}}^{\mathcal{O}})$, described in (7) and Theorem 6, take the form

$$\begin{aligned} \mathcal{L} = \{\eta \in \mathbb{R}^{3 \times 4} : 3\eta_{s_1 a_1} + 6\eta_{s_1 a_2} - 3\eta_{s_2 a_2} - 1 = 0, 6\eta_{s_2 a_1} + 6\eta_{s_2 a_2} - 3\eta_{s_3 a_2} - 1 = 0, \\ 3\eta_{s_3 a_1} + 6\eta_{s_3 a_2} - 3\eta_{s_1 a_2} - 3\eta_{s_2 a_1} - 1 = 0\} \end{aligned}$$

and

$$\mathcal{X} = \{\eta \in \mathbb{R}^{3 \times 4} : \eta_{s_1 a_1} \eta_{s_2 a_2} - \eta_{s_1 a_2} \eta_{s_2 a_1} = 0\}.$$

Thus, the reward optimization problem (ROPSA) is

$$(10) \quad \text{maximize } \eta_{s_1 a_1} + \eta_{s_1 a_2} \quad \text{subject to} \quad \begin{cases} 3\eta_{s_1 a_1} + 6\eta_{s_1 a_2} - 3\eta_{s_2 a_2} - 1 = 0 \\ 6\eta_{s_2 a_1} + 6\eta_{s_2 a_2} - 3\eta_{s_3 a_2} - 1 = 0 \\ 3\eta_{s_3 a_1} + 6\eta_{s_3 a_2} - 3\eta_{s_1 a_2} - 3\eta_{s_2 a_1} - 1 = 0 \\ \eta_{s_1 a_1} \eta_{s_2 a_2} - \eta_{s_1 a_2} \eta_{s_2 a_1} = 0 \\ \eta_{s_1 a_1}, \eta_{s_1 a_2}, \eta_{s_2 a_1}, \eta_{s_2 a_2}, \eta_{s_3 a_1}, \eta_{s_3 a_2} \geq 0. \end{cases}$$

The feasible set of this optimization problem is shown on the right in Figure 2. Comparing this to the optimization problem over the policy polytope (6) with objective function (4), now the constraints are more complex and nonlinear but the objective is linear. \square

4. COMBINATORIAL AND ALGEBRAIC COMPLEXITY OF THE PROBLEM

In this section we study the number of critical points of the reward optimization problem in the case of deterministic observations. We apply methods from polynomial optimization and in particular the theory of algebraic degrees developed in [NR09] to obtain upper bounds on the number of complex critical points. A similar approach was pursued in [MM22a] for the case of invertible observation matrix β , in which case there are linear equations and polynomial inequalities.

The description of $\Phi(\Delta_{\mathcal{A}}^{\mathcal{O}})$ obtained in Corollary 7 allows to reformulate the reward optimization problem (ROPSA) as the following constrained polynomial optimization problem:

$$(POP) \quad \text{maximize } \langle r, \eta \rangle \quad \text{subject to} \quad \begin{cases} \ell_s(\eta) = 0 & \text{for } s \in \mathcal{S}, \\ p_{sa}^o(\eta) = 0 & \text{for } o \in \mathcal{O}, a \in \mathcal{A} \setminus \{a_o\}, s \in S_o \setminus \{s_o\}, \\ \eta_{sa} \geq 0 & \text{for } s \in \mathcal{S}, a \in \mathcal{A}, \end{cases}$$

where the linear constraints ℓ_s are given in Proposition 3, the polynomial constraints $p_{sa}^o(\eta)$ are provided in (9) taking a fixed action $a_o \in \mathcal{A}$ and a fixed state $s_o \in S_o$ for each observation $o \in \mathcal{O}$, and the inequality constraints simply ensure the entries of η being nonnegative. Observe that problem (POP) is in fact a quadratically constrained linear program.

We bound the number of critical points individually for each boundary component of the feasible set. A boundary component consists of all feasible points for which a given subset of the inequality constraints are active. The boundary components of the feasible set $\Phi(\Delta_{\mathcal{A}}^{\mathcal{O}})$ are in one-to-one correspondence with the faces of $\Delta_{\mathcal{A}}^{\mathcal{O}}$ according to

$$(11) \quad \{\pi \in \Delta_{\mathcal{A}}^{\mathcal{O}} : \pi(a|o) = 0 \forall a \in A_o, o \in \mathcal{O}\} \longleftrightarrow \{\eta \in \Phi(\Delta_{\mathcal{A}}^{\mathcal{O}}) : \eta_{sa} = 0 \text{ for } a \in A_{g_{\beta}(s)}\},$$

where A_o is a proper subset of \mathcal{A} for every $o \in \mathcal{O}$, and $g_{\beta}(s)$ is the observation associated with state s . In particular, there is a boundary component associated to each tuple $(A_o)_{o \in \mathcal{O}}$ with $A_o \subsetneq \mathcal{A}$, $o \in \mathcal{O}$.

We point out the following result, which allows us to ignore high-dimensional boundary components when searching for a maximizer of the reward. Recall that for an observation $o \in \mathcal{O}$, the cardinalities of the fibers of g_{β} are denoted by $d_o = |S_o|$.

Theorem 9 (Existence of maximizers in low dimensional faces, [MR17]). *There exist $A_o \subsetneq \mathcal{A}$ with $|A_o^c| \leq d_o$, $o \in \mathcal{O}$, such that the set B described in (12) contains a (globally optimal) solution of the problem (POP).*

Remark 10. One approach to solving (POP) is to solve the critical equations over every boundary component and then selecting the critical point with the highest objective value. According to Theorem 9 there is a lower-dimensional boundary component that contains a global maximizer. This implies that, instead of considering the critical points in all $(2^{n_{\mathcal{A}}} - 1)^{n_{\mathcal{O}}}$ boundary components, it is enough to consider those in the boundary components with $A_o \subsetneq \mathcal{A}$ satisfying $|A_o| \geq n_{\mathcal{A}} - d_o$. This reduces the number of boundary components that need to be checked to

$$\prod_{o \in \mathcal{O}} \left(\sum_{l_o=\max(n_{\mathcal{A}}-d_o, 0)}^{n_{\mathcal{A}}-1} \binom{n_{\mathcal{A}}}{l_o} \right),$$

which we call *relevant* boundary components. Note that this number only depends on the number of actions $n_{\mathcal{A}}$ and d_o (the cardinality of the fibers of g_{β}).

4.1. Bounds via algebraic degrees of polynomial optimization. With the description of the boundary components of the feasible set at hand, we can deduce upper bounds on the number of critical points over each of them based on the degrees of the defining equations and the degree of the objective function.

Theorem 11 (Bound on the algebraic degree). *Consider a POMDP with deterministic observations. Fix $A_o \subsetneq \mathcal{A}$ for every $o \in \mathcal{O}$ and set $n := n_{\mathcal{S}}n_{\mathcal{A}} - n_{\mathcal{S}} - \sum_o d_o|A_o|$ and $m := \sum_o (d_o - 1)(|A_o^c| - 1)$, where we assume n is not zero. Then the number of critical points of the linear function $\eta \mapsto \langle r, \eta \rangle$ over*

$$(12) \quad B := \{\eta \in \mathcal{L} \cap \mathcal{X} : \eta_{sa} = 0 \text{ for } a \in A_{g_{\beta}(s)}\}$$

is upper bounded by $2^m \binom{n-1}{m-1}$.

Proof. Recall from Corollary 7 that $\Phi(\Delta_{\mathcal{A}}^{\mathcal{O}})$ is defined in $\mathbb{R}^{\mathcal{S} \times \mathcal{A}}$ as an intersection of $n_{\mathcal{S}}$ linear equations, $\sum_o (d_o - 1)(n_{\mathcal{A}} - 1)$ quadratic equations of the form (9), and the linear inequalities $\eta \geq 0$. It is not difficult to see that for any choice of $A_o \subsetneq \mathcal{A}$, $o \in \mathcal{O}$, the linear equations $\ell_s(\eta) = 0$, $s \in \mathcal{S}$ and $\eta_{sa} = 0$, $a \in A_o$, $s \in S_o$, $o \in \mathcal{O}$ are linearly independent. On the set B given in (12) there are $\sum_o d_o|A_o|$ active linear inequalities with $A_o \subsetneq \mathcal{A}$ for each $o \in \mathcal{O}$, and hence B is contained in an affine space of dimension $n = n_{\mathcal{S}}n_{\mathcal{A}} - n_{\mathcal{S}} - \sum_o d_o|A_o|$. Further, given these linear equations, the quadratic equations

$$p_{sa}^o(\eta) = \eta_{sa} \sum_{a' \in \mathcal{A}} \eta_{s_o a'} - \eta_{s_o a} \sum_{a' \in \mathcal{A}} \eta_{s a'} = 0$$

are redundant for all $a \in A_o$, $s \in S_o$. By choosing $a_o \in A_o^c$ in Corollary 7 for every $o \in \mathcal{O}$ there remain $n_{\mathcal{A}} - |A_o| - 1$ non-redundant quadratic equalities for every $s \in S_o \setminus \{s_o\}$. Therefore, we get $m = \sum_o (d_o - 1)(|A_o^c| - 1)$ non-redundant quadratic equalities. By Theorem 2.2 and Corollary 2.5 in [NR09] the algebraic degree for the optimization of the linear function $r \in \mathbb{R}^{\mathcal{S} \times \mathcal{A}}$ over an n -dimensional affine space subject to m non-redundant quadratic constraints is upper bounded by $2^m \binom{n-1}{m-1}$. \square

With Theorem 11 we can provide upper bounds for the number of critical points of the optimization problem (POP). Indeed, the number of critical points over the interior

$$(13) \quad \{\eta \in \mathcal{L} \cap \mathcal{X} : \eta_{sa} = 0 \text{ for all } a \in A_{g_{\beta}(s)}, \eta_{sa} > 0 \text{ otherwise }\}$$

of a boundary component is clearly upper bounded by the number of critical points over B defined in (12). This bound over the individual boundary components can be summed to obtain an upper bound on the number of critical points of the polynomial optimization problem (POP) (see also [NR09]). Note that the Zariski closure of the interior of a boundary component defined in (13) is contained in B but might be a strict subset. Similarly, a bound on the number of critical points over the relevant boundary components can be established.

Remark 12 (Tighter bounds via polar degrees). Since the problem (POP) has a linear objective, under weak assumptions the number of critical points over every boundary component is upper bounded by the polar degree of the associated variety. This approach may yield tighter bounds as demonstrated in the special case of a blind controller with two actions, i.e., a system with one observation and two actions in [MM22a]. The authors obtain an upper bound linear in $n_{\mathcal{S}}$ compared to the exponential upper bound of $n_{\mathcal{S}} \cdot 2^{n_{\mathcal{S}}-1} + 2$ that follows from Theorem 11. A refinement of Theorem 9 for the case of mean rewards was also presented in [MRA19, Theorem 2], which would be worthwhile studying from an algebraic standpoint.

4.2. Evaluation of the bounds. In Table 1 we present the upper bounds on the number of critical points for problems of different size. We compare the bound on the total number of critical points obtained by iterating Theorem 11 over all boundary components and the one iterating only over the relevant components described in Theorem 9. In addition, we report the total and relevant number of boundary components discussed in Remark 10. Both, the number of boundary components and the upper bound on the number of critical points, depend on n_S, n_A , and the tuple $(d_o)_{o \in \mathcal{O}}$. The two extreme cases for the tuple $(d_o)_{o \in \mathcal{O}}$, namely (n_S) and $(1, \dots, 1)$, correspond to a *blind controller*, i.e., all states map to the same observation, and the *fully observable case*, i.e., states and observations are in one-to-one correspondence, respectively. The bounds are independent of the specific α , so long as Assumption 4 is satisfied.

n_S	n_A	partitions of n_S : $(d_o)_{o \in \mathcal{O}}$	Number of boundary components		Bound on number of critical points	
			total	relevant	total	relevant
3	2	(3)	3	3	10	10
		(2, 1)	9	6	10	8
		(1, 1, 1)	27	8	8	8
4	3	(4)	7	7	1419	1419
		(3, 1)	49	21	2237	561
		(2, 2)	49	36	1265	153
		(2, 1, 1)	343	54	1189	81
		(1, 1, 1, 1)	2401	81	81	81
5	3	(5)	7	7	9411	9411
		(4, 1)	49	21	23745	4257
		(3, 2)	49	42	13431	4371
		(3, 1, 1)	343	63	24363	1683
		(2, 2, 1)	343	108	12159	459
		(2, 1, 1, 1)	2401	162	9195	243
		(1, 1, 1, 1, 1)	16807	243	243	243

TABLE 1. Listed are the number of boundary components and the upper bound on the number of critical points from Theorem 11 both over all boundary components and over the subset of relevant boundary components from Theorem 9 for problems of different size.

In these examples we observe that restricting to the relevant boundary components significantly reduces the upper bound. This is reflected in the last two columns in the table. The difference is most notable when the fibers of g_β have a small cardinality, i.e., only few states lead to the same observation. In the fully observable case, the relevant boundary components correspond to the vertices of $\Delta_A^\mathcal{O}$. This is consistent with the fact that in the fully observable case the feasible set $\Phi(\Delta_A^\mathcal{O})$ is a polytope [Der70] and hence the optimization problem (POP) is a linear program, for which the solutions are attained at the vertices. On the other hand, in the case of a blind controller (with a single observation o), all boundary components are relevant since $d_o = n_S$.

5. NUMERICAL METHODS FOR THE OPTIMIZATION OF DECISION RULES

In POMDPs reward optimization over the set of memoryless stochastic policies (ROP) is known to be hard in theory (NP-hard [VLB12]) and also difficult in practice as the reward function R is nonconvex and has sub-optimal strict local optima [PLT11, BR19]. In this section we discuss how the geometric description of reward optimization facilitates computational approaches based on numerical algebra. We derive polynomial systems for the critical points, globally from the Karush-Kuhn-Tucker (KKT) conditions, and separately for each boundary component of $\Phi(\Delta_{\mathcal{A}}^{\mathcal{O}})$ from the Lagrangian criterion. For different choices of $n_{\mathcal{S}}$, $n_{\mathcal{A}}$, and generic data (i.e., generic α, μ and r), we compute the complex and real solutions of the KKT and Lagrangian systems, and compare the number of solutions with the theoretical upper bounds established in Section 4. Finally, we compare these approaches with other popular methods from constrained optimization: the interior point solver `Ipopt` and convex relaxations via the moment-SOS-approach.

5.1. Critical equations and computation.

The KKT critical point equations. A standard approach for constrained optimization problems are the KKT conditions [KT51], which provide necessary conditions of stationary points under certain regularity conditions; see e.g. [Aba67, Ber97, BSS06]. If both the constraints and objective are polynomial, the KKT conditions form a polynomial system, which can be solved using various numerical algebraic methods.

Applied to our problem, the KKT conditions reduce to the following polynomial system in $\eta \in \mathbb{R}_{\geq 0}^{\mathcal{S} \times \mathcal{A}}$ with multipliers $\lambda \in \mathbb{R}^{\mathcal{S}}$, $\nu_{sa}^o \in \mathbb{R}$, $\kappa \in \mathbb{R}_{\geq 0}^{\mathcal{S} \times \mathcal{A}}$:

$$\begin{aligned}
 \text{Primal feasibility:} \quad & \ell_s(\eta) = 0 \text{ for } s \in \mathcal{S}, \\
 & p_{sa}^o(\eta) = 0 \text{ for } o \in \mathcal{O}, a \in \mathcal{A} \setminus \{a_o\}, s \in S_o \setminus \{s_o\}, \\
 (14) \quad \text{Complementary slackness:} \quad & \kappa_{soa} \eta_{soa} = 0 \text{ for all } s_o, a, \\
 \text{Stationarity:} \quad & r + \sum_s \lambda_s \nabla \ell_s(\eta) + \sum_{o,s,a} \nu_{sa}^o \nabla p_{sa}^o(\eta) + \kappa = 0,
 \end{aligned}$$

where $a_o \in \mathcal{A}$ and $s_o \in S_o$ for every $o \in \mathcal{O}$ are fixed arbitrarily. Here we have included the primal feasibility $\eta_{sa} \geq 0$ for $s \in \mathcal{S}, a \in \mathcal{A}$ and the dual feasibility $\kappa_{sa} \geq 0$ for $s \in \mathcal{S}, a \in \mathcal{A}$ in the definition of the search space for η and κ .

The number of linear constraints ℓ_s is $n_{\mathcal{S}}$, while the number of polynomial constraints p_{sa}^o is $(n_{\mathcal{A}} - 1) \sum_{o \in \mathcal{O}} (d_o - 1) = (n_{\mathcal{A}} - 1)(n_{\mathcal{S}} - n_{\mathcal{O}})$. Due to the symmetry of the effective policies, there are only $n_{\mathcal{O}} n_{\mathcal{A}}$ inequalities $\eta_{soa} \geq 0$ for each $a \in \mathcal{A}, o \in \mathcal{O}$. Hence the dimension of the square KKT system (14) is

$$n_{\mathcal{S}} n_{\mathcal{A}} + n_{\mathcal{S}} + (n_{\mathcal{A}} - 1)(n_{\mathcal{S}} - n_{\mathcal{O}}) + n_{\mathcal{O}} n_{\mathcal{A}} = 2n_{\mathcal{S}} n_{\mathcal{A}} + n_{\mathcal{O}}.$$

In this setting, we can verify that the linear independence constraint qualification is satisfied. Given an element η^* in the feasible set $\Phi(\Delta_{\mathcal{A}}^{\mathcal{O}})$, it suffices to verify the linear independence of the gradients of the active inequality constraint functions and the equality constraints at η^* . Notice that under the pullback along the birational morphism Ψ^{-1} the constraints are identified with affine-linear functions. Checking the linear independence of their gradients can be done by counting the dimension of the faces.

The Lagrange critical point equations over boundary components. Alternatively to solving the KKT system, one can compute the critical equations given by the Lagrange criterion over every boundary component individually. If there are no inequality constraints, the KKT equations specialize to the Lagrange multiplier equations. Consider a boundary component B in (12) for a choice of $A_o \subsetneq \mathcal{A}$ for every $o \in \mathcal{O}$, and consider the optimization problem over B . This amounts setting $\eta(s, a) = 0$ for $a \in A_o$ whenever $g_\beta(s) = o, o \in \mathcal{O}$, which reduces optimization to a subspace of $\mathbb{R}^{\mathcal{S} \times \mathcal{A}}$. We denote the new primal variables by $\hat{\eta}$. Similarly, we denote the restriction of ℓ_s and p_{sa}^o to this space by $\hat{\ell}_s$ and \hat{p}_{sa}^o and the projection of r onto this space (i.e., the vector obtained by dropping the indices which are set to zero in η) by \hat{r} . In the lower dimensional variables $\hat{\eta}$ for a given B the Lagrange system becomes

$$(15) \quad \begin{aligned} \text{Feasibility:} \quad & \hat{\ell}_s(\hat{\eta}) = 0 \text{ for } s \in \mathcal{S}, \\ & \hat{p}_{sa}^o(\hat{\eta}) = 0 \text{ for } o \in \mathcal{O}, a \in \mathcal{A} \setminus \{a_o\}, s \in S_o \setminus \{s_o\}, \\ \text{Stationarity:} \quad & \hat{r} + \sum_s \lambda_s \nabla \hat{\ell}_s(\hat{\eta}) + \sum_{o,s,a} \nu_{sa}^o \nabla \hat{p}_{sa}^o(\hat{\eta}) = 0, \end{aligned}$$

where $a_o \in A_o^c$ and $s_o \in S_o$ are fixed arbitrarily for every $o \in \mathcal{O}$. The dimension of the primal variable $\hat{\eta}$ is $n_{\mathcal{S}} n_{\mathcal{A}} - \sum_o d_o |A_o|$, the dimension of the Lagrange multipliers λ is $n_{\mathcal{S}}$ and of ν is $\sum_o (d_o - 1) (|A_o^c| - 1)$ (see also proof of Theorem 11). Overall, the Lagrange system (15) is a square polynomial system of dimension

$$2n_{\mathcal{S}} n_{\mathcal{A}} - (n_{\mathcal{A}} - 1)n_{\mathcal{O}} - \sum_o (2d_o - 1)|A_o|.$$

Remark 13 (Lagrange vs KKT system). It is easy to see that every real solution of the KKT system satisfying the primal and dual inequality constraints $\eta \geq 0, \kappa \geq 0$ is a solution of the Lagrange system over a boundary component, namely the boundary component defined by the zeros of η ; see Figure 3 for an illustrated example of this situation. When solving the KKT system (14), usually one solves the system of equations without the nonnegativity conditions $\eta \geq 0$ and $\kappa \geq 0$ and then selects the nonnegative solutions. Note that every solution of the Lagrange system over a boundary component appears as the solution of the KKT system without the nonnegativity constraints. Hence, solving the KKT system gives at least as many solutions as solving the Lagrange system over every boundary component.

Computation. The optimization problem (POP) can be solved using several methods:

- First, we use the numerical algebra package `HomotopyContinuation.jl` [BT18] to solve the KKT system (14) and the Lagrange system (15) of each boundary component. This automatically certifies the results [BRT21], meaning that for every returned solution it is guaranteed that there exists a unique true solution in a small neighborhood. From the returned solutions to the critical equations, we then just need to select the real ones that satisfy the primal inequality constraints $\eta_{s,a} \geq 0$, and among them the one that has the maximum objective value.
- Alternatively, we solve a convex relaxation of the polynomial optimization problem. Namely we relax the problem to a semidefinite program (SDP) via the moment-SOS-approach that is implemented in the freeware `GloptiPoly3` [HLL09], and solve the SDPs using the numerical solver `Mosek`; see [DA21] for details. We note that `GloptiPoly3` builds upon a hierarchy of moment/SOS programs (also called Lasserre hierarchy), which allows to approximate the optimal value arbitrarily close, and can be used to test optimality and extract global optimizers. We use this key feature to check if our methods reach global optimality.

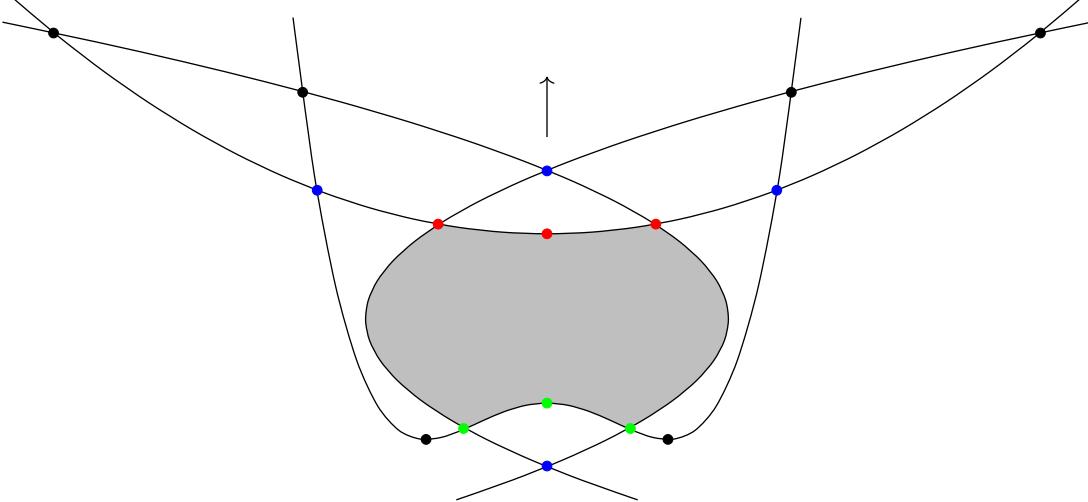


FIGURE 3. Schematic illustration of the feasible region (gray) and objective gradient (arrow) of a polynomially constrained linear program showing (i) the solutions of the KKT system checking only primal ($\eta \geq 0$) inequality constrains (red and green); and checking primal ($\eta \geq 0$) and dual ($\kappa \geq 0$) inequality constrains (red), (ii) the solutions of the KKT system without checking inequalities (red, green, black and blue), (iii) the positive ($\eta \geq 0$) solutions of the Lagrange systems over all boundary components (red and green) (iv) all solutions of the Lagrange systems over all boundary components (red, green and black).

- We may also solve the constrained optimization problem (POP) using the interior point solver Ipopt [WB06], which is a local optimization method for large-scale nonlinear optimization, an approach recently pursued in [MM22b].

5.2. Experiments.

Description of the experiments. We test our computational approach on random POMDPs of different sizes. To this end we first specify the number of states n_S , the number of actions n_A , and the number of states aggregated in each observation $(d_o)_{o \in \mathcal{O}}$ with $\sum_o d_o = n_S$. For each specification of these values, we generate 20 random problems as follows. We sample the initial state distribution μ and the transition probabilities $\alpha(\cdot|s, a)$, $(s, a) \in \mathcal{S} \times \mathcal{A}$ from a uniform distribution on the simplex Δ_S , and sample the instantaneous reward vector $r \in \mathbb{R}^{S \times A}$ from a standard Gaussian distribution. We use the same random data for each of the two approaches, KKT and Lagrange over boundary components.

Discussion of the results. In this section we discuss the experimental results on the number of solutions obtained by solving the KKT and Lagrange systems introduced above. In Table 2 we report the average and standard deviation of the number of complex, real, and positive solutions returned in each case. Note that in our setting, positive solutions (i.e., solutions satisfying $\eta \geq 0$) are (primal) feasible solutions. We also compare the performance and the computational times of these methods with convex relaxations and interior point methods.

Following the discussion in Remark 13, we start by comparing the number of solutions of the KKT and the Lagrange systems. In Table 2 we see that the KKT system has at least as many complex solutions as the Lagrange systems over all boundary components. This is consistent with our previous discussion, since, as we have pointed out, any solution of the Lagrange system over a boundary component is a solution of KKT. Moreover, we observe

n_S	n_A	$(d_o)_{o \in \mathcal{O}}$	KKT			Lagrange (all)			Lagrange (relevant)		
			complex	real	positive	complex	real	positive	complex	real	positive
3	2	(3)	6±0	4.4±1.2	2.1±0.3	6±0	4.4±1.2	2.1±0.3	6±0	4.4±1.2	2.1±0.3
		(2,1)	12±0	10.1±1.9	4.25±0.44	10±0	8.2±1.9	4.25±0.44	8±0	6.7±1.6	4.25±0.44
		(1,1,1)	20±0	20±0	8±0	8±0	8±0	8±0	8±0	8±0	8±0
4	3	(4)	45±0	17.1±4.3	4.3±1.3	45±0	17.1±4.3	4.3±1.3	45±0	17.1±4.3	4.3±1.3
		(3,1)	150±0	79±11	11±1.9	129±0	68.7±9.7	11±1.9	81±0	41.6±8.5	10.9±1.8
		(2,2)	281.6±0.75	154±16	13.9±4.7	263±0	136±16	13.9±4.7	153±0	89±10	13.65±4.3
		(2,1,1)	381.2±0.7	292±23	31.5±4.3	216±0	168±16	31.5±4.3	81±0	68±11	30.9±4.0
		(1,1,1,1)	495±0	495±0	81±0	81±0	81±0	81±0	81±0	81±0	81±0
5	3	(5)	71±0	21.4±6	3.7±0.98	71±0	21.4±6	3.7±0.98	71±0	21.4±6	3.7±0.98
		(3,2)	637.95±0.76	219±28	12.60±2.9	626±0	213±29	12.6±2.9	477±0	171±24	12.6±2.9
		(4,1)	269.85±0.49	99±20	11.9±3.3	234±0	87±18	11.9±3.3	144±0	52±13	11.55±2.6
		(3,1,1)	881.95±0.22	436±68	36±10	558±0	285±47	36±10	243±0	117±20	35.3±9.2
		(2,2,1)	1717.3±2.5	890±49	35.6±5.3	1260±0	624±56	36.5±7.1	459±0	244±25	35.7±6.6
		(2,1,1,1)	2269.9±3.9	1712±142	89±12	810±0	624±74	89.3±12.3	243±0	195±37	88.1±9.5
		(1,1,1,1,1)	3002.9±0.31	3002.9±0.3	243±0	243±0	243±0	243±0	243±0	243±0	243±0

TABLE 2. Mean and standard deviation of the number of solutions of the KKT system (14), the Lagrange system (15) over all boundary components, and the Lagrange system over the relevant boundary components, for 20 random POMDPs with the indicated number of states n_S , actions n_A , and state aggregation partition $(d_o)_{o \in \mathcal{O}}$. In our setting, positive solutions are feasible solutions.

that KKT and Lagrange over all boundary components have in general the same number of positive solutions (see Remark 13 and Figure 3).

It is also worth noting the difference between the number of complex, real, and positive solutions. That is, in Table 2 we observe in general that for the three types of systems there is a drop between the number of complex solutions and the number of real and positive solutions. However, we find an exception to this in the Lagrange system for fully observable systems ($d_o = (1, \dots, 1)$), where the number of complex, real, and positive solutions coincide. The reason for this is that in this case all boundary components are affine spaces, so only the zero-dimensional boundary components have a solution, and these correspond precisely to the $n_A^{n_S}$ vertices of the feasible set.

We also observe that the number of complex solutions has a much smaller variance than the number of real or positive ones. This is expected, since choosing the coefficients of polynomial systems randomly gives the same number of complex solutions with probability one. In fact, the number of complex solutions for the Lagrange system has no variance across the different random parameters. Still, we see a small variance in the number of complex KKT solutions, which we attribute to numerical instability which can prevent the software package `HomotopyContinuation.jl` from finding all solutions to the KKT system. In contrast to the complex case, the variance on the number of real and positive solutions is not due to numerical errors. This is a typical phenomenon in polynomial systems, and one of the possible limitations of classic algebraic methods when one wants to estimate the number of real solutions of a system.

In the following we compare the experimental results presented in Table 2 with the theoretical upper bounds shown in Table 1 and highlight two particular facts. First notice that

in most cases the theoretical bound is significantly larger than the number of solutions of the Lagrange system. Moreover, this gap becomes particularly pronounced for problems where the fibers of g_β are large. This clearly indicates that there is a discrepancy between the theoretical bounds and the algebraic degree of the optimization problem. Indeed, our bounds are based on theory for generic polynomials and hence we do not expect that they provide a tight estimate of the algebraic degree for the particular polynomials we are dealing with. Here we also observe a particular behavior in the case of fully observable systems where the number of critical points of the Lagrange systems agree with our bounds. On the other hand, we see that in some cases the number of solutions of KKT is larger than the bound, which agrees with our discussion on solutions of KKT and Lagrange systems in Remark 13.

In addition to analyzing the number of solutions of the KKT and Lagrange systems, we are interested in comparing the different solution methods for the optimization problem. Therefore, we compare the optimal solution found by solving these systems with `HomotopyContinuation.jl` with the one found by `Ipopt` and `GloptiPoly3`. Although `HomotopyContinuation.jl` is not guaranteed to find all solutions to the KKT and Lagrange systems, we observe that this approach yields a reward that is at least as high as the one obtained by the interior point method `Ipopt` and in a few instances strictly higher. In fact, solving the optimization problem with `GloptiPoly3` returns a certificate for the optimality of the result, which in all computed instances coincides with the optimal value obtained by solving the KKT and Lagrange systems with `HomotopyContinuation.jl`. That is, `GloptiPoly3` offers numerical evidence that they always provide globally optimal solutions. It is noteworthy that in all computed instances using `GloptiPoly3`, the optimal value of the optimization problem is already attained at the first order relaxation of the Lasserre hierarchy [Las01]. We conjecture that objective value exactness for the first order relaxation of (POP) holds with high probability for generic input data. Since the size of the SDP depends very sensitively on the order of the relaxation, this conjecture would remedy one of the major drawbacks of the SDP relaxation method.

Finally, we observed that solving the Lagrange equations only over the relevant boundary components is up to two orders of magnitude faster than solving them over all boundary components. The improvements become more pronounced when g_β has small fibers in which we can exclude more faces by the means of Theorem 9; see also Table 1. The computation times for the solution of the KKT system are on the same order to magnitude as the computation time of the solution of the Lagrange systems over all boundary components. KKT was slightly faster when g_β has small fibers and slightly slower when g_β has large fibers.

Reproducibility statement. The computer code for our experiments is publicly available at <https://github.com/marinagarrote/Algebraic-Optimization-of-Sequential-Decision-Rules>. We conducted our experiments using Julia [BEKS17] version 1.7.0, an open source programming language under the MIT license. We used the Julia package `HomotopyContinuation.jl` version 2.6.3, which is freely available for personal use under the MIT license, and `Ipopt.jl` version 0.7.0, which is a Julia interface to the COIN-OR nonlinear solver `Ipopt` available under the EPL (Eclipse Public License) open-source license. The convex relaxation is computed via the freeware `GloptiPoly3` implemented in Matlab, for which there also exists an Octave implementation.

6. CONCLUSION AND OUTLOOK

Reward optimization in infinite-horizon discounted MDPs with state aggregation and stationary stochastic policies is equivalent to a polynomially constrained linear objective problem. We obtained a characterization of the feasible region of state-action frequencies for this problem as the intersection of a product of affine varieties of rank one matrices, an affine space and a simplex. Using this description as a polynomial optimization problem, we derived an upper bound on the number of critical points of the reward optimization problem.

To solve the polynomial optimization problem, we considered KKT equations and the Lagrange system over individual boundary components, where we leveraged knowledge about the location of maximizers on lower dimensional boundary components. The relatively small number of solutions observed in the experiments indicate that there is room for refining the theory either to obtain tighter estimates of the algebraic degree or also tighter descriptions of the possible number of feasible solutions. Using a convex relaxation to an SDP we obtained empirical evidence that our approach of solving the critical equations provides global maximizers of the reward. This is in strong contrast to naive gradient optimization, which yields only locally optimal solutions for this problem. In our experiments, the first order relaxation produced exact objective values, which indicates an interesting direction for further research.

ACKNOWLEDGMENT

The authors are grateful to Bernd Sturmfels for insightful discussions. GM and JM have been supported by ERC Starting Grant 757983 and DFG SPP 2298 Grant 464109215. GM has been supported by NSF CAREER Award DMS-2145630. JM acknowledges support from the International Max Planck Research School for Mathematics in the Sciences (IMPRS MiS) and the Evangelisches Studienwerk Villigst e.V.

REFERENCES

- [Aba67] J. Abadie, *On the Kuhn-Tucker theorem*, Nonlinear Programming (NATO Summer School, Menton, (1964), North-Holland, Amsterdam, 1967, pp. 19–36. MR 0218116
- [ABZ06] C. Amato, D.S. Bernstein, and S. Zilberstein, *Solving POMDPs using quadratically constrained linear programs*, Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems, 2006, pp. 341–343.
- [AYA18] K. Azizzadenesheli, Y. Yue, and A. Anandkumar, *Policy gradient in partially observable environments: Approximation and convergence*, 2018, Preprint, arXiv:1810.07900.
- [BEKS17] J. Bezanson, A. Edelman, S. Karpinski, and V.B. Shah, *Julia: A fresh approach to numerical computing*, SIAM review **59** (2017), no. 1, 65–98.
- [Bel57] R. Bellman, *A Markovian decision process*, Journal of mathematics and mechanics **6** (1957), no. 5, 679–684.
- [Bel66] R. Bellman, *Dynamic programming*, Science **153** (1966), no. 3731, 34–37.
- [Ber97] D.P. Bertsekas, *Nonlinear programming*, Journal of the Operational Research Society **48** (1997), no. 3, 334–334.
- [BR19] J. Bhandari and D. Russo, *Global optimality guarantees for policy gradient methods*, Preprint, arXiv:1906.01786 (2019).
- [BRT21] P. Breiding, K. Rose, and S. Timme, *Certifying zeros of polynomial systems using interval arithmetic*, 2021, Preprint, arXiv:2011.05000.
- [BSS06] M.S. Bazaraa, H.D. Sherali, and C.M. Shetty, *Nonlinear programming*, third ed., Wiley-Interscience [John Wiley & Sons], Hoboken, NJ, 2006, Theory and algorithms. MR 2218478
- [BT18] P. Breiding and S. Timme, *HomotopyContinuation.jl: A package for homotopy continuation in julia*, International Congress on Mathematical Software, Springer, 2018, pp. 458–465.

- [Che68] H. Chernoff, *Optimal stochastic control*, Sankhyā: The Indian Journal of Statistics, Series A (1968), 221–252.
- [DA21] J. Dahl and E. D. Andersen, *A primal-dual interior-point algorithm for nonsymmetric exponential-cone optimization*, Mathematical Programming (2021).
- [Der70] C. Derman, *Finite state Markovian decision processes*, Tech. report, 1970.
- [DTLR⁺19] R. Dadashi, A.A. Taiga, N. Le Roux, D. Schuurmans, and M.G. Bellemare, *The value function polytope in reinforcement learning*, International Conference on Machine Learning, PMLR, 2019, pp. 1486–1495.
- [HLL09] D. Henrion, J.B. Lasserre, and J. Löfberg, *Gloptipoly 3: moments, optimization and semidefinite programming*, Optimization Methods & Software **24** (2009), no. 4-5, 761–779.
- [How60] R.A. Howard, *Dynamic programming and Markov processes*, John Wiley, 1960.
- [JKM04] C. Jones, E.C. Kerrigan, and J. Maciejowski, *Equality set projection: A new algorithm for the projection of polytopes in halfspace representation*, 45.
- [Kal94] L.C.M. Kallenberg, *Survey of linear programming for standard and nonstandard Markovian control problems. Part I: Theory*, Zeitschrift für Operations Research **40** (1994), no. 1, 1–42.
- [KT51] H.W. Kuhn and A.W. Tucker, *Nonlinear programming*, Second Berkeley Symposium on Mathematical Statistics and Probability, 1951, pp. 481–492.
- [Las01] J.B. Lasserre, *Global optimization with polynomials and the problem of moments*, SIAM J. Optim. **11** (2000/01), no. 3, 796–817.
- [LHPT08] J.B. Lasserre, D. Henrion, C. Prieur, and E. Trélat, *Nonlinear optimal control via occupation measures and lmi-relaxations*, SIAM journal on control and optimization **47** (2008), no. 4, 1643–1666.
- [MGZA15] G. Montúfar, K. Ghazi-Zahedi, and N. Ay, *Geometry and determinism of optimal stationary control in partially observable Markov decision processes*, 2015, Preprint, arXiv:1503.07206.
- [MM22a] J. Müller and G. Montúfar, *The geometry of memoryless stochastic policy optimization in infinite-horizon POMDPs*, International Conference on Learning Representations, 2022.
- [MM22b] J. Müller and G. Montúfar, *Solving infinite-horizon POMDPs with memoryless stochastic policies in state-action space*, 5th Multi-disciplinary Conference on Reinforcement Learning and Decision Making, 2022.
- [MR17] G. Montúfar and J. Rauh, *Geometry of policy improvement*, International Conference on Geometric Science of Information, Springer, 2017, pp. 282–290.
- [MRA19] G. Montúfar, J. Rauh, and N. Ay, *Task-agnostic constraining in average reward POMDPs*, ICLR 2019 Workshop on Task-Agnostic Reinforcement Learning, 2019.
- [Ney03] A. Neyman, *Real algebraic tools in stochastic games*, Stochastic games and applications, Springer, 2003, pp. 57–75.
- [NR09] J. Nie and K. Ranestad, *Algebraic degree of polynomial optimization*, SIAM J. Optim. **20** (2009), 485–502.
- [NT21] J. Nie and X. Tang, *Convex generalized nash equilibrium problems and polynomial optimization*, Mathematical Programming (2021), 1–34.
- [PLT11] P. Poupart, T. Lang, and M. Toussaint, *Analyzing and escaping local optima in planning as inference for partially observable domains*, Joint European Conference on Machine Learning and Knowledge Discovery in Databases, Springer, 2011, pp. 613–628.
- [PS22] I. Portakal and B. Sturmfels, *Geometry of dependency equilibria*, 2022, Preprint, arXiv:2201.05506.
- [Put14] M.L. Puterman, *Markov decision processes: discrete stochastic dynamic programming*, John Wiley & Sons, 2014.
- [SMSM99] Richard S Sutton, David McAllester, Satinder Singh, and Yishay Mansour, *Policy gradient methods for reinforcement learning with function approximation*, Advances in Neural Information Processing Systems, vol. 12, MIT Press, 1999.
- [VLB12] N. Vlassis, M.L. Littman, and D. Barber, *On the computational complexity of stochastic controller optimization in POMDPs*, ACM Transactions on Computation Theory (TOCT) **4** (2012), no. 4, 1–8.

- [WB06] A. Wächter and L.T. Biegler, *On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming*, Mathematical programming **106** (2006), no. 1, 25–57.
- [WDL22] Y. Wu and J.A. De Loera, *Geometric Policy Iteration for Markov Decision Processes*, Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (New York, NY, USA), KDD '22, Association for Computing Machinery, 2022, p. 2070–2078.
- [Whi88] D.J. White, *Further real applications of Markov decision processes*, Interfaces **18** (1988), no. 5, 55–61.
- [WKZ⁺22] Kaixin Wang, Navdeep Kumar, Kuangqi Zhou, Bryan Hooi, Jiashi Feng, and Shie Mannor, *The geometry of robust value functions*, Proceedings of the 39th International Conference on Machine Learning, Proceedings of Machine Learning Research, vol. 162, PMLR, 17–23 Jul 2022, pp. 22727–22751.

MAREIKE DRESSLER, SCHOOL OF MATHEMATICS AND STATISTICS, UNIVERSITY OF NEW SOUTH WALES, SYDNEY, NSW 2052, AUSTRALIA.

MARINA GARROTE-LÓPEZ, DEPARTMENT OF MATHEMATICS, UNIVERSITY OF BRITISH COLUMBIA, VANCOUVER, V6T 1Z2 BC, CANADA.

GUIDO MONTÚFAR, DEPARTMENTS OF MATHEMATICS AND STATISTICS, UNIVERSITY OF CALIFORNIA, LOS ANGELES, 90095 CA, USA.

GUIDO MONTÚFAR, JOHANNES MÜLLER, KEMAL ROSE, MAX PLANCK INSTITUTE FOR MATHEMATICS IN THE SCIENCES, 04103 LEIPZIG, GERMANY